

https://dspace.org/

Application of stochastic programming techniques to airline scheduling

Silverwood, Mark William

2012-02-01
http://hdl.handle.net/10539/11226

Downloaded from a DSpace repository. https://dspace.org/

APPLICATION OF STOCHASTIC PROGRAMMING TECHNIQUES TO AIRLINE SCHEDULING

Mark William Silverwood

A research report submitted to the Faculty of Engineering and the Built Environment, University of the Witwatersrand, Johannesburg, in partial fulfillment of the requirements for the degree of Master of Science in Engineering.

Johannesburg, 2011

DECLARATION

I declare that this research project is my own unaided work. It is being submitted to the Degree of Master of Science to the University of the Witwatersrand, Johannesburg. It has not been submitted before for any degree or examination to any other University.

.....

(Signature of Candidate)

..... day of ..... year .....

ABSTRACT

The goal of this project was to evaluate the effectiveness of stochastic programming techniques when applied to the airline scheduling process to reduce the effect of stochastic flight delays. A variety of traditional and stochastic programming models were developed for generating flight schedules. The resultant flight schedules were tested using simulations to evaluate their performance in real-world conditions with regard to flight delays, and their effects on the schedule's operations. It was found that stochastic programming techniques were able to improve the delay recovery performance of the schedules at the cost of decreasing the schedule's profit; and that flight schedules which are more dense with flight activity are affected more by the stochastic programming techniques. The use of stochastic programming techniques is recommended for the cases where an airline's flight schedule has a high density of activity and the negative effects of flight delays needs to be minimized.

ACKNOWLEDGEMENTS

I would like to acknowledge the assistance and guidance of the project's supervisor, Mr Ian Campbell.

CONTENTS **PAGE**

DECLARATION		**2**
ABSTRACT		**3**
ACKNOWLEDGEMENTS		**4**
LIST OF FIGURES		**7**
LIST OF TABLES		**9**
NOMENCLATURE		**11**
1	**INTRODUCTION**	**12**
1.1	**Motivation**	**12**
1.2	**Literature Review**	**13**
1.2.1	**Airline operations**	**14**
1.2.2	**Airline scheduling**	**16**
1.2.3	**Stochastic programming**	**24**
1.2.4	**Stochastic airline scheduling**	**27**
1.3	**Objectives**	**29**
2	**METHODOLOGY**	**31**
3	**STOCHASTIC FLIGHT DELAYS**	**35**
3.1	**Overview Of On Time Performance**	**35**
3.2	**Departure Delays**	**37**
3.3	**Arrival Delays**	**40**
4	**OPTIMIZATION MODELS**	**44**
4.1	**Basic Fleet Assignment Model**	**44**
4.2	**Fleet Assignment With Variable Time Windows Model**	**48**
4.3	**Fleet Assignment With Time Windows And Route Selection Model**	**50**
4.3.1	**Schedule formation from model solution**	**52**
4.4	**Stochastic Models**	**54**
4.4.1	**Expected value (EV) model**	**55**
4.4.2	**Scenario generation with chance constrained programming**	**56**

5 MONTE CARLO SIMULATION 62
5.1 Implementation Of Monte Carlo Simulation 62

6 TEST SITUATION DATA SETS 67
6.1 Base Situation 67
6.2 Additional Test Situations 70

7 STOCHASTIC MODEL PARAMETERS 72
7.1 Model Assumptions 73

8 OBSERVATIONS 75
8.1 Test Equipment Specifications 75
8.2 Observations 76
8.2.1 Gurobi solver solution process 76
8.2.2 Schedule generation 77
8.2.3 Monte Carlo Simulation 78

9 RESULTS AND DISCUSSION 80

10 SUMMARY CONCLUSIONS AND RECOMMENDATIONS 90
10.1 Summary Of Work Completed 90
10.2 Conclusions 90
10.3 Recommendations 90
10.3.1 Recommendations for further research 91

REFERENCES 92

APPENDIX A COMPLETE OBSERVATIONS 97
APPENDIX B ZIMPL MODEL INDEX FILE 105
APPENDIX C GENERATED FLIGHT SCHEDULE FOR BASE PRM-1 MODEL 108
APPENDIX D REAL-WORLD TEST CASE (JETBLUE AIRWAYS) 109

LIST OF FIGURES

Figure	Page
1-1: Delay Departures Against Number Of Flights For The Years Of 2001 To 2010	16
1-2: Taxonomy of SP problems	25
3-1: Percentage Breakdown Of Total Delay Time Per Attributed Flight Delay Reason	37
3-2: Histogram Of Delayed Departures	38
3-3: Probability Plot Of Lognormal Distribution Versus Delayed Departures Distribution	39
3-4: Histogram Of All Departure Times	40
3-5: Histogram Of Difference Between Scheduled And Actual Arrival Times	41
3-6: Histogram Of Adjusted Arrival Distribution	42
3-7: Histogram Of Additional Arrival Delays Distribution	42
4-1: Flow Network Example	45
4-2: Ground And Wrap-Around Arcs	46
4-3: Basic Fleet Assignment Model	47
4-4: Fleet Assignment With Variable Time Windows Model	49
4-5: Fleet Assignment With Time Windows And Route Selection model	52
4-6: FARS Model Schedule Formation Step 1	54
4-7: Example Schedule Represented On A Time-Space Network	54
4-8: CCP Visual Example	58
4-9: Scenario Generation With Chance Constrained Programming Model	59
5-1: Monte Carlo Simulation Flow Diagram	63

5-2: Example Cumulative Distribution Function 64

6-1: Flight Network For Test Situations 69

9-1: Scatter Plot of Delay Recovery Versus Schedule Density 81

9-2: Scatter Plot of Average Propagated Delay Versus Delay Recovery 81

9-3: Scatter Plots of Average Delay Time Versus Schedule Density 82

9-4: Plots of Schedule Density Versus Delay Time Setting For Each Test Situation 83

9-5: Plots of Delay Recovery Versus Delay Time Setting For Each Test Situation 84

9-6: Plots of Schedule Profit Versus Delay Time Setting For Each Test Situation 85

9-7: Bar Chart Of Average Profit Per Delay Time Setting 86

9-8: Bar Chart Of Average Percentage Of Propagated Delayed Flights Per Delay Time Setting 87

9-9: Bar Chart Of Average Delay Recovery Percentage Per Delay Time Setting 87

9-10: Stacked Bar Chart Of The Relative Percentages Of Profit And Delay Recovery For Each Delay Time Setting 89

9-11: Bar Chart of Compromised Performance Rating For Each Delay Time Setting 89

LIST OF TABLES

Table	Page
3-1: Reasons For Flight Delays	36
3-2: Descriptive Statistics of Departures Distribution (minutes)	40
3-3: Descriptive Statistics Of Arrival Distributions (minutes)	43
4-1: Notation For The Basic Fleet Assignment Model	47
4-2: Notation For Fleet Assignment Model With Variable Time Windows Model	49
4-3: Additional Notation For Fleet Assignment With Time Windows And Route Selection Model	51
4-4: Additional Notation For The Scenario Generation With Chance Constrained Programming Model	59
5-1: Description Of Monte Carlo Simulation Table	65
6-1: List Of Airports Served	67
6-2: Passenger Demand Table	67
6-3: Aircraft Fleet Specifications	69
6-4: Additional Parameters For The Base Situation	70
6-5: Additional Test Situations	71
7-1: Delay Time Parameter Settings	73
8-1: Computer Specifications	76
8-2: Sample Solver Observations	77
8-3: Schedule Observations	78
8-4: Monte Carlo Simulation Observations	79
9-1: Relative Percentage Performance Measures	88
A-1: Complete Solver Observations	97

A-2: Complete Schedule Observations 99

A-3: Complete Monte Carlo Simulation Observations 102

D-1: JetBlue Model Observations 109

NOMENCLATURE

F	set of flights
K	set of fleets
S_k	number of aircraft of fleet k
G_k	set of ground arcs in fleet k's network
L_k	set of nodes in fleet k's network
c_{ik}	cost to fly flight i with fleet type k
x_{ik}	1 if flight i flown by fleet k, 0 otherwise
y_{gk}	number of aircraft on ground arc g in fleet k's network
$b1_{lik}$	1 if flight i begins at node l in fleet k's network -1 if flight i ends at node l in fleet k's network, 0 otherwise
$b2_{lgk}$	1 if ground arc g begins at node l in fleet k's network -1 if ground arc g ends at node l in fleet k's network, 0 otherwise
$d1_{ik}$	1 if flight i crosses the count time in fleet k's network, 0 otherwise
$d2_{gk}$	1 if ground arc g crosses the count time in fleet k's network, 0 otherwise
N_{ik}	number of arc copies of flight i in fleet k's network
c_{nik}	cost to fly copy n of flight i with fleet type k
x_{nik}	1 if copy n of flight i flown by fleet k, 0 otherwise
$b1_{lnik}$	1 if copy n of flight i begins at node l in fleet k's network -1 if copy n of flight i ends at node l in fleet k's network, 0 otherwise
$d1_{nik}$	1 if copy n of flight i crosses the count time in fleet k's network, 0 otherwise
p_i	penalty variable for flight i
q_i	seat price for flight i
dm_i	passenger demand for flight i
z_{nik}	profit of copy n of flight i with fleet type k
h_k	aircraft seat capacity of fleet type k
w	Number of Scenarios Generated
v_r	1 if flight duplicate r is flown, 0 otherwise
R_{nik}	Set of duplicate scenario flight arcs for copy n of flight i flown by fleet type k
Δ	Set of scenario levels
δ	Scenario level
α	CCP constraint tolerance
$j1_{r\delta l}$	1 if duplicate flight r begins at node l in scenario δ's network -1 if duplicate flight r ends at node l in scenario δ's network, 0 otherwise
$j2_{r\delta gl}$	1 if ground arc g begins at node l in scenario δ's network -1 if ground arc g ends at node l in scenario δ's network, 0 otherwise

1 INTRODUCTION

The structure of this report is as follows. Section 1 offers an introduction to the topics relevant to the project, presents the objectives of the project and gives a motivation for the project. Section 2 describes the methodology used during the completion of the project. Section 3 discusses the properties of flight delays and their distributions. Section 4 presents the various operations research models considered and developed for use during the project; both traditional mixed-integer network flow models and stochastic programming models are present. Section 5 describes the Monte Carlo Simulation used during the project to evaluate the operational performance of the schedules generated. Section 6 presents the variety of test situations that were generated for the purpose of running and evaluating the models used for the project. Section 7 presents the variety of parameter settings which were used for the stochastic programming models. Section 8 presents the observations of the project, namely the solver performance measures, schedule characteristics and the Monte Carlo Simulation outcomes. Section 9 presents and discusses the results obtained from analyzing the observations. Section 10 concludes the report with a summary of the work undertaken, the conclusions of the project, and recommendations.

1.1 Motivation

The creation of the flight schedules is a key part of an airline's operations and an interesting problem for the application of operations research techniques. The objective in mind when generating a flight schedule is to assign the most cost effective (or profitable) schedule whilst considering the operational constraints and issues, such as: maintenance schedules, fluctuating passenger demand, strict legislation, and stochastic events such as delays. The airline scheduling problem is a hard combinatorial problem to solve given the complexity of the constraints that have to be satisfied and the huge solution search space that has to be explored when searching for the optimal solution. Much research has been conducted in the field of airline scheduling optimization over the past few decades and the application of the techniques developed have yielded significant improvements. However, there is still room for improvement and a technique which has started to gain popularity in the airline scheduling field is the application of stochastic programming techniques which are suited to the handling of unpredictable variables which follow stochastic behavior. The majority of work conducted on stochastic programming within the airline scheduling field

has been focused on the sub-problem of crew scheduling. This project will attempt to evaluate the applicability of stochastic programming for the route scheduling and fleet assignment sub-problems of the airline scheduling process.

1.2 Literature Review

Carrier airlines operate in a highly uncertain environment yet very few airline planning and scheduling models consider this uncertainty (Schaefer et al. 2005; Rosenberger et al. 2002). The planning of airline operations is a long-term decision that has to be taken at least three to six months before the schedule is in operation. These planning decisions need to be made before all the relevant data is available and thus forecasts and intuition are used to predict the future operating characteristics such that a decision can be made. These predictions are inherently uncertain and as a result, there is often a notable discrepancy between an airline's planned and actual performance (Rosenberger et al. 2002). Uncertain factors include forecasted data such as passenger demand and unforeseen disruptions such as flight delays. Incorrectly forecasted passenger demand can lead to massive losses for the airlines because of either empty flights or lost customers on flights which are already fully booked. Unforeseen disruptions on the day of operation can result in the flight schedules becoming infeasible (Stojkovic 2002), thus requiring that they be updated. This rescheduling is costly and results in a loss of traveler goodwill (Stojkovic 2002) because it can lead to delayed or cancelled flights, or swapping aircraft among flights, which in turn can affect future deployment of aircraft and crews (Stojkovic 2002).

Traditional methods model airline scheduling problems using deterministic operating conditions (Yen & Birge 2006; Burke et al. 2010), which discards the potential disruptions and uncertainty of the predicted values used in the model. In addition, Burke et al. (2010) argues that 'current models have a strong focus on minimising the operating cost' and 'improving the utilisation of resources', this often has the effect of making the schedule less robust to disruptions and thus highly susceptible to flight delays and cancellations.

Stochastic programming is an approach for modeling optimization problems that involve uncertainty. Stochastic programming allows the use of parameters which are unknown at the time a decision has to be made, which is invariably the case with most real world problems. Airline scheduling operations are a good candidate for the application of stochastic programming due to the inherent

uncertainty of the parameters and long-term decision process. The application of stochastic programming in the airline scheduling sector offers various potential advantages over the deterministic approach, including: creation of more robust schedules which results in fewer delays, reduced costs of actual airline performance and more feasible schedules.

The remainder of this literature review will focus on: the current and forecasted state of the airline industry in general, predominate scheduling techniques used by airlines, stochastic programming techniques, and the current application of stochastic programming in the airline industry.

1.2.1 Airline operations

The airline industry operates in a highly competitive market; competition is ever increasing as airlines attempt to secure and retain customers. This increased competition calls for more streamlined operations that reduce costs as price is used as the main competitive factor for winning customers; however, there is little difference between the ticket prices of competing airlines and all airlines have matching frequent flyer and discount programs (now considered order qualifiers). As the market nears, or is already operating at, perfect competition airline operators' profits are dropping (especially as fuel prices continue to rise and stock markets are failing to recover quickly from the recent financial crisis). Additionally, Jones and Sasser (1995) states that airlines are relatively efficient in responding to competitors' price changes. This leads to the airline operators needing to reduce costs through the better use of their resources and by improving scheduling decisions (Yen & Birge 2006) to remain competitive, and differentiate their service from the competitors such that their market share increases or they can charge a premium for their services.

One such competitive dimension for airliners to target, for obtaining a competitive advantage, is service quality. Ostrowski et al. (1993) showed that continuing to provide perceived high quality services would help airlines acquire and retain customer loyalty. Empirical studies of demand for airline services conducted by the *Bureau of Transport and Communication Economics* (BTCE) (1994), showed that service quality is central to the choice of airlines for both business and leisure travelers.

Parasuraman et al. (1988) identified 10 potentially overlapping dimensions of service quality; which were: tangibles, reliability, responsiveness, communication, credibility, security, competence, courtesy, understanding or knowing the customer, and access. Reliability and responsiveness are service quality dimensions which are directly affected by the selection of the flight schedule because flight delays and cancellations fall under those dimensions. This makes the construction of more robust schedules, which are less likely to contain flights to be delayed or cancelled, 'an absolute necessity for airlines' (Burke et al. 2010). Stochastic programming techniques are a candidate for improving the robustness of airline schedules.

Rosenberger et al. (2002) suggests that 'major airline carriers almost never experience a day without disruptions'. The need for robust schedules would appear to be increasing as the data suggests that on-time performance is deteriorating (Schaefer et al. 2005). Rosenberger et al. (2002) cites Dobbyn (2000) whom states that 'average daily flight delays increased 20% from 1998 to 1999'. Consequently customer complaints also increased by 130%. Delay analysis by Eurocontrol (2005), showed that the European airline industry was suffering from an indisputable increase in delays. In 2005, as much as 20% of flights were delayed by 15 minutes or more.

Using data provided by the Bureau of Transportation Statistics (BTS) (2010) for the period between 2002 and 2010, it was found that on average 17% of flights were delayed (with an increase of 4% between 2002 and 2010) and 1,6% of flights were cancelled (with an increase of 0,4% over the period). Rosenberger et al. (2002) offer the following causes of the delays: air traffic controllers are struggling to accommodate the number of flights wanted by the airlines, mechanical failures, and in particular weather all account for three-quarters of the delays. Burke et al. (2010) concluded, after investigating the causes of the delays, that 50% are due to airline related issues, whereas 19% are due to airport operations.

Schaefer et al. (2001) states that for each 1% increase in air traffic, an expected 5% increase in flight delays will occur. Burke et al. (2010) and Lan et al. (2006) concur that increased air travel will result in the increased congestion of airports and airspace, which will 'likely' result in a further increase in the number of delays. Air traffic over the period between 2002 and 2009 increased by 22% in

Europe (number of passengers increased by 27%) (AEA 2010) whilst air traffic in America increased by 22% over the same period (BTS 2010). Analyzing the flight statistics provided by BTS for the period of 2001 to 2010 for American flight operations, the following observation that departure delays increases with number of flights can be made and is presented below in Figure 1-1.

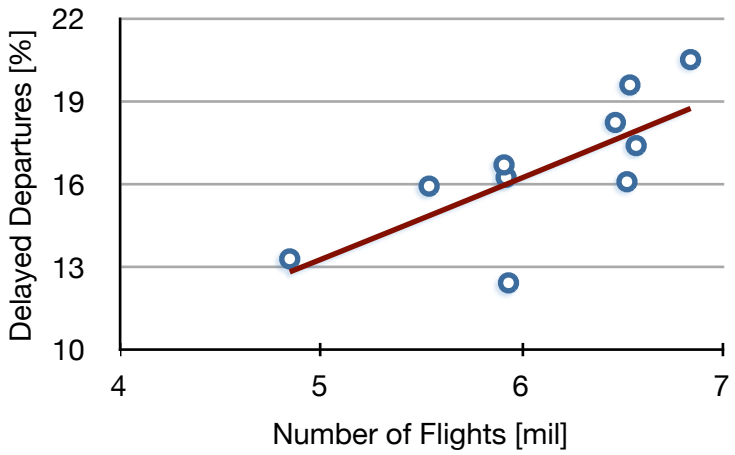


Figure 1-1: Delay Departures against Number of Flights for the years of 2001 to 2010

The growth trend of the airline industry is expected to continue, the International Air Transport Association (IATA) is forecasting an 18% increase in international and domestic air traffic for the period 2010 until 2013 (IATA Industry Fact Sheet December 2010). Boeing Commercial Airplanes current market outlook for the period 2010 to 2029 is forecasting a yearly growth rate of 4,2% for number of passengers and 5,3% for air traffic. Which equates to an increase of 118% for the number of passengers and 167% of air traffic for the period of 2010 to 2029, thus within the next two decades air traffic will more than double its current level. Boeing is also forecasting a 92% increase in the number of aircraft operating worldwide. These forecasts suggest that flight delays are to become an even more important issue in the future.

1.2.2 Airline scheduling

Aircraft seats are the airline's product. As with any product, 'a larger quantity secures sales, while extra inventory incurs costs' (Sherali et al. 2006). Larger quantities (capacity of aircraft) results in higher operating costs, additionally aircraft seats are 'perishable' (Sherali et al. 2006) which means that any unsold seats on a flight are lost or wasted. The ideal strategy for an airline is to provide

the exact number of seats to the passengers at exactly the right price, this is clearly a challenge in an environment that is not deterministic. The use of operations research techniques in the planning and scheduling of airline operations attempts to provide the ideal strategy using the information available.

Flight scheduling has been a primary focus of airlines because it has a great bearing on a 'carrier's profitability, its level of service and its competitiveness in the market' (Yan, Tang, Fu 2008). The creation and optimization of flight schedules in the operations research field generally consists of creating a mathematical model which describes the scheduling problem and considers important operational factors such as: passenger demand, ticket prices, operating costs, operating constraints (e.g. specifications and limitations of aircraft types, fleet size, airport constraints), aircraft maintenance requirements, legislation and crew scheduling constraints. Using optimization techniques, this model is then solved to satisfactory level with regard to the theoretical optimal solution. The solution of the operations research model is then converted into the flight schedule (containing crew and maintenance schedules).

The scheduling process for most airlines involves different departments, namely: the network department, the operational plan management department and operations control (Burke et al. 2010). The responsibility of the network department is to create the initial feasible flight schedule; the main goals of this department are the 'maximization of the market share, maximization of passenger revenue and minimization of operating costs' (Burke et al. 2010). The majority of the work detailed in this report is conducted within this department. Burke et al. (2010) mentions that at KLM, schedule construction takes approximately 2 months to complete and Rosenberger et al. (2002) states that typically this process takes place at least 3 to 12 months in advance of the scheduled flights and that it is largely driven by market considerations. Once the schedule is constructed by the network planning department then it is passed on to the operational plan management department (OPMD). The OPMD job is to make minor adjustments and changes to the schedule such that operational performance is increased and to anticipate changes in the market (Burke et al. 2010). The operations department receives the schedule approximately two weeks before the day of the scheduled flights; it is their duty to make last minute adjustments and manage the schedule on the day of operation (Burke et al. 2010). Included in the operations departments responsibilities is the

implementation of recovery processes or strategies to manage disruptions on the day of operation.

The construction of the flight schedule by the network department is generally composed of the following optimization sub-problems: route selection, fleet assignment (FA), aircraft maintenance routing (MR), and crew scheduling (CSP) (Papadakos 2009). The flight scheduling problem is divided into these various sub-problems because the full optimization problem is considered computationally intractable. This will be further discussed later. The subproblems are optimized sequentially such that the output of one sub-problem is the input to the following sub-problem; this is known as the sequential approach. Airlines initially commence the scheduling process with the tactical planning process of schedule generation, route selection, (Papadakos 2009) which produces a timetable of the most profitable flight legs (a *leg* is a flight from one specific origin to another specific destination at a given time) which the airline would wish to provide. After this, the next step is to solve the fleet assignment stage which decides which fleet (grouping of similar aircraft) should fly each scheduled leg while using the available aircraft and maximizing profits. The next stage is the maintenance routing stage, also known as the aircraft rotation stage (Rosenberger et al. 2002), which is solved such that each aircraft is periodically scheduled for maintenance in accordance with legislation. The output of the MR stage is an individual schedule for each aircraft in the fleet. A rotation generally takes many days to fly, however (particularly in europe and domestic flights) it is also common to have a rotation occurring within the same day (known as a daily route). Essentially the scheduling of the aircraft is now complete which leads to the final stage of scheduling which is crew pairing, or crew scheduling. The crew pairing stage devises the series of legs crew have to fly whilst respecting labour rules, minimizing crew costs (Papadakos 2009) and considering the fleets types that each pilot and crew are qualified to fly.

Fleet assignment

Lets now consider the scheduling steps in further detail beginning with the fleet assignment problem. Large airline carriers typically have more than one type of aircraft; a fleet is a set of planes of the same type. Sherali et al. (2006) offers the following description: ‘the fleet assignment problem (FAP) deals with assigning aircraft types, each having a different capacity, to the scheduled flights, based on equipment capabilities and availabilities, operational costs, and potential

revenues'. Gu et al. (1994) states that the driving force for this optimization is that different fleets produce different revenues when assigned to specific flight legs because of different seating capacities and operating costs; thus the fleeting decision highly impacts an airlines revenues. Assigning a smaller aircraft than needed for a particular flight will result in lost revenues because customers will be 'lost' whom can't be placed on the flight due to insufficient capacity; conversely, assigning a larger aircraft than required would result in unsold seats, and generally higher operating costs.

As stated above, most of the traditional approaches solve the FAP in isolation using the sequential method; additionally Sherali et al. (2006) notes that traditional approaches often operate under 'restrictive assumptions' such as: considering the same-every-day schedule and using point forecasts for flight demands. Fleet assignment problems are solved using a mixed integer multi-commodity model (Rosenberger et al. 2002) based on the airlines flight network (Sherali et al. 2006). Sherali et al. (2006) observed that there are two principal trends that are adopted in constructing the network for the model, namely connection networks and time-space networks. Connection networks use arcs (connection of nodes in the network) to represent connections whilst time-space networks use the arcs to represent flight legs. These two network constructs are similar in that they share the following main constraints: 1) balance constraints for the conversation of flow, 2) 'cover constraints' so that each leg is assigned to only a single fleet; and 3) aircraft availability constraints which limits the usage of each fleet depending on the amount of available aircraft.

Sherali et al. (2006) states that the time-space network has largely become the network of choice in solving the FAP. The time-space network construction focuses on representing flight legs and leaves the model to determine the connections to use whilst taking into account the constraints and the time and space considerations. Hane et al. (1995), Gu et al. (1994), Barnhart et al. (1998), Boland et al. (2000), Kniker (1998), and Sherali et al. (2006) describe the fleet assignment model in more detail.

Maintenance routing

Maintenance routing is strongly linked to fleet assignment and is the next step in the scheduling procedure. MR assigns each flight leg to a specific individual aircraft within the fleet such that costs are minimized and maintenance

requirements for each aircraft are meet. According to regulations, the maintenance period for each aircraft has to occur at least every three or four days and this maintenance period typically lasts for around eight hours. The critical requirement for the MR is that the resultant flight schedule is feasible for operation.

Crew scheduling

Crew scheduling (CSP) is the allocation of pilots and crew to specific aircrafts for specific flight legs and is not necessarily the same as the aircraft assignments in MR. Crew scheduling is subject to many constraints such as regulations, airline specific rules and labour agreements (Burke et al. 2010). There are two components to the CSP: long-term planning and short-term planning. Yen et al. (2006) notes that traditionally these have been seen as two separate problems. Short-term planning involves making the crew assignments under short-term time constraints and is usually done at the operational level (Stojkovic 1995). Long-term planning is the traditional approach where crew schedulers create a set of optimal crew itineraries to be later assigned to specific crews. Itineraries are usually selected by the crews in order of seniority. Schaefer et al. (2005) states that the CSP can be very difficult to solve because of the many 'governmental and contractual regulations' regarding pilots and that problems found in practice often have more than a billion possible solutions due to the enormous number of pairings.

Once the crew scheduling has been completed then in essence the scheduling process is complete. However, in reality disruptions occur that result in portions of the schedule (generally the current day of operation) becoming infeasible and thus quick decisions have to be made to adjust the immediate schedule such that delays and cancellations of flights can be reduced. Recovery is the process of reacting to a disruption (Rosenberger et al. 2002); the optimal recovery decision is hard to determine and is rarely clear. Lettovsky (1997) reports that irregular operations can be responsible for as much as 3% of an airline's operating expenses. Different recovery policies will give different results, however Lettovsky (1997) states that in practice, airlines make recovery decisions manually with little decision support. A recent trend in research in the airline operations research sector has been the investigation of optimizing this process and the generation of more robust schedules that are less susceptible to disruptions.

Decomposed sequentially approach

Returning to the issue of deterministic and decomposed models for the airline scheduling process. The reason why the airline scheduling problem is decomposed in the sequential method is that the complete integrated problem is currently computationally intractable (Papadakos 2009; Burke et al. 2010; Sherali et al. 2006). The main drawbacks of solving the scheduling problem using the sequential approach are that an optimal solution for one stage of the process may not necessarily be the optimal solution for the entire system and that the solution from one stage may not be a feasible input for the subsequent stages (Papadakos 2009; Burke et al. 2010; Sherali et al. 2006). Thus this process results in sub-optimal results for the entire system. Burke et al. (2010) states 'applying deterministic and decomposed approaches is know to result in sub-optimal schedules with many tightly interconnected resources, leaving less flexibility in the schedule to recover from delays'.

The interdependence of the various stages of the decomposed approach has motivated researchers to focus on the area of integrated models. Integrated models 'simultaneously consider several of these processes (stages) so as to achieve a better solution for the entire system' (Sherali et al. 2006). Currently no attempts have been made to integrate the entire flight scheduling system due to the computational effort of the resulting model. The current trend are the so called semi-integrated models, which integrate two or more sub-problems but not the entire system. The aim of these semi-integrated models was to achieve better quality results and lay the framework for further research into the integration of the entire system. Papadakos (2009) and Sherali et al. (2006) state that significant savings have been achieved using semi-integrated models and the techniques has generated optimal or near-optimal solutions resulting in higher revenues for the airlines. Papadakos (2009) states that the best performing model available from the literature is a combination of semi-integrated models executed in a sequential approach. Papadakos (2009) offers this description of the method: 'This method consists of initially solving the integrated FA with MR problem and feeding the acquired solution of each fleet into the integrated MR with CP problem'.

The obvious question is: if the fully integrated method offers the best results then why is it not used in practice? The answer is that the advantage of the integrated method (improved solutions) is counted and outweighed by its disadvantage of

slower runtimes and complexity. As previously stated, the full airline scheduling problem is considered computationally intractable (Papadakos 2009; Sherali et al. 2006; Burke et al. 2010; Schaefer et al. 2005) thus the problem is solved in the sequentially manner so as to reduce the computational complexity. Even with a network of modern computers, the complete problem could take months to be solved and thus its current application isn't feasible. It is worth noting that even each step of the sequentially method is a computationally difficult problem and that takes a significant amount of time to solve. Gu et al. (1994) has shown that the fleet assignment problem, even without the availability constraints, is NP-hard (non-deterministic polynomial-time hard) for three fleet types. Mirabi et al. (2010) offer the following explanation of the properties of NP-hard problems: it 'means that an efficient algorithm to solve the problem to optimality is unavailable. Therefore, solving the problem by an exact algorithm is time consuming and computationally intractable'.

Reducing the computational time of airline scheduling models is a clear target of research in the area. The reduction of computational time will allow the formulation of more complex models that result in improved solutions. Hane et al. (1995) proposed a series of preprocessing steps that aimed to reduce the size of the network in the FA problem and thus it's computational difficulty. These steps have become a standard practice (Sherali et al. 2006). Hane et al. (1995) reported that the size of a typical problem was reduced from 48 982 rows and 66 942 columns to 7703 rows and 20 464 columns. Similarly, advances in the techniques used to solve the models have resulted in successful gains in the area. In general, airline scheduling models have a poor LP relaxation in the sense that the gap between the optimal integer programing (IP) and linear programming (LP) objectives can be very large. This optimality gap is used to estimate the optimality of the current solution and thus whether the current solution is optimal; it determines when the model is considered 'solved'. The large gap between the IP and LP objectives results in longer solving periods. A typical solution algorithm solves the branch-and-bound integer program by searching depth-first and terminating upon the first encountered solution (Papadakos 2009). This heuristic method is common in the airline industry because experiments have shown that the initial solutions obtained within a short period of solving the model are often very close to the final optimal solution obtained after the model has been completely solved (after a long period) (Papadakos 2009). Mirabi et al. (2010) adds that "all previous researchers

preferred heuristic methods to exact algorithms" to tackle the complexity of the problems (and the resulting time consuming solving). Hybridizations of genetic algorithms with multiple local search operators - as known as multi-meme memetic algorithms - are becoming increasing popular solution algorithms for the airline scheduling problems (Burke et al. 2010). Genetic algorithms are valued for their ability to locate promising and diverse regions of the solution search space quickly (Sastry et al. 2005), however they often have difficulty locating the local optimal solutions for that region. Local search methods are coupled with the genetic algorithm to overcome this problem.

Airline scheduling is usually content with near-optimal solutions in practice and solution algorithms are usually terminated when the optimality gap falls within a user-defined constant (Mercier et al. 2005; McDaniel & Devine 1977). This leads us to the topic of what is considered a "good" schedule in the airline industry. Current models focus on 'minimising the operating cost, improving the utilisation of resources and decomposing the schedule into several independent sub-problems' (Burke et al. 2010) however, most models discard the influence of potential disruptions on the day of operation and use deterministic operating values in the model. The increase in airline delays has shifted the focus of airline schedulers from the goal of simply maximizing the profit to maximizing the real world performance of the schedule. Schaefer et al. (2005) adds that the quality of the schedules is not measured by its performance in operations but rather by its performance on paper assuming that everything goes to plan (which it rarely does, as discussed earlier). This represents a clear flaw in the scheduling process and it has been proposed that the evaluation of schedules should consider these factors.

Another major flaw in the current scheduling models is the estimation of the model parameters (such as demand, fares, flight times, maintenance times, etc.) into deterministic values. The resultant accuracy of the solution is fully dependent on these values. Sherali et al. (2006) states that in the current airline market, 'demands change drastically'. This coupled with the fact that schedules are constructed on a long-term process (typically more than 3 months ahead of the scheduled flights) calls for models which consider the distributions of these values more accurately and result in more robust schedules. Stochastic programming is a potential method to achieve these goals. Sherali et al. (2006) further adds that 'robust models are imperative so as to obtain solutions that are

good in not only the ideal or expected instances, but even in the unexpected cases'.

1.2.3 Stochastic programming

Optimization under the assumption that the input data is not completely available at the decision time has, in recent years, received increasing attention (Grothklags & Lorenz 2006). Grothklags & Lorenz (2006) conjecture that the world is becoming more dynamic and that data from the past, and forecasts, is having a decreased predictive power for planning and deciding. Stochastic programming merges the models of operations research and the models of statistical randomness (probabilities, distributions) to form a robust decision-making tool.

'Stochastic programming is an approach for modeling optimization problems that involve uncertainty' (Shapiro & Philpott 2007). Classical deterministic optimization problems are formulated using known values for the model parameters, however real world problems are far more likely to contain parameters which are unknown at the time a decision needs to be made. Stochastic programming models try to take advantage of the fact that the probability distributions for the data are known or can be estimated. In essence, stochastic programming is the replacement of deterministic values in an optimization problem with random variables or probability distributions describing the true nature of the parameter; thus it allows management decisions which are usually made in uncertain environments to be considered more accurately with fewer assumptions. Benisch et al. (2004) add that in most optimization applications stochastic information about the shape of the future is readily available in the form of probabilistic models built from historical data.

Stochastic programming has become a more viable decision-making tool in recent times due to the advances of computer hardware, software techniques and solution methods (Domenica et al. 2007; Benisch et al. 2004). Benisch et al. (2004) state that 'there are considerable advantages to taking account of stochastic information'.

Traditional operations research techniques have approached the issue of uncertainty and its effects on the solution through the use of sensitivity analysis. However, it has been shown by Higle & Wallace (2002) that this approach had a

number of limitations and may provide misleading conclusions in respect of the nature of the solution. Domenica et al. (2007) adds that, in general, sensitivity analysis is not a suitable approach for the understanding of uncertainty and random variables within the model parameters.

There are a variety of stochastic programming techniques that have been developed to suit different applications and purposes. The first method for dealing with stochastic parameters in stochastic programming is the *expected value model*, which optimizes the expected objective function subjected to some expected constraints (Baoding & Liu 1997). The second method called the *sample average approximation* (SAA) was developed by Shapiro (2002). Third, *chance constrained programming* (CCP) was developed by Charnes and Cooper (1959) which operates by specifying a confidence level at which the desired stochastic constraints should be held. The next method is called *dependent-chance programming* (DCP) pioneered by Baoding and Liu (1997) which tackles the problem that occurs when 'a decision system undertakes multiple certain tasks called events, and the decision maker wishes to maximize the probabilities of satisfying these events' (Baoding & Liu 1997). Another approach to solving SP problems involves scenario-based analysis in which various different scenarios are constructed and solved in a mostly deterministic manner. The solutions for each scenario are analyzed together to determine the most appropriate solution. Domenica et al. (2007) offers a taxonomy of SP problems which is presented in Figure 1-2 below.

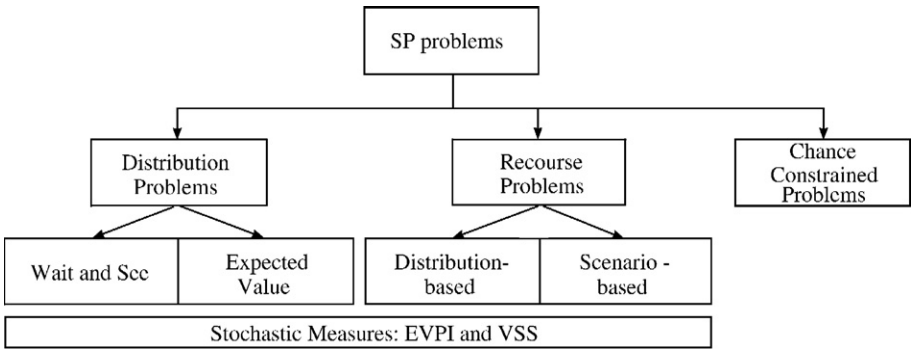


Fig. 4. Taxonomy of SP problems.

Figure 1-2: Taxonomy of SP problems (Domenica et al. 2007)

The 'most widely applied and studied' stochastic programming models are the two-stage linear programs (Shapiro & Philpott 2007). Two-stage models are

applicable for cases where a decision needs to be made before a random event occurs and after the event has occurred then adjustments to the decision are possible. The first stage decisions need to be taken before 'experiment' has taken place and the decision needs to consider all the possible outcomes of the experiment. Once the experiment has taken place and its effects have affected the outcome of the first-stages decision then a second decision (known as the recourse decision) can be taken that compensates for the outcome of the random event and determine the optimal response. The optimal policy from such a model is a single first-stage decision and a collection of recourse decisions, resulting in a decision rule, which define the second-stage action to be taken in response to each random outcome (Shapiro & Philpott 2007). Two-stage models are applicable for the operations department of an airline that needs to make quick decisions on the day of operation as disruptions affect the initial schedule.

The expected value model is constructed by replacing the stochastic random parameters with their expected values (e.g. mean values for distributions) such that it becomes a linear program. The uncertainty is dealt with before the selection of the parameters for the optimization model. The EV problem is often used to gain insight into the decision problem (Domenica et al. 2007). Unfortunately, the EV model largely ignores the stochastic properties of the information provided (Benisch et al. 2004) and is often inappropriate (Liu 2002). The sample average approximation (SAA) method has been shown to outperform the EV model (Benisch et al. 2004) and offer an approach which considers the stochastic information to a greater degree.

Chance-constraint problems incorporate the uncertainty of certain events into the model through the use of probability constraints. Bravo and Gonzalez (2009) state that in precise conditions some simple problems can be formulated using CCP but that the treatment of most problems by CCP becomes 'too cumbersome'. Man (2008) states that the CCP approach can lead to problems with non-convex constraint sets which are difficult to solve.

Scenario based analysis consists of constructing many diverse scenarios according to the stochastic data at hand which can be in the form of discrete distributions, or limited samples, or approximation methods, or by expert opinion (Petkov 1997). Man (2008) notes that scenario-based analysis approaches

provide a ‘relatively straightforward’ way to account for uncertainty, however they may rely on a very large number of scenarios.

The next section will discuss the application of stochastic programming to airline scheduling operations.

1.2.4 Stochastic airline scheduling

Floudas and Pardalos (2005) describe the field of stochastic scheduling as being motivated by the ‘design and operational problems arising in systems where scarce service resources must be allocated over time to jobs with random features vying for their attention’. The random features are modeled by specifying their probability distributions which are assumed to be known.

Traditional methods in the airline industry model the scheduling problem as a deterministic model and do not explicitly consider disruptions and variance in the operations (Yen & Birge 2006; Rosenberger et al. 2002) even though disruptions occur frequently. An airlines performance can be significantly different to the planned performance according to the schedule. Traditional models measure the performance of a schedule according to the assumption that every flight will take off and land as scheduled and experience no disruptions in operation (Rosenberger et al. 2002); however, this scenario very rarely takes place. Since the traditional methods neither consider stochastic variables in the model or during the evaluation of the schedule, the result is the formulation of schedules that are not robust and highly susceptible to disruptions in the actual operations.

Yan, Tang & Fu (2008) discuss the stochastic variations that occur for the market demand of airline flights; Yan et al. (2008) mention that market demand usually varies on a daily basis and that market share may vary with passenger choice behaviors. Market share may also vary according to the perceived convenience of the schedule for the passengers. Yan et al. (2008) conclude that passenger demand fluctuations arising from stochastic market conditions have to be taken into account when modeling the scheduling problem.

Aside from passenger demand fluctuations, the other key reason for the application of stochastic programming in airline scheduling is the desire to build more robust schedules that handle disruptions to the schedule in a better manner. Burke et al. (2010) offers the following definition for the robustness of a

schedule: ‘the robustness of a schedule is influenced by its sensitivity to stochastic events, the flexibility within the schedule and its stability’. Flexibility is related to the number of recovery options available in the schedule to deal with the effects of a disruption (can be defined according to the number of single point swap opportunities between aircraft in the schedule). Stability is the measure of the probability that a delay will propagate through the rest of the schedule and cause further delays and disruptions. Bian et al. (2005) formally presented an analytical approach for the evaluation of the robustness of airline schedules using features called *robustness objectives*. Burke et al. (2010) has proposed the use of Pareto optimization (Deb 2001; Landa-Silva et al. 2004; Deb 2005) for the improvement of robustness objectives in airline scheduling models. Burke et al. (2010) found that schedule reliability was dominant through the use of sensitivity analysis and that increased flexibility could result in improved performance; they recommend ‘schedule operators to primarily focus on schedule reliability, and take schedule flexibility into account while building the schedule’. The traditional approach of maximizing the profit results in schedules which maximize the amount of time which the aircraft is in the air (utilization). This results in a decrease in the amount of slack time between connecting flights in the schedule. Chiraphadhanakul (2010) defines slack as ‘additional time allocated beyond the minimum time required for each aircraft connection, passenger connection, or expected flight duration’. The reduction of slack time reduces the robustness of the schedule, since slack is desirable in robust schedules because it can potentially absorb delays (Chiraphadhanakul 2010). AhmadBeygi et al. (2008) proposed a method that reduces delay propagation by redistributing the existing slack in the flight schedule to the points where it is needed the most. Lan et al. (2006) developed a robust aircraft routing model to minimize the expected propagated delay along aircraft routes; they used an approximate delay distribution to model the delay propagation along each string. By adjusting flight times without changing aircraft routing, Wu [24] revealed that significant delay (cost) savings can be achieved via robust scheduling.

Burke et al. (2010) offers the following recommendations for the modeling of airline schedules in stochastic programs: 1) The stochastic nature of flight times and departure handling can be modeled using γ-distributions, 2) arrival handling can be modeled using deterministic distributions, and 3) Changing operating conditions (such as weather and congestion patterns) can be modeled by

defining different distributions depending on the time of the day, time of the year, the origin/destination of the flight and the scheduled times for the activities.

Schaefer et al. (2005) and Rosenberger et al. (2002) both suggest the use of Monte Carlo simulation for the evaluation of the operational performance of the schedules produced from the stochastic models. SimAir is a Monte Carlo simulation of airline operations that can be used for this purpose and a more detailed description of the underlying stochastic model can be found in Rosenberger et al. (2000b). Burke et al. (2010) made use of the KLM simulation model for the evaluation of the operational performance of his schedules.

It has been shown that the use of stochastic programming in airline scheduling results in improvements of the operating performance of the schedules constructed (Yan et al. 2008; Schaefer et al. 2005; Rosenberger et al. 2002; Burke et al. 2010). However, most research into this area has been focused on the crew scheduling problem and stochastic attempts in the remaining portions of the airline scheduling process have been reduced to simple problems (using such assumptions as single-fleet and non-stop flight operations) due to the computational difficulties present. It has been shown by Grothklags and Lorenz (2006) that the stochastic variant of the fleet assignment problem is PSPACE-complete and thus very difficult to solve. There is scope to research the creation of a hybrid model that attempts to incorporate the advantages of stochastic programming into the traditional mixed-integer scheduling models whilst keeping the computational requirements low; such research would most likely draw from the following areas of research: robustness objectives, simple stochastic solution algorithms (such as SAA and EVM), and the traditional scheduling models.

1.3 Objectives

1. Investigate the appropriateness of stochastic programming techniques for the airline scheduling problem.
2. Evaluate the performance of various operations research techniques in formulating and solving of the airline scheduling problem with respect to computational time, optimality of solution w.r.t. the objective value, and the performance of the solution during simulation.
3. Offer recommendations for the selection of an appropriate technique for the airline scheduling problem.

4. Identify avenues for the further improvement of the models presented and offer recommendations for further research.

2 METHODOLOGY

The goal of this research project is to evaluate the effectiveness of stochastic programming techniques when applied to the route scheduling and fleet assignment sub-problems of the airline scheduling process under the effect of stochastic flight delays. The distinct challenges of this project are: ensuring that the generated schedules are applicable and feasible in real-world conditions, and that the formulation and solution of the operations research model is computationally tractable and completes within a reasonable time frame.

The research conducted did not involve the participation of human participants as informants or subjects and all information was obtained from freely available public sources.

The approach taken in completing this project is summarized as follows: A literature review was conducted, then the relevant data required for the project was gather from various sources, the data was then analyzed to form the required information for the projects models, a small test model was developed upon which the larger and more complex stochastic programming models were built, these model's performance were evaluated using a selection of test situations, and, finally, the observations collected were used to form conclusions.

The methodology utilized during the completion of this project is as follows:

- *Conduct a literature review.* The first step of the project was to conduct a literature review such that the current state of research in the relevant fields could be established and utilized during the completion of the project. The main areas of focus for the literature review were: airline scheduling techniques, stochastic programming, and mixed-integer programming.
- *Locate and obtain the relevant data required for the project regarding flight delays and general information regarding the commercial airline industry*. Up-to-date information regarding flight delays and the airline industry were gathered from the bureau of transportation statistics in the United States of America. Detailed information regarding every commercial flight operated in the USA is freely available. Additionally, information regarding the various aircraft technical specifications and the composition of the airlines fleets were obtained from various additional sources.

- *Analysis of the obtained data*. The obtained raw data was analyzed such that the relevant information required for the project could be extracted. The key information required for the project included: departure delay distributions, arrival delay distributions, the causes and their relative contribution to the delays, airport locations (GPS) used for calculating distances, typical passenger demand for a selection of routes, aircraft passenger capacities, aircraft cruising speeds and range and aircraft operating costs.
- *Generate a small test situation*. A small test situation was generated for the use of developing and testing a small operations research model. The test situation contained all the necessary properties that would be relevant for the larger models to be developed at the later stage.
- *Generation of a small and simple naive operations research model*. A small and simple operations research model was formulated to solve the route scheduling and fleet assignment problem; this small model was used for testing and developing the model such that the model was stable, solved quickly and produced the desired schedules. Once the small model achieved the desired level of performance (subjective evaluation by the researcher) then it was used as a base upon which the more advanced models would be built.
- *Develop new stochastic programming models*. The next step was to build more complex models which incorporated stochastic programming techniques. This process is further explained in the mathematical models section (Section 4.4).
- *Generate a test situation*. For the purpose of running and evaluating the models, a test situation was created. The test situation was modeled on a local low-cost airline operator. Further information on this test situation is included in Section 6.1.
- *Generate additional test situations.* Additional test situations were generated by modifying certain key parameters of the base test situation. The key parameters adjusted were: number of fleets, number of aircraft and length of the flight operations day. The other parameters remained unchanged, the different situations generated simulate different degrees of airline capacity and as such the density of the resultant flight schedules.
- *Run and evaluate the various test situations with the various models developed*. The developed models are run on each of the various test situations created and their performance is recorded. The key performance

measures include: computational solving time, profit of schedule and the ability of schedule to absorb delays. Monte Carlo simulation is used to evaluate the operational performance of the schedules generated whilst the solving time and profit are returned with the solution from the solver.

- *Analyze the observations and drawn conclusions*. The observations made during the running and evaluating of the developed models using the various test situations is now analyzed such that results can be obtained and the results can be used to form conclusions.

A mathematica application was developed for the purpose of generating the models to be solved, generating the schedule from the solution obtained from the solver, evaluating the performance of the schedule using Monte Carlo simulation, and the general analysis of the raw data and results obtained. The mathematica program has the following features:

- Loading and processing of the raw data obtained from the various sources into useful information. This information can be saved to disk for future use. Some of the information extracted included: location of airports, distances between airports, passenger demand for routes per month, aircraft specifications, cost per hour of operation for each aircraft type, and airline fleets.
- Formulation of delay distributions. Accurate empirical delay distributions can be extracted from the raw performance data of airlines operating in the USA. These delay distributions can be fitted to standard distribution forms such as the exponential distribution or they can remain as empirical distributions. The distributions can be saved to disk for future use.
- Calculation of model parameters and creation of models. The program can use the relevant data to form the developed models for use by an external solver. The application generates a mathematical model in the ZIMPL format (Koch 2004) which includes the formulation and all the values required. The ZIMPL model can then be converted into either a MPS or LP file for use with various solvers. The index file (excludes data set files) of the ZIMPL model is shown in Appendix B. Calculations include: flight times, cost price per seat per flight, and profit per passenger for each flight. The model created forms and considers all the constraints within the developed models. All necessary decision variables, ground variables, constraints and coefficients of the model are created. The resultant model is written to disk.

- Formulation of flight schedule from outputted solver solution file. The application can read in the solution file generated by the external solver and translate it into the correct selection of flights and form the flight schedule. The flight schedule can be outputted into a human-readable format as well as additional file formats including: KML (for use in Google Earth, with animated flight arcs depicting the schedule) and GEXF (for use in network graphing applications).
- Monte Carlo Simulation. The program can evaluate the performance of various schedules using Monte Carlo simulation and gather the relevant results. The results and the entire Monte Carlo simulation log table can be saved to disk in a human-readable format.
- The application contains a graphical user interface (GUI) for the control of the application. The GUI provides an easy to use interface for generating and evaluating scheduling models.

The completed mathematica application consists of over 4000 lines of code and is included in the digital appendix. Mathematica is required for the installation and running of the application package.

3 STOCHASTIC FLIGHT DELAYS

A critical dimension of service quality for the airline industry is the “on time performance” of the airline operator; it is a measure of the airlines ability to operate according to their flight schedule and to minimize the number of delayed flights. The consequences of a flight delay are significant for the airlines; not only will it decrease customer goodwill, potentially delay other flights, and cause flight cancellations, but it is also becoming a critical measure by which customers select airlines (there are various websites which offer detailed information for the specific flights and routes which a customer may be considering). Flight delays are of stochastic nature and thus are difficult to predict accurately for any given day of flight operations; however there are large quantities of historic data which can be used to obtain distributions and information regarding flight delays. Since flight schedules are susceptible to flight delays, there is an advantage to considering these flight delay distributions during the flight scheduling modeling phase such that more robust schedules can be created. Presented below, in the remainder of this section, is an analysis of over half a million flights which occurred during June 2010 throughout the United States of America. The data was obtained from the TranStats[1] service provided by the Research and Innovative Technology Administration (RITA) and the Bureau of Transportation Statistics (BTS).

3.1 Overview Of On Time Performance

Presented in this sub-section is a summary of the on time performance of all the flights operated in the US during the month of June 2010. A total of 551 688 flights were scheduled, of which 317 795 flights departed either on or before their scheduled takeoff time and 333 019 flights landed at their destination on or before their scheduled arrival time. Thus 57,6% of departures and 60,4% of arrivals were on time respectively. The increase of the on time performance between departures and arrivals can be attributed to the following reasons: 1) pilots can decrease the flight time (in order to recover lost time) by increasing the aircraft’s speed at the cost of additional fuel consumption, 2) scheduled flight times are often slightly overstated such that there is a slight buffer in the schedule and the airlines on time performance is improved on paper.

[1] http://www.transtats.bts.gov/

The previous figures were determined with the presumption that any flight that operates later than its scheduled time is deemed to be delayed (even if the delay was only a single minute); however the general practice adopted in the industry is to only consider a flight delayed if the delay is 15 minutes or longer. This general practice was used when attributing the delays to various causes below. Table 3-1 presents a breakdown of the reasons attributed to the flight delays. The reasons are categorized as follows:

- *Carrier Delay:* a general delay attributed to the operations of the airline carrier,
- *Weather Delay:* a flight delay attributed to inclement weather conditions at either the origin or destination airports,
- *National Air System (NAS) Delay:* a flight delay attributed to the operations of the National Air System (air traffic control),
- *Security Delay:* a flight delay attributed to security issues,
- *Late Aircraft Delay:* a flight delay attributed to the late arrival of the aircraft required for the flight.

Table 3-1: Reasons For Flight Delays

Reason	# of Flights	Total Delay Time (minutes)	Average Length of Delay (minutes)
Cancelled	8279		
Carrier	58104	2023990	34,8
Weather	7227	324997	45,0
NAS	65427	1721650	26,3
Security	404	7033	17,4
Late Aircraft	62691	2854720	45,5

Presented in Figure 3-1 is a pie chart of the total delay time per attributed flight delay reason. It is a fair assumption that the airline operator has direct influence over the carrier delays and the late aircraft delays. These delays together account for 70% of the flight delays experienced.

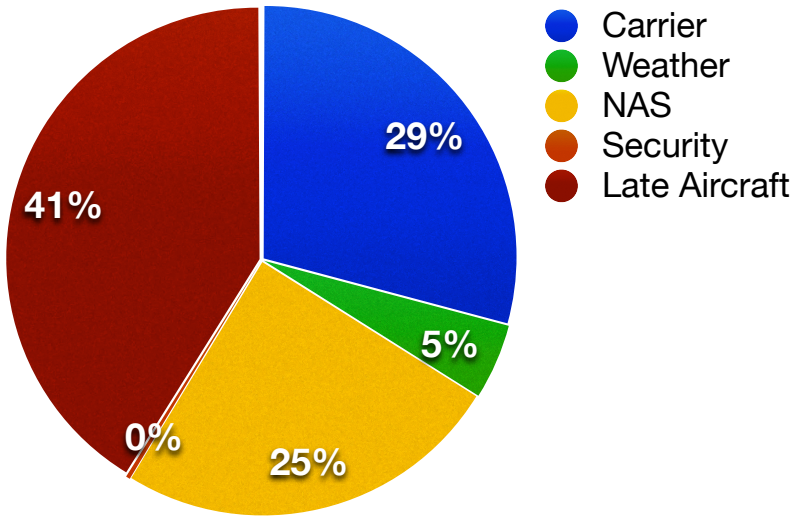


Figure 3-1: Percentage Breakdown Of Total Delay Time Per Attributed Flight Delay Reason

Flight delays can be split into two sections: departure delays and arrival delays. The following two sub-sections will analyze these sections in more detail.

3.2 Departure Delays

As mentioned above, only 57,6% of flights depart on or before their scheduled departure time thus 42,4% of flight departures are delayed; however a single measure is insufficient to fully describe the nature of the flight departure delays thus a more detailed analysis was conducted. Presented in Figure 3-2 is a histogram of the delay departures, the bin size is 5 minutes. The shape of the distribution can be represented using the Lognormal distribution with the μ and σ parameters set to 2,58 and 1,41 respectively. The probability plot (plots the cumulative distribution function, CDF, of the distribution versus the CDF of the fitted distribution) of the Lognormal distribution versus the delayed departures distributions is shown in Figure 3-3 to illustrate the fitness of the Lognormal distribution (the dashed line represents the target of a perfect fit).

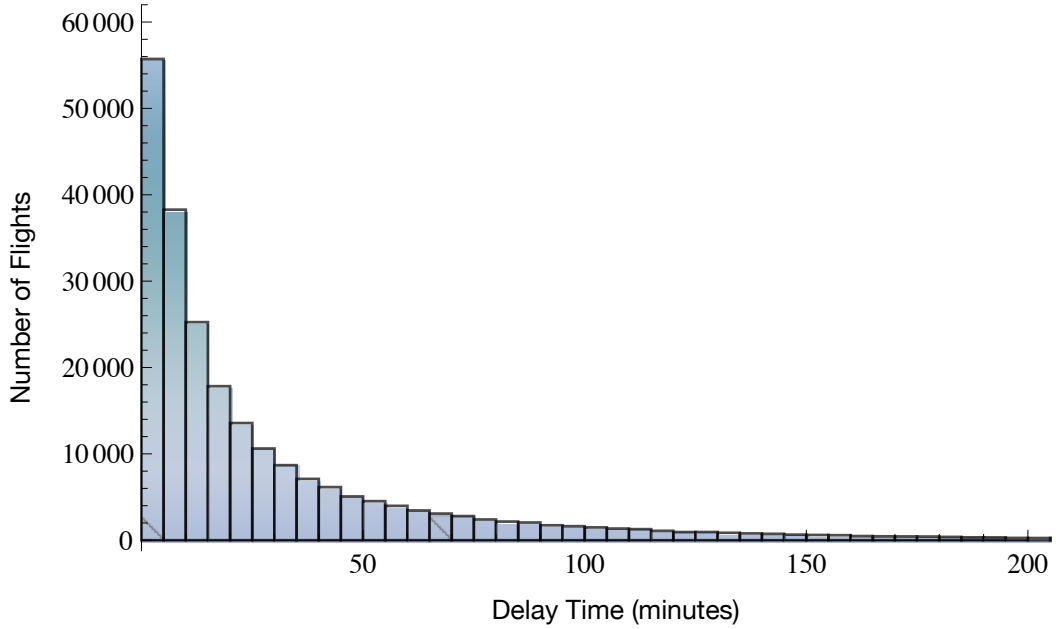


Figure 3-2: Histogram Of Delayed Departures

The distribution displayed in Figure 3-2 is useful (after the probability of a delay occurring is considered) for the Monte Carlo simulation used to determine the operating performance of the created flight schedules. It is, however, more appropriate to use the full distribution of departure times (including early and on-time departures) for the calculation of parameters for the operations research models. These parameters are calculated using one or more descriptive statistics measures. Below is a brief description of the descriptive statistic measures used to describe the departures distribution (shown in Figure 3-4) in Table 3-2.

- *Mean:* arithmetic mean is the sum of all the values divided by the number of values, commonly thought of as the average. The means presented are population means for the data.
- *Median:* the median separates the data into two halves of equal counts of values. It is found by arranging the values from lowest to highest value and picking the middle value such that there is an even number of values on either side of the median value.
- *Root Mean Square:* The root of the mean of the squares of the values, i.e. each number is squared and added together, the total is divided by the count of numbers in the set and this mean is rooted. Useful when the variates of the distribution are positive and negative.
- *Trimmed Mean:* the arithmetic mean with a small faction of the smallest and largest elements truncated from the data set (the outliers). A 20% trimmed mean was used.

- *Geometric Mean:* indicates the central tendency of a set of numbers (just like arithmetic mean); calculated by multiplying the numbers together and then taking the *n*th root (where *n* is the count of numbers in the set). Useful for data values that are meant to be multiplied together or are exponential in nature.
- *Variance:* variance is the measure of how far a set of numbers in a data set is spread out from the mean value. It is calculated by obtaining the mean of the squares of the distance between each value and the mean of the data set.
- *Standard Deviation:* is a measure of the variability or diversity of the values within a data set. It is the square root of the variance of a set.
- *Mean Absolute Deviation:* is the mean of the absolute deviation of each value from the mean.
- *Median Absolute Deviation:* is the median of the absolute deviations of each value from the median.
- *Maximum:* Largest value in the data set.
- *Minimum:* Smallest value in the data set.

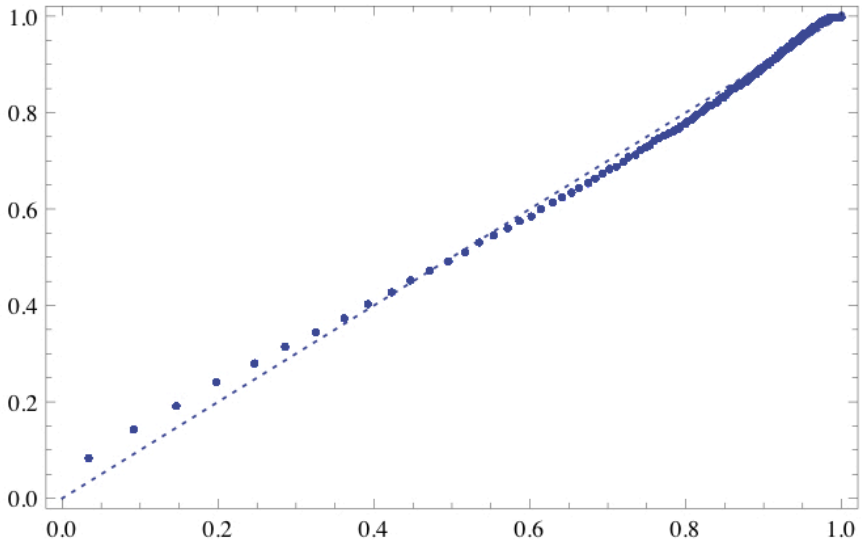


Figure 3-3: Probability Plot Of Lognormal Distribution Versus Delayed Departures Distribution

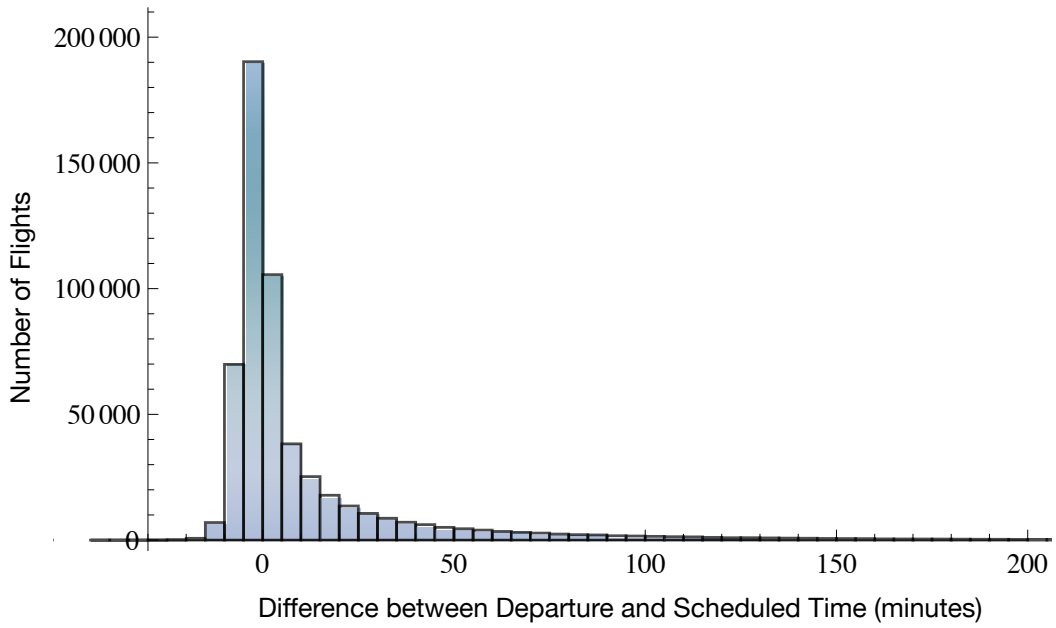


Figure 3-4: Histogram Of All Departure Times

Table 3-2: Descriptive Statistics of Departures Distribution (minutes)

Mean	11,3
Median	0
Root Mean Square	37,4
Trimmed Mean	1,0
Geometric Mean	0
Variance	1271,5
Standard Deviation	35,7
Mean Absolute Deviation	19,8
Median Absolute Deviation	5
Maximum	1324
Minimum	-59

3.3 Arrival Delays

This sub-section deals with arrival delays, the difference between scheduled landing time and actual landing time. Arrival delays are caused by a number of different factors including: 1) departure delays, 2) air traffic delays and 3) route adjustments due to weather, among others. Figure 3-5 presents a histogram (bin size of 5 minutes) of the distribution of the difference between the arrival time and scheduled arrival time for the analyzed flights, a negative number represents an early arrival. The possible reasons for early arrivals was mentioned in the opening paragraph of section 3.1 and is dependent on the length of the flight so that longer flights have a greater probability of arriving early. The distribution

presented in Figure 3-5 will be referred to as the *Standard Arrival Distribution* in the later comparisons.

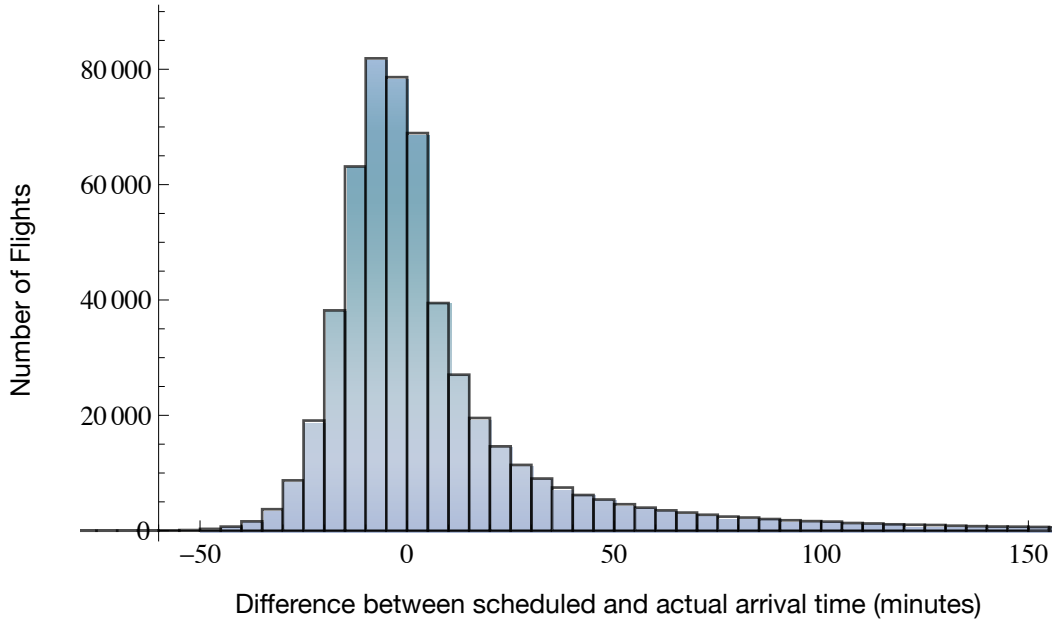


Figure 3-5: Histogram Of Difference Between Scheduled And Actual Arrival Times

Since the previous distribution is affected by the departure time of the aircraft, a useful distribution to derive is the distribution of the delay caused solely by the actual flight and the landing procedures and excluding the departure delay. This distribution is calculated by taking the arrival delay minus the departure delay for each flight. The resultant distribution will be known as the *Adjusted Arrival Distribution* and is presented in Figure 3-6 below. From the Adjusted Arrival Distribution it is found that 23% of flights incur an additional delay during the period after takeoff. Figure 3-7 represents the distribution of the additional arrival delays (the 23% of flights) which would be applicable for the Monte-Carlo simulation, this distribution will be referred to as the *Additional Arrival Delays Distribution*.

The descriptive statistics measures for the three arrival distributions is displayed in Table 3-3.

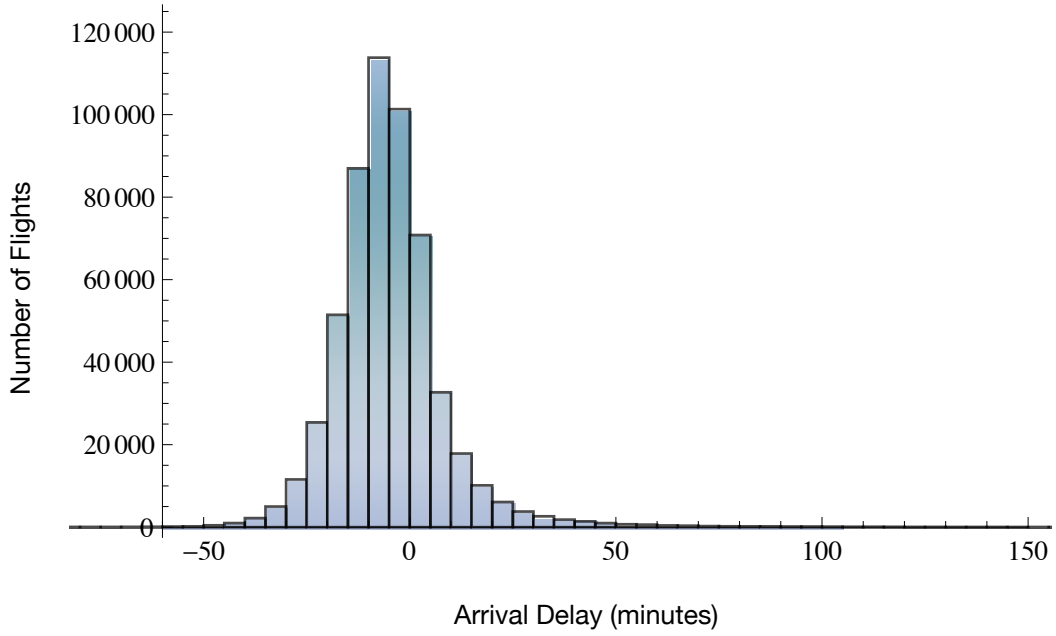


Figure 3-6: Histogram Of Adjusted Arrival Distribution

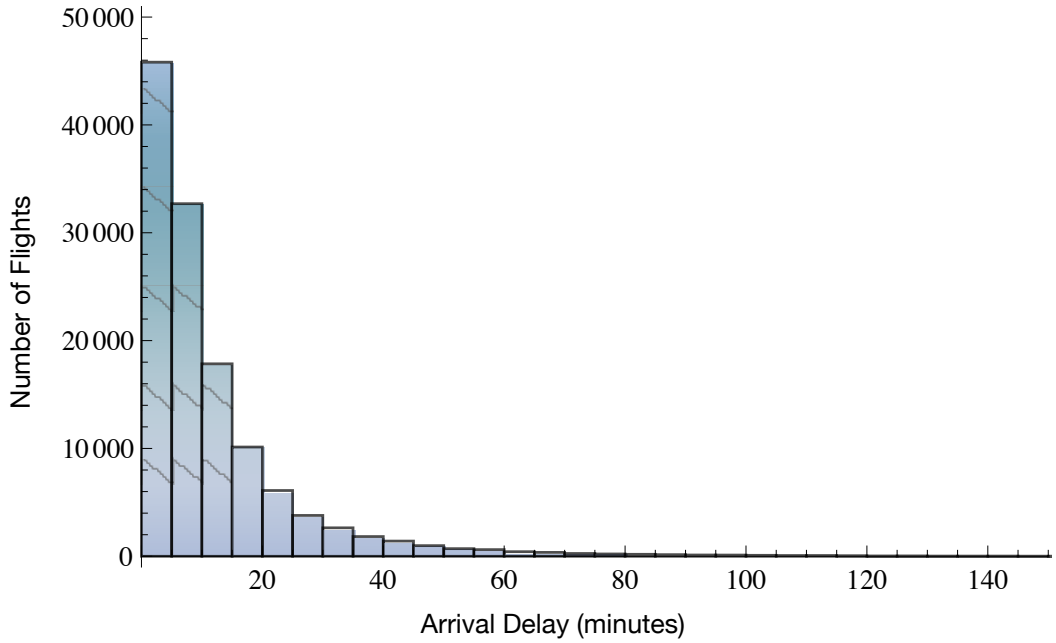


Figure 3-7: Histogram Of Additional Arrival Delays Distribution

Table 3-3: Descriptive Statistics Of Arrival Distributions (minutes)

	Standard Arrival Distribution	Adjusted Arrival Distribution	Additional Arrival Delays Distribution
Mean	7,9	-5,6	11,4
Median	-2	-6	7
Root Mean Square	38,9	15,0	18,2
Trimmed Mean	-1,0	-6,4	7,6
Geometric Mean	0	0	6,5
Variance	1449,4	194,9	200,1
Standard Deviation	38,1	14,0	14,1
Mean Absolute Deviation	22,1	9,2	9,1
Median Absolute Deviation	10	6	5
Kurtosis	49,1	43,0	
Maximum	1315	250	250
Minimum	-79	-646	1

4 OPTIMIZATION MODELS

This section presents and discusses the operations research mathematical models formulated to solve the airline scheduling problem. The discussion begins by introducing the traditional basic fleet assignment model.

4.1 Basic Fleet Assignment Model

The basic fleet assignment model is an application of the multi-commodity network flow (MCNF) problem formulation. A network, in operations research, is a directed graph containing nodes and arcs, which join the nodes together. Each arc has a flow capacity which limits the amount of flow that can travel along that arc for a specific period of time. The flow throughout the network must satisfy the *conservation of flow* constraint which restricts that the amount of flow entering into a node must be equal to the amount of flow leaving that node. The exceptions to the conservation of flow constraint are the *source* and *sink* nodes which are the start and end of the network respectively; however the amount of flow leaving the source node must be equal to amount of flow entering the sink node thus by artificially joining these nodes in the model the conservation of flow constraint is applied throughout the network. An example of a network is presented below in Figure 4-1. The dashed lines represent a particular feasible flow through the network. The conservation of flow constraint can be explained using node *a* as an example: the amount of flow entering node *a* is 5 and the amount of flow exiting node *a* is the sum of 3 and 2, hence the amount of flow entering and exiting the node is the same and the constraint is satisfied.

A multi-commodity network flow problem seeks to optimize the selection of which arcs to use to 'transport' each commodity whilst keeping the solution feasible with respect to the constraints of flow conservation, arc capacities and the demand satisfaction (ensuring that the required goods is transported through the network). It should be noted that a commodity's collection and delivery nodes need not be the source and sink nodes of the network structure; this is achieved by creating an additional source node linked to each required node (collection node) and an additional sink node linked to each delivery node. These additional source and sink nodes are linked to super-source and super-sink nodes respectively such that the basic required network structure is maintained.

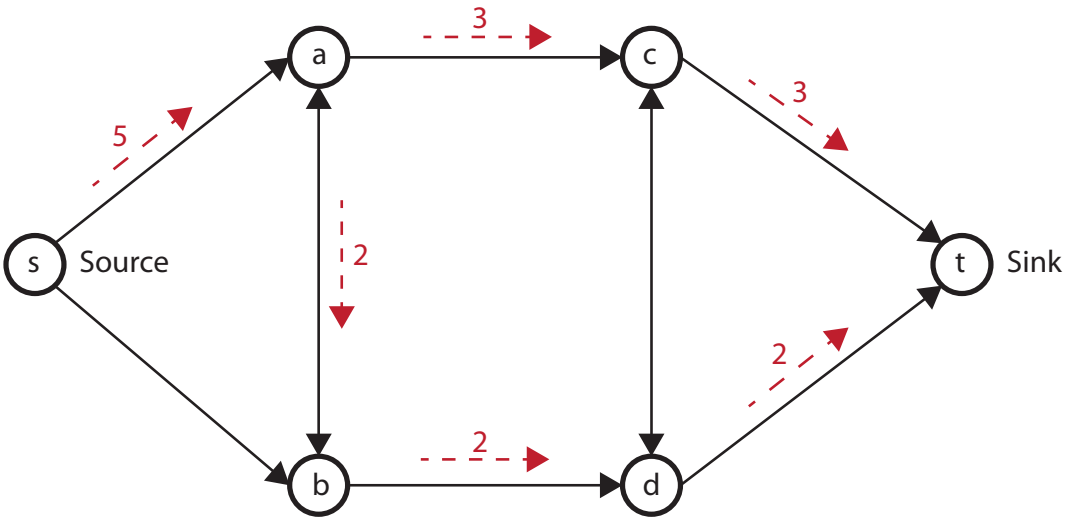


Figure 4-1: Flow Network Example

The airline's flight network is the foundation of the basic fleet assignment model. The use of a network structure in the formulation of the model ensures that only feasible flight connections are created, there are a balanced number of aircraft of each type at each airport at the beginning and end of each schedule (source and sink nodes), and that the conservation of aircraft flow is maintained. In the model the network is represented as follows. A node in the network represents a specific location at a specific time (for example, the location could be Johannesburg International Airport and the time could be 7:05 AM), nodes are created for each applicable point in the chosen schedule thus there will be a node for each event such as a departure or landing at a specific airport. For each flight in the schedule there will be a *flight arc* which connects the appropriate departure and arrival nodes in the network. *Ground arcs* represent the case when an aircraft remains on the ground at an airport between the event nodes, ground arcs are only connected between event nodes which represent the same airport and are connected sequentially as illustrated in Figure 4-2 below. Wrap-around arcs are special case ground arcs that connect the last event node of the schedule to the first event node for each airport, this replaces the source and sink nodes such that flow balance is guaranteed and that the final schedule is repeatable and cyclic in nature.

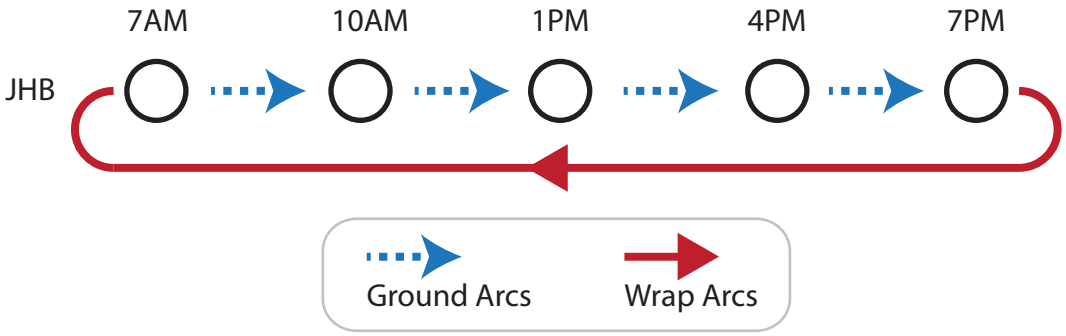


Figure 4-2: Ground And Wrap-Around Arcs

To ensure that only feasible flight connections are made, arrival nodes are placed at the flight's arrival time plus its turn-around time; thus the arrival node represents when the aircraft will be ready for its next flight. If the turn-around time is ignored then it is possible that the solution may require that a flight depart immediately after it has just landed which is infeasible because the aircraft will still need to be refueled and exchange passengers among other tasks.

The basic fleet assignment model considers a schedule containing flights (which have departure and arrival locations and specific departure times) and optimizes the selection of which fleet should be used for each flight in the schedule such that the total cost is minimized. The network formulated for the model (discussed above) is duplicated for each different fleet in the model, and costs are assigned for each flight arc. The cost for each flight arc is calculated using the operating costs for that specific fleet (and may consider factors such as fuel consumption and crew requirements).

Once the flight networks for each fleet are created, the basic fleet assignment model is formulated as shown in Figure 4-3, whilst the notation used is given in Table 4-1. The formulated model seeks to minimize the total cost to cover each flight in the schedule with only one aircraft type (first constraint), it ensures that the aircraft flow is conserved at each node with the second constraint and the third constraint ensures that the total number of aircraft used in each fleet does not exceed the amount available. The third constraint ensures the amount of aircraft used is feasible by counting the amount of aircraft at a specific count time (usually at the time of least activity, perhaps 3 AM). The sum of all the aircraft leaving the count time nodes (this includes both ground arcs and flight arcs) plus the sum of all the aircraft on flight arcs that occur during the count time, results in

a measure of the amount of aircraft used which can be constrained to be less than the amount available.

Decision variables *x* define which fleet type covers each flight, this variable is defined as a binary variable and those may only take the value of 0 or 1. If the solution value of a variable x_{ik} is 1 then that flight *i* is to be flown by that fleet *k*. The zero-cost ground arc variables, *y*, need not be integer or binary variables because more than one aircraft of a fleet type may remain on the ground at a specific airport at the same time and the integrality of the network structure ensures that the value will an integer; however *y* must be non-negative.

Table 4-1: Notation For The Basic Fleet Assignment Model

F	=	set of flights
K	=	set of fleets
S_k	=	number of aircraft of fleet *k*
G_k	=	set of ground arcs in fleet *k*'s network
L_k	=	set of nodes in fleet *k*'s network
c_{ik}	=	cost to fly flight *i* with fleet type *k*
x_{ik}	=	1 if flight *i* flown by fleet *k*, 0 otherwise
y_{gk}	=	number of aircraft on ground arc *g* in fleet *k*'s network
$b1_{lik}$	=	1 if flight *i* begins at node *l* in fleet *k*'s network
		-1 if flight *i* ends at node *l* in fleet *k*'s network, 0 otherwise
$b2_{lgk}$	=	1 if ground arc *g* begins at node *l* in fleet *k*'s network
		-1 if ground arc *g* ends at node *l* in fleet *k*'s network, 0 otherwise
$d1_{ik}$	=	1 if flight *i* crosses the count time in fleet *k*'s network, 0 otherwise
$d2_{gk}$	=	1 if ground arc *g* crosses the count time in fleet *k*'s network, 0 otherwise

$$\text{Minimize} \quad \sum_{k \in K} \sum_{i \in F} c_{ik} x_{ik}$$

subject to

$$\sum_{k \in K} x_{ik} = 1 \qquad \forall i \in F \qquad (1)$$

$$\sum_{i \in F} b1_{lik} x_{ik} + \sum_{g \in G_k} b2_{lgk} y_{gk} = 0 \qquad \forall i \in L_k, \forall k \in K \qquad (2)$$

$$\sum_{i \in F} d1_{ik} x_{ik} + \sum_{g \in G_k} d2_{gk} y_{gk} \leq S_k \qquad \forall k \in K \qquad (3)$$

$$x_{ik} \in \{0, 1\} \qquad \forall i \in F, \forall k \in K \qquad (4)$$

$$y_{gk} \geq 0 \qquad \forall g \in G_k, \forall k \in K \qquad (5)$$

Figure 4-3: Basic Fleet Assignment Model

4.2 Fleet Assignment With Variable Time Windows Model

The basic fleet assignment model works with the premise that a flight schedule is already available and that the only task is to assign specific fleets to each flight in an optimal manner. Whilst this simplifies the fleet assignment problem, it does so by reducing the opportunity of locating better solutions. The notable weakness is the requirement that flight departure times are fixed and thus the flexibility of the model is reduced. An example of where these fixed departure times may lead to an inferior answer is: Suppose that flight A lands at airport 1 at 10 AM and flight B departs from the same airport at 9:45 AM; it is clear that the same aircraft cannot be used for both flights because flight B departs before flight A lands. However, if the model is formulated such that the departure times for each flight can be adjusted then it is possible for a single aircraft to serve both flights as the departure time of flight B can be moved, for example, from 9:45 AM to 10:15 AM; this can lead to significant cost savings as it is possible that fewer aircraft will be required. Levin (1971) was the first to propose the use of variable time windows in a scheduling and fleet routing model. The time windows were modeled to allow departure times to occur at discrete intervals within a time window. For example, if a flight is required to depart some time between 8 AM and 10 AM to maintain passenger demand (if the flight is moved too far from its original schedule time then it could lose passenger demand) for that flight then the time window will be 8 AM to 10 AM. This time window can be divided into a set number of discrete intervals such that possible departure times could be 8 AM, 8:15 AM, 8:30 AM, etc... The model will be able to determine the optimal departure time for each flight such that the objective function is optimized. The objective function could be formulated in such a manner that factors such as gate availability, passenger connection times, and robustness can be improved. Desaulniers et al. (1994) extended Levin's model to consider more than a single fleet. Desaulnier's model is the basis for the fleet assignment with variable time windows model (FAVTW) presented here.

The time windows can be modeled by creating copies of the flight arc at a specified interval (such as 5 minutes) within that flight's time window and then adding the constraint that only one of the flight copies need to be selected in the solution. This allows the model to choose the departure time of the flight. When considering the interval to be selected for each problem, it is important to note that using a narrow interval will result in a more flexible model but will also result in an 'explosion' in the model size. A general rule of thumb is that the time taken

to solve a model grows exponentially with the size of the problem (number of variables, constraints, etc...) thus the selection of the interval must be small enough to ensure flexibility whilst wide enough to allow acceptable solving times. This selection is mostly determined through trial and error.

The formulation of the FAVTW model is almost the same as with the basic fleet assignment model; the differences are the size of the resulting model and the constraints used. The constraints work exactly as before, in the basic fleet assignment model, however they just consider more nodes in the network due to the extra flight arcs that were created. The notation that is different from the basic fleet assignment model is given in Table 4-2 and the model formulation is shown in Figure 4-4.

Table 4-2: Notation For Fleet Assignment Model With Variable Time Windows Model

N_{ik}	=	number of arc copies of flight *i* in fleet *k*'s network
c_{nik}	=	cost to fly copy *n* of flight *i* with fleet type *k*
x_{nik}	=	1 if copy *n* of flight *i* flown by fleet *k*, 0 otherwise
$b1_{lnik}$	=	1 if copy *n* of flight *i* begins at node *l* in fleet *k*'s network -1 if copy *n* of flight *i* ends at node *l* in fleet *k*'s network, 0 otherwise
$d1_{nik}$	=	1 if copy *n* of flight *i* crosses the count time in fleet *k*'s network, 0 otherwise

$$\text{Minimize} \sum_{k\in K}\sum_{i\in F}\sum_{n\in N_{ik}} c_{nik}x_{nik}$$

subject to

$$\sum_{k\in K}\sum_{n\in N_{ik}} x_{nik} = 1 \qquad \forall i \in F \qquad (1)$$

$$\sum_{i\in F}\sum_{n\in N_{ik}} b1_{nlik}x_{nik} + \sum_{g\in G_k} b2_{lgk}y_{gk} = 0 \qquad \forall i \in L_k, \forall k \in K \qquad (2)$$

$$\sum_{i\in F}\sum_{n\in N_{ik}} d1_{nik}x_{nik} + \sum_{g\in G_k} d2_{gk}y_{gk} \le S_k \qquad \forall k \in K \qquad (3)$$

$$x_{nik} \in \{0,1\} \qquad \forall i \in F, \forall k \in K, \forall n \in N_{ik} \qquad (4)$$

$$y_{gk} \ge 0 \qquad \forall g \in G_k, \forall k \in K \qquad (5)$$

Figure 4-4: Fleet Assignment With Variable Time Windows Model

The first constraint ensures that only one flight arc among all the copies and every fleet is selected for each flight required. This ensures that the model picks a single departure time for each flight. Since each flight copy has its own variable, it is possible to assign different costs to each flight copy; this is useful if one wants to make it more expensive to fly a flight the further away it is moved from its original scheduled time.

4.3 Fleet Assignment With Time Windows And Route Selection Model

The previous two models presented rely on the prior selection of routes to be flown and to a lesser extent the time at which the flight should occur. For the purpose of this research project the following model was developed to address additional considerations in the flight scheduling problem. The model is an extension of the FAVTW model and will be referred to as the fleet assignment with time windows and route selection (FARS) model. The FARS model doesn't use a proposed flight schedule as its starting point but rather the forecasted demand for each possible flight route that the airline is considering; the FARS model then proceeds to determine the best selection of routes to fly (to maximize profit), the fleet type which should fly each flight and the departure time of the flight. There are clear advantages to this approach because the optimization model is now considering a larger part of the decision process and thus will output a superior solution at the cost of additional computational time. It is in an airlines best interest to only fly routes which will profitable and not to fly routes which will incur a loss of profit; the FARS model considers this fact and optimizes the objective function which considers profit. It is possible that an unprofitable route is selected in the solution but only if that route allows additional profitable routes to be flown such that the total profit is increased. The time windows within the FARS model can be adjusted such that the time window encompasses the entire day, this allows the FARS model full flexibility to adjust the departure times of the flights. However, it still allows the time windows to be constrained to a smaller period of time thus ensuring that certain flights only occur at certain times during the day (useful for the high demand of morning flights).

The following paragraph describes the changes that were made to the previous models to form the FARS model. Firstly, for each possible route there is a demand value assigned; this demand value represents the average amount of passengers expected to fly that route if the capacity of the aircraft wasn't an issue. Secondly, in addition to the number of aircraft available for each fleet, the capacity of the aircraft in each fleet is introduced into the model (for example, fleet 1 may have aircraft capacities of 150 passengers per flight). The FARS model changes the structure of the first constraint of the other models. In the other model, the constraint requires that each flight on the schedule be flown once and only once out of all the possible flight copies and fleet types. The FARS model does not enforce that each route be flown and allows the route to be flown

more than once. A specific flight arc can still only be selected once because of the use of binary variables, this is intentional as it is undesirable to have two flights with the same destination departing at the same time at a single airport (this would cause much confusion among passengers as to which aircraft they should be on). The first constraint in the FARS model adds together the total capacity of flights selected for a route minus a new penalty variable (real number, non-negative) and constrains this total to be less than or equal to the passenger demand for that route. In essence, this constraint is recording the over-capacity of the flights for that route. For example, if a route has a demand of 350 passengers and two aircraft are selected to fly that route with capacities of 180 and 200 respectively then the penalty variable will be equal to 30, corresponding to the number of empty seats for that route. These penalty variables are used in the objective function. The objective function for the FARS model now considers the expected profit of the schedule instead of the expected cost and as such now wishes to maximize that profit. The profit is calculated by calculating the revenue for each flight arc if the flight was completely full minus the cost of the flight minus the product of the penalty variables and the respective seat price for that route. The seat prices (amount that the passenger would pay to be on that flight) are assumed to be fixed for each route regardless of the fleet type used; this assumption is applicable because a passenger wouldn't expect to pay more or less for their ticket depending on which aircraft is used, instead they would expect the price to be affected by route and class of seat.

The notation that is different from the previous models is given in Table 4-3 and the model formulation is shown in Figure 4-5.

Table 4-3: Additional Notation For Fleet Assignment With Time Windows And Route Selection Model

p_i	=	penalty variable for flight *i*
q_i	=	seat price for flight *i*
dm_i	=	passenger demand for flight *i*
z_{nik}	=	profit of copy *n* of flight *i* with fleet type *k*
h_k	=	aircraft seat capacity of fleet type *k*

$$\text{Maximize} \quad \sum_{k\in K}\sum_{i\in F}\sum_{n\in N_{ik}} z_{nik}x_{nik} - \sum_{i\in F} p_i q_i$$

subject to

$$\sum_{k\in K}\sum_{n\in N_{ik}} h_k x_{nik} - p_i \le dm_i \qquad \forall i \in F \qquad (1)$$

$$\sum_{i\in F}\sum_{n\in N_{ik}} b1_{nlik}x_{nik} + \sum_{g\in G_k} b2_{lgk}y_{gk} = 0 \qquad \forall i \in L_k, \forall k \in K \qquad (2)$$

$$\sum_{i\in F}\sum_{n\in N_{ik}} d1_{nik}x_{nik} + \sum_{g\in G_k} d2_{gk}y_{gk} \le S_k \qquad \forall k \in K \qquad (3)$$

$$x_{nik} \in \{0,1\} \qquad \forall i \in F, \forall k \in K, \forall n \in N_{ik} \qquad (4)$$

$$y_{gk} \ge 0 \qquad \forall g \in G_k, \forall k \in K \qquad (5)$$

$$p_i \ge 0 \qquad \forall i \in F \qquad (6)$$

Figure 4-5: Fleet Assignment With Time Windows And Route Selection model

Without changing the mathematical formulation of the model it is still possible to tailor the model to the specific application required. Adjusting parameters, variable coefficients and the creation of variables can allow additional control of the resultant schedule. It is possible to exclude certain routes for certain fleets or for all the fleets, specify when certain flights occur and even fix their departure time, specify that for a specific route that one flight always occur during a certain period of the day and if the demand is high enough then any more flights for that route can be added at any time during the schedule, specify the operating hours for each fleet or airport or flight, specify the number of passengers that represent a break-even cost point for each flight, by adjusting the profit or cost coefficients of the variables it is possible to favor specific flights and departure times, and there is potentially even more control of the schedule possible.

4.3.1 Schedule formation from model solution

The solution to the FARS model would contain a list of the selected flight arcs in the model, the following section will describe the process of formulating the final flight schedule from this solution.

The solution to the FARS model will be outputted by the solver in the form of a list of selected variables; this list of selected variables contains no further information regarding the flights chosen. The first step for forming the schedule is to cross-reference the list of selected variables with a table created during the formulation and calculation of the model and its parameters; this will result in a table listing the selected flights and containing information such as departure time, origin and destination airports, landing time and fleet type among others. A visual example of this first step is presented in Figure 4-6 below.

Once the table containing the information about the selected flights has been created, it is still difficult to determine the final schedule, especially with larger models, since the flights are not in order or assigned to specific aircraft (tail assignment). There is a simple iterative procedure that will form the final schedule and it is as follows:

1. Sort the table by departure time.
2. Split data in a table for each fleet type (whilst maintaining the sorted order).
3. Select one of the fleet tables created in step 2.
4. Select the first flight on the table and assign it to an unused aircraft of the correct fleet type, and remove the entry from the table.
5. Take note of the previously selected flights destination and landing time.
6. Search through the table in order and select the first flight that satisfies the following criteria: origin location is the same as the previously selected flights destination and its departure time is after the landing (and turn-around) time of the previous flight. Assign the newly selected flight to the current aircraft and remove the flight from the table.
7. Repeat steps 5 and 6 until no more flights can be selected.
8. If table is not empty then proceed to step 4.
9. If table is empty then proceed to next table for another fleet and proceed to step 4.
10. If all tables are empty then scheduling is complete.

The resultant schedule will maintain the flow of aircraft correctly and the schedule will be repeatable in a cyclical manner. An example of a final schedule represented within a time-space network is presented in Figure 4-7 for a two aircraft fleet with a time window interval of 3 hours (for space considerations).

Solution Variables
x_{25}
x_{215}
x_{245}
x_{524}
x_{1045}
x_{1755}

Variable	Origin	Destination	Departure Time	Landing Time	Fleet
1	JHB	CPT	8:00 AM	10:05 AM	1
2	JHB	CPT	8:15 AM	10:20 AM	1
3	JHB	BLM	10:00 AM	11:25 AM	2
4	DUR	CPT	1:00 PM	2:30 PM	3
5	CPT	BLM	10:00 AM	12:30 PM	1

Cross-Reference

Origin	Destination	Departure Time	Landing Time	Fleet
JHB	DUR	7:00 AM	8:30 AM	Airbus A319
BLM	CPT	4:00 PM	6:00 PM	Airbus A321
BLM	JHB	1:00 PM	2:30 PM	Airbus A319
DUR	BLM	10:00 AM	11:15 AM	Airbus A319
CPT	JHB	7:00 AM	9:30 AM	Airbus A321
JHB	BLM	10:00 AM	11:30 AM	Airbus A321

Figure 4-6: FARS Model Schedule Formation Step 1

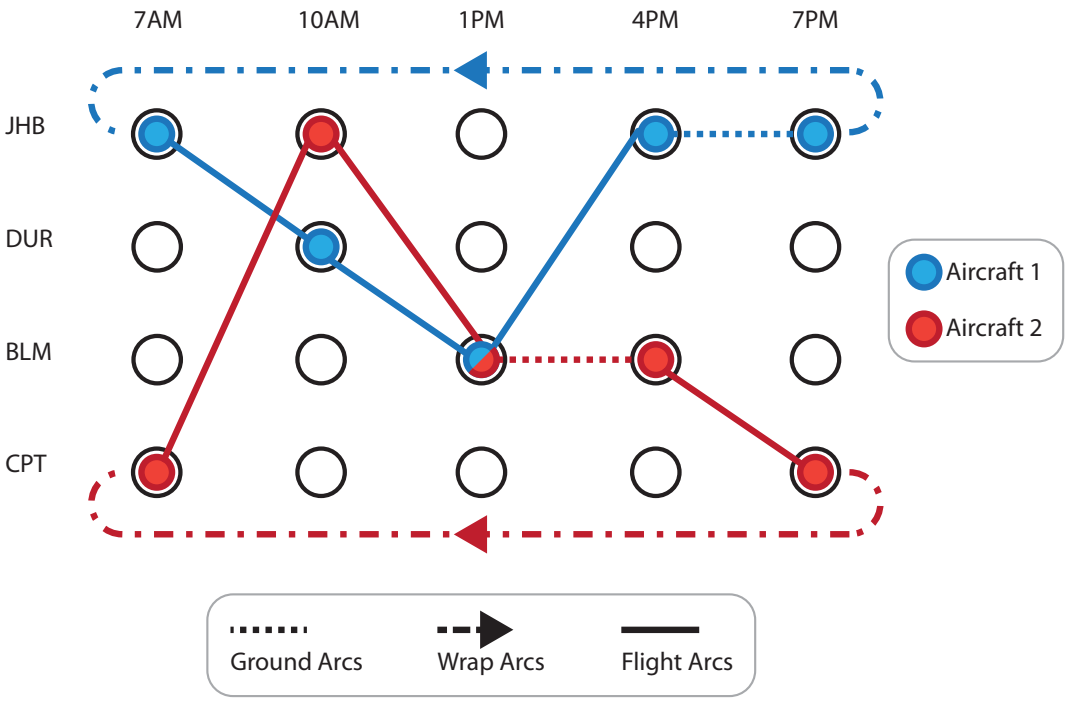


Figure 4-7: Example Schedule Represented On A Time-Space Network

4.4 Stochastic Models

This section will introduce and discuss the adjustments made to the FARS model such that the stochastic nature of the flight delays is considered during the formulation and solving of the model. The calculation and specification of the flight time for each flight variable in the model is conducted during the formulation

stage of the model (generating the MIP program to be solved by an appropriate solver) and as such the flight times are fixed once the model is generated. It is not beneficial to model the problem in such a manner that flight times can be adjusted during the solving of the model for various reasons. Each change in a flight time would require a re-specification of the network flow constraints (which would make the problem more computationally difficult to solve), additional constraints would need to be introduced to handle the limits of the flight times, and the constant changing of flight times would affect the objective values which would introduce problems with improving the solution and selecting an optimal solution. An effective way of handling stochastic flight delays is to consider them as part of the flight time for each variable in the model. The consequence of this technique is that the stochastic behavior of the flight delays are handled during the formulation stage of the model and thus the model becomes deterministic in nature (in respect to the solver). The main advantage of this consequence is that computational time required to solve the model will be drastically reduced.

The FARS model discussed earlier, is in essence a naïve model which handles the stochastic nature of flight delays by simply ignoring them and assuming that no flight delays will occur. The consequence of this assumption is that potentially the final schedule generated will not be robust in handling flight delays and may not even be suitable for real-world applications. To improve the robustness of the final schedule, the flight delays should be considered.

4.4.1 Expected value (EV) model

The simplest stochastic programming technique is the Expected Value (EV) method. The EV method replaces the stochastic distribution of a particular variable with its expected value. The expected value of a random variable, whose behavior is described by a distribution, is the weighted average (weighted by probability of occurrence) of all the possible values that the random variable could be; thus it is equivalent to the mean value, or first moment, of a statistical distribution. Integrating the Expected Value method into the FARS model is a simple process. There are two flight delays to be considered for each flight, namely: the departure delay (time difference between the scheduled departure time and the actual departure time of the flight) and the arrival delay (time difference between the scheduled arrival time and the actual arrival time of the flight). The arrival time delay is strongly influenced by the departure delay thus an independent arrival delay distribution is required; a flight duration delay time

distribution is calculated to address this issue. The flight duration delay time is the difference between the departure delay time and the arrival delay time; it is, in other words, the difference between the scheduled flight duration time (from gate to gate, not actual flying time) and the actual flight duration time. The overall flight delay is the sum of the departure delay and the flight duration delay. The calculation of the expected values for these distributions is a trivial exercise and thus an expected flight delay can be obtained (for example: 15 minutes). The introduction of the expected flight delay into the FARS model is rather simply: the sum of the expected flight delay and the flight time is used instead of solely the flight time for each variable.

4.4.2 Scenario generation with chance constrained programming

The following model makes use of two stochastic programming techniques, namely: scenario based analysis and chance constraint problems (CCP). The scenario based approach introduces the stochastic variation into the model whilst the CCP constraints ensure that the model remains feasible even if an extreme value is sampled from the stochastic distribution. First, the general concept of how the model works is discussed, then the mathematical formulation of the model is introduced, and finally a deterministic heuristic which greatly simplifies the model is presented.

The scenario and CCP based model (S-CCP) is based on the FARS model introduced earlier; the FARS model is included (and unchanged) in the S-CCP model, however a variety of additional variables and constraints are introduced into the model. The decision variables of the S-CCP model are identical to that of the FARS model and no additional decision variables are added; additionally all the constraints of the FARS model are still present, such as the network flow constraints and the aircraft cover constraints. The flight times of the decision variables in the FARS model portion of the S-CCP model do not include any flight delays (essentially the naive model). The first new addition to the model is that of the generation of multiple scenarios. Each flight variable in the FARS model is duplicated a specified number of times with only a single change made, namely: the flight delay. The duplicates have the same take-off time, profit, origin and destination as the original variable; however, because of the flight delay the duplicates now have different arrival times. The scenarios can be thought of as independent levels of the model. Each scenario (level) needs to be constrained by the network flow such that the scenario is feasible. The scenario duplicate

variables are selected in the model by linking them directly to the base level's decision variables; such that if a certain flight variable is selected then all the appropriate (linked) scenario duplicate variables also have to be selected. For example, if the flight variable which represents a flight from Johannesburg to Durban departing at 8 AM is selected then all the duplicate variables (one per scenario) which share the same departure time, origin, aircraft, and destination must also be selected. Supposing that the example flight from Johannesburg to Durban takes one hour to complete (including the turn-around time), then it is possible that the same aircraft could be used for another flight from Durban to Cape Town which departs at 9:15 AM. In the naive FARS model the network flow constraint will be satisfied and the schedule will be feasible; however, consider that one of the scenarios generated has a flight delay of 30 minutes for the flight from Johannesburg to Cape Town then the same aircraft would not be able to be used for the following flight because the network flow constraint for that scenario's level would not be satisfied. Such a case would require that the solver choose another solution so that all the scenario's network flow constraints are satisfied (such as moving the first flight earlier, or moving the second flight later, or not flying the flights). If all the constraints are to be satisfied then this approach can be simplified to only considering the worst-case of each flight variable (and its duplicates). The problem with this worst-case structure is that a single scenario flight can sample an extreme value from the distribution, at the end of the tail, that effectively eliminates that flight from the feasible region of the solution space. This is an undesirable situation.

A possible solution to this problem is the use of CCP constraints. A CCP constraint can be formulated in such a manner that only a certain percentage of the network flow constraints on the scenario levels for a particular node (time-space destination) need to be satisfied. The CCP constraint will only be used for the scenario levels and the normal base level will retain its own network flow constraints which must still be satisfied; this ensures that a feasible schedule is always generated. To explain the principle further the following example is presented: Flight 1 departs at 8 AM and the scheduled flight time is one hour, flight 2 departs at 9:30 AM from the destination of the first flight; 5 scenarios are generated and the CCP constraint is formulated such that only 1 of the scenarios may not pass their network flow constraint. A basic visual representation is presented in Figure 4-8 below.

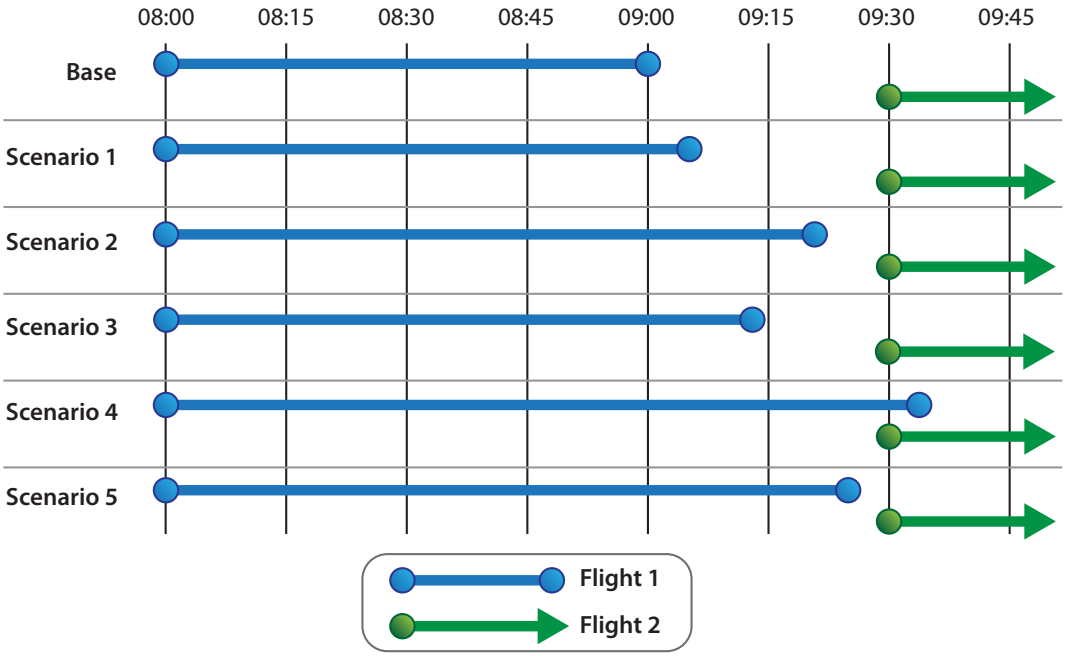


Figure 4-8: CCP Visual Example

The situation depicted in Figure 4-8 would be a feasible solution if the CCP constraint was set such that at most only one network flow constraint may be violated, in this case scenario 4 is violating its network flow constraint because flight 2 departs before its aircraft is available. If more than one network flow constraint was violated then the solution would be infeasible. In the example depicted, the CCP constraint ensures that the flight schedule is feasible and operates without a propagated delay 80% of the time.

The mathematical formulation of the deterministic equivalent of the S-CCP model is now presented. Presented in Table 4-4 is the model notation that is different from the previous models and the model formulation is shown in Figure 4-9.

Equation 4 for the S-CCP model ensures that all the scenario duplicate variables are selected when the "linked" decision variable is selected. Equation 5 is the CCP constraint which enforces the network flow constraints for each node for each scenario but also allows a specific number of network constraints to not hold which is equal to the value assigned to α. The duplicate scenario variables are also required to be binary.

Table 4-4: Additional Notation For The Scenario Generation With Chance Constrained Programming Model

w	=	Number of Scenarios Generated
v_r	=	1 if flight duplicate r is flown, 0 otherwise
R_{nik}	=	Set of duplicate scenario flight arcs for copy *n* of flight *i* flown by fleet type *k*
Δ	=	Set of scenario levels
δ	=	Scenario level
α	=	CCP constraint tolerance
$j1_{r\delta l}$	=	1 if duplicate flight *r* begins at node *l* in scenario δ's network -1 if duplicate flight *r* ends at node *l* in scenario δ's network, 0 otherwise
$j2_{r\delta gl}$	=	1 if ground arc *g* begins at node *l* in scenario δ's network -1 if ground arc *g* ends at node *l* in scenario δ's network, 0 otherwise

$$\text{Maximize} \quad \sum_{k\in K}\sum_{i\in F}\sum_{n\in N_{ik}} z_{nik}x_{nik} - \sum_{i\in F} p_i q_i$$

subject to

$$\sum_{k\in K}\sum_{n\in N_{ik}} h_k x_{nik} - p_i \le dm_i \qquad \forall i \in F \qquad (1)$$

$$\sum_{i\in F}\sum_{n\in N_{ik}} b1_{nlik}x_{nik} + \sum_{g\in G_k} b2_{lgk}y_{gk} = 0 \qquad \forall i \in L_k, \forall k \in K \qquad (2)$$

$$\sum_{i\in F}\sum_{n\in N_{ik}} d1_{nik}x_{nik} + \sum_{g\in G_k} d2_{gk}y_{gk} \le S_k \qquad \forall k \in K \qquad (3)$$

$$wx_{nik} - \sum_{r\in R_{nik}} v_r = 0 \qquad \forall i \in F, \forall k \in K, \forall n \in N_{ik} \qquad (4)$$

$$\sum_{\delta\in\Delta}\left|\sum_{i\in F}\sum_{r\in R_{nik}} j1_{r\delta l}v_r + \sum_{g\in G_k} j2_{r\delta gl}y_{gk}\right| \le \alpha \qquad \forall i \in L_k, \forall k \in K \qquad (5)$$

$$v_r \in \{0,1\} \qquad \forall r \in R_{nik} \qquad (6)$$

$$x_{nik} \in \{0,1\} \qquad \forall i \in F, \forall k \in K, \forall n \in N_{ik} \qquad (7)$$

$$y_{gk} \ge 0 \qquad \forall g \in G_k, \forall k \in K \qquad (8)$$

$$p_i \ge 0 \qquad \forall i \in F \qquad (9)$$

Figure 4-9: Scenario Generation With Chance Constrained Programming Model

If the CCP constraint tolerance variable (α) is set to zero then effectively the model is constrained such that all the network flow constraints are required to hold for a feasible solution. Since all the flight departure times are fixed for each trial solution (as shown in Figure 4-8, above), the only scenario case we need to worry about for each flight is the worst case (longest delay time) since if the worst case fits then all the others will fit too. Thus, if α is set to zero then the model can be simplified such that it is equivalent to the FARS model but with the flight times equal to the respective worst case value in the S-CCP model. This simplification can be modified such that it applies to any value of α. Consider the case when the S-CCP model consists of 5 generated scenarios for each flight and the CCP

constraint tolerance value α is set to one (i.e. one scenarios network constraint may be violated for each node in the network); it can be deduced that the model can be simplified such that the only flight time that is important for each flight is the fourth worst since if that scenario is valid then at least four of the five scenarios will be valid and the CCP constraint will be satisfied. Thus the simplification is that for each flight, the α highest flight times can be ignored and that all of the remaining flight times besides the highest remaining flight time can also be ignored.

Consider the case when the number of scenarios generated for the S-CCP model approaches infinity. As the number of scenarios approaches infinity then there will be an almost infinite amount of flight duplicate variables generated which each have a randomly sampled flight delay from a specific flight delay distribution. Due to the large number of samples for each flight, the distribution of flight delays for each flight will approach identity with the distribution from which they are sampled from. A large number of scenarios generated would result in a very large (and difficult to solve) S-CCP model and this is undesirable; however, taking advantage of the fact that the model effectively contains the identical distribution of flight delays as the sample distribution and this coupled with the CCP constraint simplification is particularly significant. If α is set to a small number such as one, then it is apparent that the model will effectively become a worst-case scenario model; this is not of much use since the worst case scenario is typically a flight cancellation. Thus the schedule would be empty; however if α is set to be a certain percentage of the number of scenarios generated then the constraint becomes useful. The CCP constraint now becomes a direct probability constraint that allows us to set the desired level of operational performance with respect to the handling of flight delays. For example, if α is set to be equal to 80% then the effective requirement is that at least 80% of the time the flights in schedule will not cause the next flight to occur a propagated delay; or, in other words, that at most 20% of the times that a flight is operated will it be subjected to a large enough flight delay that it will also delay the following flight in the schedule.

The above simplifications of the CCP constraint and the introduction of an infinite number of scenarios result in a deterministic model which can be identical to the FARS model in mathematical formulation. The only adjustment that needs to be made to the FARS model to make it equivalent to the infinitely large S-CCP

model is the selection of the flight time (which includes the flight delays). The calculation of the flight time is the sum of the expected flight time and the critical value of the flight delay distributions. The critical value of the flight delay distributions is equivalent to using the inverse cumulative distribution function of the flight delays distribution evaluated at *α.* From probability theory, the cumulative distribution function CDF is defined by:

$$CDF(x) = \Pr[\alpha \leq x]$$

The value of evaluated expression is the corresponding area underneath the curve (probability). Its inverse specifies for each probability level, the point for which the integral equals the probability level. This value is equivalent to a delay time which ensures that the CCP constraint holds for the particular level set by *α.*

The infinite S-CCP simplification model will be used for the project as it can be adjusted to represent the naive model, the expected value model and the S-CCP model.

5 MONTE CARLO SIMULATION

Monte Carlo Simulation, otherwise known as Stochastic Simulation, is a simulation technique which relies on repeated random sampling of distributions in order to approximate the stochastic behavior of the actual system. Monte Carlo Simulation is used when a deterministic algorithm is unavailable for the evaluation of stochastic system. Monte Carlo Simulation is being used in this project to evaluate the “real-world” performance of the various schedules generated with the stochastic events being the departure delays and flight length (time between take-off and actual landing). The advantages of using Monte Carlo Simulation for this application are that the processing time is very short, it makes use of the actual stochastic distributions, and offers a reliable approximation of the actual performance.

5.1 Implementation Of Monte Carlo Simulation

The basic program flow of the Monte Carlo Simulation implemented is shown below in Figure 5-1. The programmed Monte Carlo Simulation is supplied with the following inputs: a complete flight schedule, a saved random number stream, departure delay distribution, flight delay distribution, and the following user entered parameters: the chance of a departure delay occurring and the chance of a flight delay occurring. The simulation makes use of the saved random number streams so that each schedule evaluated can use the same set of random numbers, the advantage of this is that each schedule gets an equal, and fair, selection of delays such that the results obtained from the simulation are comparable. Each schedule Monte Carlo Simulation consists of multiple “runs”; a single run is equivalent to one actual day in operation for the flight schedule (a single simulation pass through the schedule). Multiple runs are used to gain a more accurate approximation of the actual performance of the schedule. The results from all the runs are averaged to create the final Monte Carlo Simulation results.

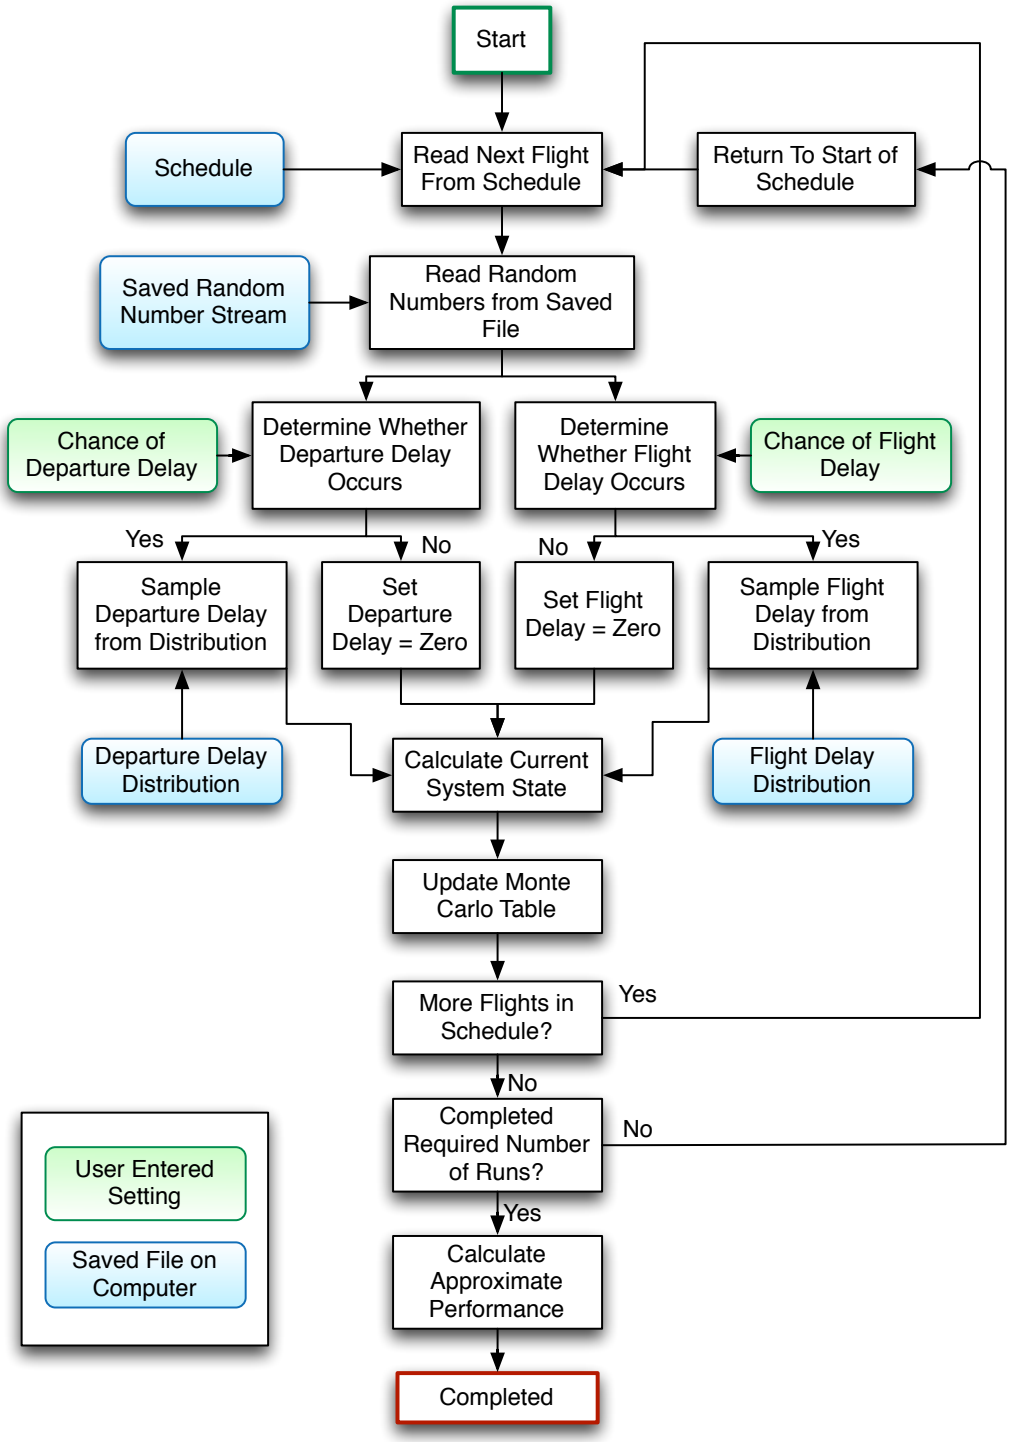


Figure 5-1: Monte Carlo Simulation Flow Diagram

A single pass simulation of a schedule is described in this paragraph. The flight schedule is a selection of flights to be operated during a particular day. The first step of the simulation process is to divide this schedule into the selection of flights for each particular aircraft in the fleet. The Monte Carlo Simulation is conducted separately for each aircraft and each flight is simulated in order. The

Monte Carlo Simulation is captured within a table which records and calculates the various measures which are of interest. The table headings are described in Table 5-1 below. A simulation of each flight in the schedule is conducted in order and a row of the table is completed for each flight, the values entered in the table are affected by the previous state of the system (the previous row in the table). The process is repeated for each aircraft in the fleet. The random number stream is used for the sampling of the distributions and for the determination of whether a delay will occur in the first place. Each random number is a Real number between 0 and 1. The process of determining whether a delay will occur using a random number is conducted as follows: a random number is drawn from the random number stream (for example: 0.5193), this random number is compared against the user-entered parameter which describes the chance of a delay (for example: 0.424); if the random number is larger or equal to the delay chance parameter then a delay is chosen to occur (for example: 0.5193 is greater than 0.424 thus a delay will occur). The process of sampling from a distribution is similar and is conducted as follows: a random number is drawn from the random number stream (for example: 0.723) and used to sample from the Inverse of the Cumulative Distribution Function (CDF) of the distribution to obtain a value representing the delay (for example: 72.5 minutes). Using Figure 5-2 as a reference, the random number represents a value on the Y-axis. Draw a horizontal line across the figure until it meets the CDF function line, at this point draw a vertical line down towards the X-Axis and where it crosses the X-axis is the representative delay.

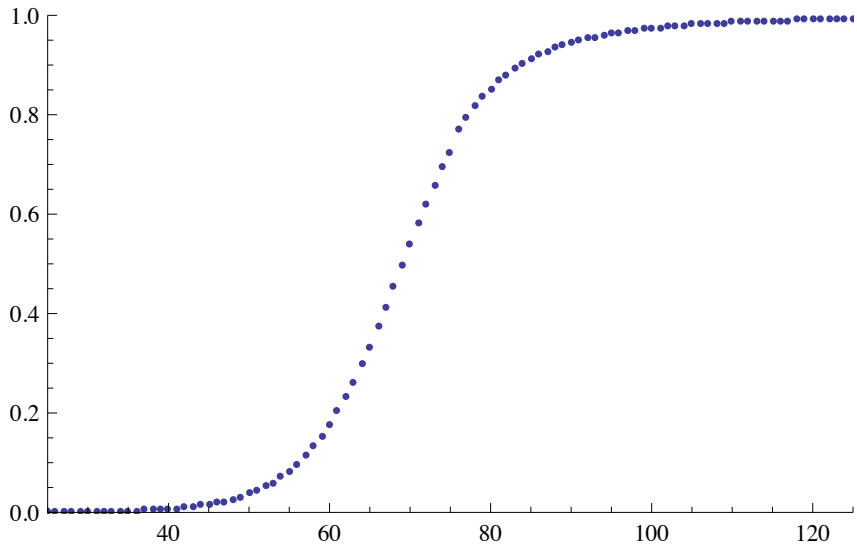


Figure 5-2: Example Cumulative Distribution Function

Once the Monte Carlo Simulation has completed the specified number of runs then the final step of the process is to collect the information regarding the performance of the schedule. The measures collected are an average of the respective measures for each run. The measures collected are:

- Number of Flights,
- Number of Departure Delays,
- Number of Arrival Delays,
- Number of Propagated Delays (delays which affect the following scheduled flight),
- Number of times the operated flights exceeded the length of the scheduled flight day,
- Number of time the operated flights exceeded the length of an entire day (24 hours),
- Average Departure Delay,
- Average Arrival Delay,
- Average Propagated Delay Time,
- Average Lateness of Departures,
- Average Lateness of Arrivals,
- Percentage of Flights Affected by Propagated Delays,
- Delay Recovery Percentage (Total Arrival Delay Time as a percentage of Total Lateness of Arrivals), Represents the ability of the schedule to absorb delays.

Table 5-1: Description Of Monte Carlo Simulation Table

	Heading	**Description (Calculation)**
1	Scheduled Departure Time	*Read from the Flight Schedule*
2	Propagated Delay	*Previous Flight Ready Time (17) subtracted from Scheduled Departure Time (1). If negative then set to zero.*
3	Earliest Possible Departure Time	*Scheduled Departure Time (1) + Propagated Delay (2)*
4	Random Number # 1	*Drawn from Random Number Stream*
5	Departure Delay?	*YES, if (4) >= Departure Delay Chance parameter; else NO.*
6	Random Number # 2	*Drawn from Random Number Stream*
7	Departure Delay Time	*If Departure Delay? (5) is YES, then sample from distribution;else set to zero*
8	Actual Departure Time	*Earliest Possible Departure Time (3) + Departure Delay Time (7)*
9	Scheduled Arrival Time	*Read from Schedule*
10	Expected Arrival Time	*Actual Departure Time (8) + Flight Time*
11	Random Number # 3	*Drawn from Random Number Stream*

12	Flight Delay?	*YES, if (11) >= Flight Delay Chance parameter; else NO.*
13	Random Number # 4	*Drawn from Random Number Stream*
14	Flight Delay Time	*If Flight Delay? (12) is YES, then sample from distribution;else set to zero. Flight Delay may be a negative value representing a faster than expected flight*
15	Actual Arrival Time	*Expected Arrival Time (10) + Flight Delay Time (14)*
16	Turn Around Time	*Read from the Flight Schedule. Typically 30 minutes.*
17	Ready Time	*Actual Arrival Time (15) + Turn Around Time (16)*
18	Departure Time Difference	*Actual Departure Time (8) - Scheduled Departure Time (1)*
19	Arrival Time Difference	*Actual Arrival Time (15) - Scheduled Arrival Time (9)*

The use of Monte Carlo Simulation allows for the comparison of the performance of various generated schedules in an approximation of the real world conditions.

6 TEST SITUATION DATA SETS

A variety of test situations were generated for the purpose of running and evaluating the operations research models developed for the project. A base test situation was created which was strongly based on a local low-cost airline operating in South Africa; the base situation was then modified to form additional test situations which would have a different degree of schedule density (amount of flying time per aircraft schedule).

6.1 Base Situation

The base situation strongly resembles the operating parameters of a low-cost airline operating in South Africa. There are eight airports served in the base situation and they are listed in Table 6-1 below. Presented in Table 6-2 are the projected passenger demand quantities for each route, the number given in the table refers to the number of people wishing to fly that route per day. The code "NF" refers to the situation were no flight is offered for that particular route.

Table 6-1: List Of Airports Served

O.R. Tambo International Airport, Johannesburg	JNB
Cape Town International Airport, Cape Town	CPT
King Shaka International Airport, Durban	DUR
Bloemfontein International Airport, Bloemfontein	BFN
Port Elizabeth International Airport, Port Elizabeth	PLZ
East London Airport, East London	ELS
George Airport, George	GRJ
Nelspruit Airport, Nelspruit	NLP

Table 6-2: Passenger Demand Table

		Destination							
	Airport	*BFN*	*CPT*	*DUR*	*ELS*	*GRJ*	*JNB*	*NLP*	*PLZ*
Origin	*BFN*		300	200	NF	NF	450	NF	NF
	CPT	275		450	NF	NF	650	NF	120
	DUR	200	450		NF	NF	350	NF	100
	ELS	NF	NF	NF		NF	95	NF	NF
	GRJ	NF	NF	NF	NF		105	NF	NF
	JNB	470	650	350	95	105		150	250
	NLP	NF	NF	NF	NF	NF	150		NF
	PLZ	NF	130	100	NF	NF	250	NF	

There are additional constraints placed upon the flights between Nelspruit and Johannesburg; the flight from Johannesburg to Nelspruit is constrained to only depart between 8:30 AM and 11:00 AM, whilst the flight from Nelspruit to Johannesburg is constrained to only depart between 9:45 AM and 1:30 PM. These additional time constraints on these flights are included to simulate the real-world case when a flight's passenger demand is strongly affected by its operation times and in some cases the flight is only feasible from a profit perspective if the flight occurs in a certain time frame. It is also beneficial to have flights between major business cities occur early in the morning and late in the afternoon; this is to meet the demand of the business traveler. However, it is undesirable to limit the flight operation times for those routes to only the times which are convenient to the business traveler thus a different approach is taken in handling this situation. The approach is to incentivize flights occurring during the peak business traveler period but still allow flights to occur at any time during the schedule. The incentive given is a 5% increase in profit for the applicable flights which will affect the solver's optimal solution such that those flight times are preferentially chosen. If additional flights between those cities are required then their flights at another time of the day may also be selected. The flights between Cape Town & Durban, Cape Town & Johannesburg, and Durban & Johannesburg have preferred departure times of 6:00 AM till 7:45 AM and 6:00 PM till 8:30 PM. The flight network of the base situation is shown in Figure 6-1. The size of the airport node is relative to the amount of passengers served at that airport for the situation and the width of the flight arcs is proportional to the passenger demand for that particular route.

The aircraft fleet being operated in the base situation is shown in Table 6-3. The selection of aircraft and their costs are based upon a local low-cost airline operator. The passenger capacity, cruising speed, and range are the performance specifications quoted by the various aircraft manufacturers. The hourly flight cost was calculated using reported operating costs for those particular aircraft. These costs were obtained from detailed information retrieved from the Bureau of Transportation Statistics in the USA. The operating costs include fuel costs, maintenance costs, depreciation, repairs, crew costs, and taxes among many other factors, These costs were averaged across all the aircraft of that specification and divided by the amount of flight time. The costs may not be accurate for the South African environment but the relative costs

between aircrafts types should be accurate and that is the measure by which the optimal solution is found.

Table 6-3: Aircraft Fleet Specifications

Aircraft	# Active	Passenger Capacity	Cruising Speed (km/h)	Range (km)	Hourly Flight Cost
Boeing 737-400	5	168	780	4204	R35 497,89
Boeing 737-800	3	189	828	5665	R30 729,04
McDonnell Douglass MD-81	1	172	811	2910	R29 489,77

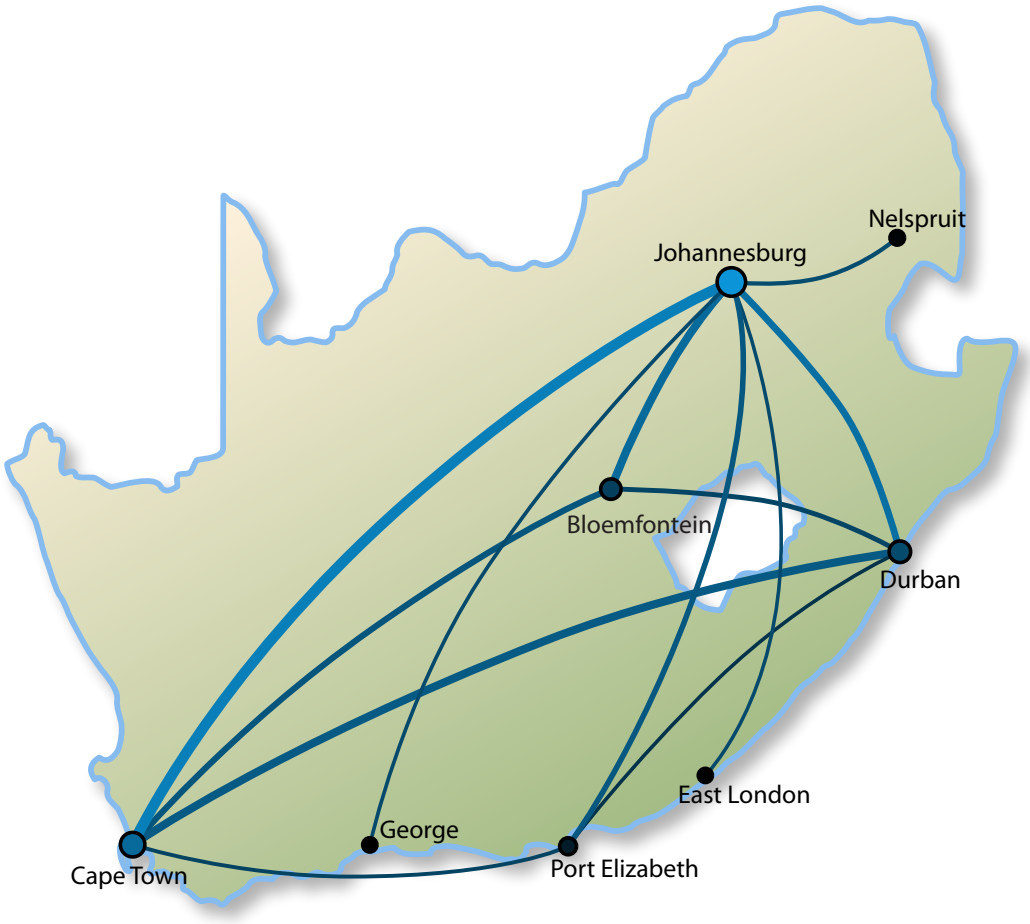


Figure 6-1: Flight Network For Test Situations

The additional parameters for the base situation are presented in Table 6-4 below. The turn-around time is the amount of time required after the aircraft has landed before it is available and ready for its next flight. Break-even passenger load percentage is used to calculate the seat prices for each flight; the cost of the flight is first calculated and then seat price is set such that if the passenger load

on that flight is equal to the break-even percentage then the revenue will be equal to the cost for that flight. Any more passengers over the break-even percentage will result in a profit for that flight whilst any passengers less than the break-even point will result in a loss for that flight. In such a case depending on the combination of the other flights, the flight would probably not be chosen for the schedule. The frames per hour parameter sets the amount of discrete time points used for the time-space network in the model. A higher number would result in a larger time-space network which would increase the size and complexity of the problem. This parameter controls the compromise between the accuracy of the model and its solving time. With the parameter set to four the time windows are of 15 minute length. This was found to be a good compromise, resulting in acceptable solving times without adversely affecting the solution accuracy. The schedule length defines the length of the operating day for the aircraft in the schedule and thus the schedule's operating day if all the aircraft schedules start at the same time. A schedule length of 16 hours allows enough time for maintenance operations which are typically less than eight hours in length. The schedule start time is used for defining the specific time periods which have additional constraints such as time-constrained flights and preferred flight times.

Table 6-4: Additional Parameters For The Base Situation

Turn-around Time	*30 minutes*
Break-Even Passenger Load	*65%*
Frames per Hour	*4*
Schedule Length	*16 hours*
Schedule Start Time	*6:00 AM*

6.2 Additional Test Situations

Additional test situations were generated such that the models could be tested over a range of different schedule densities (percentage of operating time per schedule). The parameters of the base situation was modified to form these additional test situation data sets. The parameters which were adjusted were the aircraft quantities for each fleet and the schedule length. A situation with a long schedule length and many additional aircraft (e.g. situation ATS-7) would offer the theoretical optimal solution because the time constraints would have very little effect on the solution whilst a situation with fewer aircraft and a shorter schedule length would create a densely packed schedule. Such a schedule would have very little idle slack time and thus would be highly susceptible to flight delays. The schedule's susceptibility to flight delays is highly influenced by the schedule's

density, this why the situations have been modified as described above. It is important to evaluate the stochastic models ability to handle delays with a variety of schedules having different degrees of susceptibility to flight delays. The additional test situations (ATS) are presented in Table 6-5, all the other parameters remain unchanged from the base situation.

Table 6-5: Additional Test Situations

Situation	Schedule Length (hours)	# 737-400s	# 737-800s	# MD-81s
Base	16	5	3	1
ATS-1	18	4	3	1
ATS-2	14	3	1	1
ATS-3	12	2	1	1
ATS-4	20	4	3	2
ATS-5	16	0	1	1
ATS-6	20	0	1	1
ATS-7	24	6	7	4

7 STOCHASTIC MODEL PARAMETERS

The developed stochastic programming model is a simplified deterministic S-CCP model. Its behavior is controlled by adjusting the CPP constraint tolerance factor (α). As explained in Section 4.4.2, the adjustment of α affects the probability that a flight delay will be large enough to affect additional flights in the schedule (a propagated delay). There is a compromise between reducing the probability of propagated flight delays and the operating profit of the schedule. In order to understand the relationship between the operating profit, the ability of the schedule to absorb delays, and the CCP constraint tolerance factor, the model was run on the test data sets using a variety of settings for α. The setting for the model will be specified in delay time (minutes). Table 7-1 presents the various parameter settings which were used for testing. Listed in the table are the equivalent statistical measure, the equivalent on-time percentage, the equivalent on-time percentage assuming that delays less than 15 minutes are not considered a delay, and the breakdown of the delay time into departure delay and flight delay. The CCP constraint tolerance percentage is equivalent to the on-time percentage for arrivals.

The naive model ignores the effect of flight delays and assumes that all flights will experience zero flight delays thus the delay time setting is set to zero. Since 60% of the flights within the delay distribution operate on-time, without a delay in arrival time, the equivalent CCP constraint tolerance factor is 60%. The naive model is parameter setting #1 (PRM-1). PRM-2 was the expected value (EV) method which makes use of the expected value of the distribution. In this case it is the statistical mean of the distribution. PRM-3 uses the statistical mean plus the standard deviation of the distributions to obtain the delay time. PRM-4 uses just the standard deviation of the distributions to obtain the delay time. PRM-5 uses the mean absolute deviation (MAD) of the distribution to obtain the delay time. PRM-6 uses the root of the standard deviation of the distributions to obtain the delay time. Parameters settings 7 & 8 (PRM-7 and PRM-8) do not use a statistical measure of the distribution to obtain the delay time but rather makes use of a desired on-time arrival performance percentage (95% and 90% respectively). Statistical measures are used as the basis for the majority of the parameter settings because the statistical measures will take more consideration of the properties of the distribution rather than just a specified level of

performance. The goal is that the best statistical measure could be confidently applied to many similar situations which have different distributions of data.

Table 7-1: Delay Time Parameter Settings

	Delay Time	Statistical Measure	Departure Delay Time	Flight Delay Time	Departure On-Time %	Departure On-Time % (15 Minute)	Arrival On-Time %	Arrival On-Time % (15 Minute)
PRM-1	0		0	0	57,6	80,8	60,4	79,7
PRM-2	5,7	Mean	11,3	-5,6	76,7	85,8	68,3	83,0
PRM-3	55,4	Mean + Standard Deviation	47	8,4	91,4	93,7	93,3	95,1
PRM-4	49,7	Standard Deviation	35,7	14	88,7	91,9	92,3	94,5
PRM-5	29	Mean Absolute Deviation	19,8	9,2	82,5	88,4	87,2	91,3
PRM-6	9,7	Root Standard Deviation	6,0	3,7	70,9	83,5	73,6	85,1
PRM-7	68						95,0	96,3
PRM-8	39						90,0	93,1

7.1 Model Assumptions

The assumptions for all the models are:

- Aircraft travel at a block speed equal to its cruising speed. This assumption simplifies the calculation of the flight times. An airline with detailed flight times could easily substitute the actual flight times into the model.
- Flight costs are independent of the number of passengers on board flight.
- An Airline will wish to use as many of their aircraft in their fleet as possible to reduce the schedule density without affecting the profit of the flight schedule. Thus they will spread flights among all the aircraft of the specified fleet type. This assumption is adjustable in the application. The alternative option is to use as few aircraft as possible.
- The Schedule to be created is for a single day. The day length parameter in the application can be adjusted to consider flight schedules which last longer than one day.

- The flight delay distribution is identical for each airport and fleet type. The application allows for different flight delay distributions (both departure and arrival) for each airport.
- Seat prices are directly related to flight costs. Actual airline seat pricing policies are complex and seats on the same aircraft can have different costs depending on factors such as: day of purchase, current load factor of aircraft (passenger density), and discounts.

8 OBSERVATIONS

8.1 Test Equipment Specifications

The models were solved and evaluated using the Gurobi solver (version 4.0.1) running on a Macintosh computer. The Gurobi solver was running on an academic license. The specifications of the computer used are given in Table 8-1 below. The Gurobi solver was chosen for its relative performance against the following alternatives: CPLEX 12.2 and SCIP 1.2.0. The Gurobi solver outperformed the alternative solvers with respect to the solving time of the models. It was five times faster in solving the base situation model. Upon further analysis of the solution log files produced by the solver, it was found that the optimal solution was obtained relatively quickly and that the majority of the solving time was spent trying to prove the optimality of the solution by tightening the bounds of the solution. The Gurobi solver allows for the adjustment of the high-level solution strategy for MIP models. By default, the Gurobi MIP solver strikes a balance between finding new feasible solutions and proving that the current solution is optimal. The Gurobi solver can be made to focus more attention on proving optimality, which is useful when the optimal solution is being found quickly. This was the case with the model in this project, thus the focus of the MIP solver was adjusted. The parameter in question is MIPFocus and it was set to a value of 2. Another parameter which was changed was the aggressiveness of the pre-solver; the aggressiveness was increased. The parameter Presolve was set to a value of 2. The adjustment of parameters for the Gurobi solver resulted in a significant decrease of solving time from 325 seconds to 131 seconds for the naive base situation model. Finally, a solving time limit was imposed of 1200 seconds. The solver time limit was imposed to speed up the solving of the models. The optimal solution is generally found in less than 100 seconds and the rest of the solving time is spent on proving optimality. In some cases this can take over 1200 seconds. The time limit is used to stop the solving process early.

Table 8-1: Computer Specifications

Computer Brand	*Apple*
Model	*MacBook*
Processor	*Intel Core 2 Duo*
Processor Speed	*2,2 GHz*
Number of Processors	*1*
Number of Cores	*2*
L2 Cache	*4 MB*
RAM Size	*4 GB*
Memory Type	*DDR2 SDRAM*
Memory Speed	*667 MHz*
Bus Speed	*800 MHz*
Operating System	*Mac OS X 10.6.6 (Snow Leopard)*
Architecture	*64-bit*

8.2 Observations

8.2.1 Gurobi solver solution process

Table 8-2 offers a sample of the observations whilst the complete observations table is included in Appendix A, Table A-1. 'Model Code' refers to the combination of test situation data set used and the model parameters used to generate that specific model. These codes correspond to the tables given in Sections 6 and 7. The size of the models are reported for both the original model (as inputted into the solver) and the resultant pre-solved, simplified, model. The size of the model is reported by the number of rows, columns, and nonzero variables present in the model. The optimality gap is the maximum possible percentage difference between the best solution found and the current best known theoretical optimal solution (which may not exist, but hasn't been proven not to exist yet). The solver defaults to stopping the solving process if the optimality gap is less than 0,01%.

Table 8-2: Sample Solver Observations

Model Code	Solving Time (s)	Original			Pre-solved			Optimality Gap (%)
		Rows	Columns	NonZero	Rows	Columns	NonZero	
BASE PRM-1	66,83	1563	5445	14907	1211	5096	14286	0
BASE PRM-2	1200	1563	5425	14847	1193	5058	14135	0,5587
BASE PRM-3	213,89	1563	5209	14199	1062	4709	13174	0,0042
BASE PRM-4	40,52	1563	5229	14259	1070	4737	13227	0
BASE PRM-5	65,69	1563	5317	14523	1139	4895	13623	0,0041
BASE PRM-6	1200	1563	5415	14817	1191	5046	14103	0,6752
BASE PRM-7	115,18	1563	5153	14031	1020	4610	12876	0
BASE PRM-8	276,49	1563	5283	14421	1106	4828	13446	0,0094
ATS-1 PRM-1	11,87	1755	6165	16875	1379	5792	16206	0
ATS-1 PRM-2	14,86	1755	6145	16815	1361	5754	16055	0
ATS-1 PRM-3	30,09	1755	5929	16167	1230	5405	15094	0
ATS-1 PRM-4	64,73	1755	5949	16227	1238	5433	15147	0,0054
ATS-1 PRM-5	123,22	1755	6037	16491	1309	5593	15547	0,0057
ATS-1 PRM-6	12,34	1755	6135	16785	1359	5742	16023	0
ATS-1 PRM-7	15,32	1755	5873	15999	1188	5306	14796	0
ATS-1 PRM-8	74,51	1755	6003	16389	1275	5525	15368	0

8.2.2 Schedule generation

Once the solver has produced a solution (list of variables and their values) then the next stage of the process is to generate a schedule using that solution. The generated flight schedule for the BASE PRM-1 model is shown in Appendix C. Table 8-3 offers a sample of the observations whilst the complete observations table is available in Appendix A, Table A-2. ‘Aircraft Used’, ‘Unused Aircraft’, and the number of ‘Flights’ within the schedule are listed. The schedule is spread out, if possible, amongst all the aircraft of the same fleet type to reduce the schedule density as much as possible. Therefore the amount of unused aircraft generally corresponds to the amount of aircraft in a specific fleet type which was not used for the schedule. The schedule density measure is the percentage of time for

which the aircraft are either flying or in the process of “turning-around” (i.e. all activities besides idle time) compared with the total schedule time for each aircraft. This total schedule time for each aircraft excludes any idle time before the first flight of the day and after the last flight of the day for each aircraft. The expected net profit for operating the schedule is given. The load factor is an airline industry measure which records the passenger density of the flights. It is the number of passengers on the flight divided by the number of seats available on that flight (RPK / ASK). A value of 1 for the load factor means that every seat in the aircraft has a passenger allocated to it.

Table 8-3: Schedule Observations

Model Code	Aircraft Used	Unused Aircraft	Flights	Schedule Density (%)	Profit (R)	Load Factor
BASE PRM-1	3	4	26	81,66	R419 766,63	0,99692
BASE PRM-2	3	4	26	81,22	R419 086,43	0,996898
BASE PRM-3	7	0	24	55,56	R427 816,07	1
BASE PRM-4	5	2	24	52,42	R425 757,64	1
BASE PRM-5	5	2	24	78,60	R416 349,17	1
BASE PRM-6	3	4	26	78,84	R419 086,43	0,996898
BASE PRM-7	7	0	24	54,90	R431 418,48	1
BASE PRM-8	6	1	26	61,19	R430 108,76	0,996856
ATS-1 PRM-1	4	4	28	58,99	R417 730,13	0,974925
ATS-1 PRM-2	4	4	28	60,06	R417 985,63	0,974925
ATS-1 PRM-3	4	4	26	55,96	R420 073,51	0,99693
ATS-1 PRM-4	4	4	26	56,00	R420 073,51	0,99693
ATS-1 PRM-5	4	4	28	62,85	R417 857,54	0,974925
ATS-1 PRM-6	4	4	28	62,33	R417 730,13	0,974925
ATS-1 PRM-7	4	4	26	59,46	R419 766,63	0,99692
ATS-1 PRM-8	4	4	26	60,07	R416 887,58	0,994179

8.2.3 Monte Carlo Simulation

A Monte Carlo Simulation was conducted, described in Section 5, on the resultant schedule for each model. The observations of the simulation are given in Appendix A, Table A-3, whilst a sample of the observations is presented below in Table 8-4. ‘Percentage of Propagated Delayed Flights’ is the percentage of the flights simulated which experienced a knock-on delay from the flight which preceded it in the schedule. ‘Delay Recovery’ is a measure of how well the schedule is able to recover from flight delays (i.e. not affecting the rest of the schedule). It is calculated by dividing the total time of flight delays experienced by the total time difference between scheduled and actual arrival times. ‘Average

Propagated Delay' is the average propagated delay experienced by a flight delayed due to previous flight delays. 'Total Average Propagated Delay' is the average propagated delay per flight when considering all the flights including the flights which don't experience that type of delay. 'Average Time Difference Departures' and 'Average Time Difference Arrivals' refer to the average time difference between the scheduled and actual departure or arrival times respectively. They are the average departure and arrival delays experienced by the passengers.

Table 8-4: Monte Carlo Simulation Observations

Model Code	Percentage of Propagated Delayed Flights	Delay Recovery %	Average Propagated Delay	Total Average Propagated Delay	Average Time Difference Departures	Average Time Difference Arrivals
BASE PRM-1	36,54	23,16	51,73	18,901	29,98	24,60
BASE PRM-2	23,46	36,52	42,78	10,037	21,19	15,81
BASE PRM-3	2,50	83,87	36,93	0,923	10,81	5,73
BASE PRM-4	3,54	76,68	43,86	1,553	11,74	6,66
BASE PRM-5	12,50	50,04	46,07	5,759	16,61	11,53
BASE PRM-6	20,96	38,43	44,13	9,250	20,40	15,03
BASE PRM-7	2,08	86,20	38,62	0,804	10,91	5,83
BASE PRM-8	9,62	58,34	37,99	3,653	14,14	8,77
ATS-1 PRM-1	7,86	67,11	43,00	3,379	15,01	10,27
ATS-1 PRM-2	11,07	58,29	46,13	5,107	16,98	12,24
ATS-1 PRM-3	5,00	72,84	37,56	1,878	12,29	6,91
ATS-1 PRM-4	3,85	78,97	31,84	1,225	11,20	5,82
ATS-1 PRM-5	10,89	59,91	43,98	4,790	16,68	11,95
ATS-1 PRM-6	8,93	63,58	44,55	3,977	15,66	10,92
ATS-1 PRM-7	4,42	77,19	35,48	1,569	12,26	6,88
ATS-1 PRM-8	5,58	72,16	35,64	1,988	12,52	7,14

9 RESULTS AND DISCUSSION

No correlation could be identified between the various measures of the models, such as number of row, number of columns, number of variables, parameter setting, or even test situation data set, and the time required to solve them. It was expected that the size of the model would have some degree of correlation to the solving time of the model. No such correlation was found and the solving time appeared to be random in nature. This apparent random solving time could be attributed to the solver's use of heuristic methods and pseudo-random numbers during the solving process. Therefore the solving time is dependent on how "lucky" the heuristic and random numbers were at finding good solutions. Since the Gurobi solver's source code isn't freely available and its techniques for solving models is secret and proprietary any further investigation into the solving process is not possible. However, it can be seen (by referring to Table A1 in Appendix A) that the solving time is generally less than 150 seconds (87,5% of the models solved in under 150 seconds) which demonstrates that for the test situations the model is solvable in a reasonable amount of time. Thus the model is suitable for the application to the low-cost airline operating in South Africa. The model was tested on a real scenario (JetBlue Airlines operating in the United States of America) and performed well, solving the model in 8,8 seconds. JetBlue operates a large fleet in relation to the flight schedule thus the scenario resulted in a low schedule density. A description of the JetBlue scenario and the observations are included in Appendix D.

The inter-relationships between the density of the schedule, actual departure and arrival delays, the ability of the schedule to recover from delays, and the average propagated delay time will have a direct effect on the success of an implementation of this model. Relationships between these measures is discussed below. The approach taken in achieving this project's objectives assumes that there is a correlation between the density of the flight schedule and the ability of the schedule to recover from delays. It is expected that as the density of the schedule decreases, the ability to recover from delays will increase. Figure 9-1 shows the scatter plot of delay recovery percentage versus schedule density percentage. The dashed line is the best fit polynomial trend-line. There is a relationship between schedule density and delay recovery and that delay recovery ability increases as the density of the schedule decreases. Schedule densities less than 35% appear to have almost zero propagated delays.

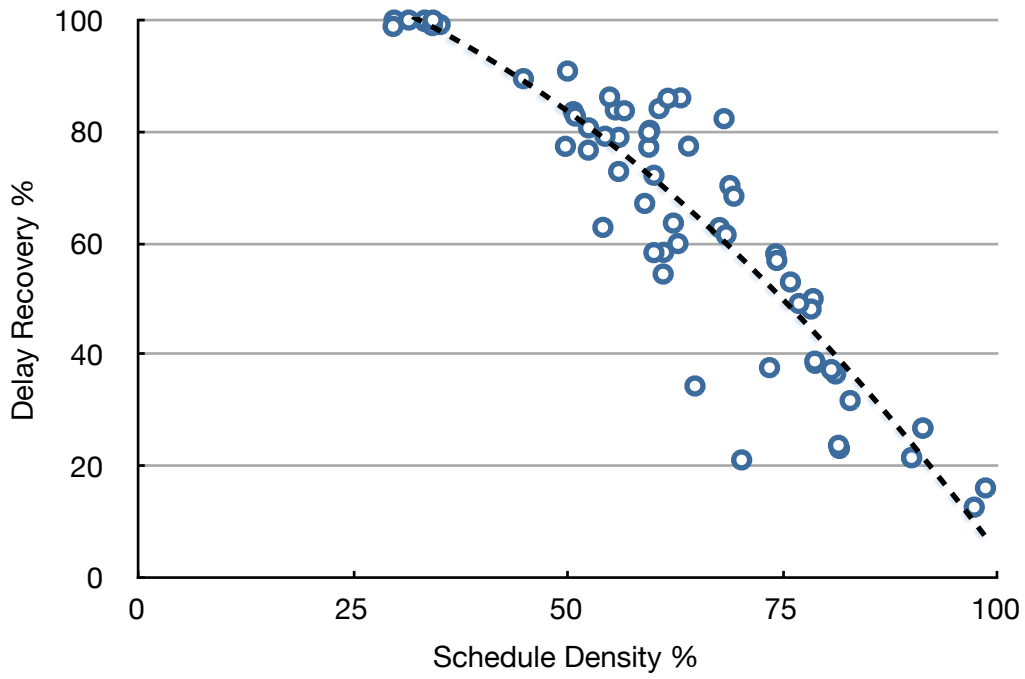


Figure 9-1: Scatter Plot of Delay Recovery Versus Schedule Density

Figure 9-2 shows a scatter plot of average propagated delay time per flight against delay recovery percentage. The dashed line is the best fit exponential trend-line. There is an exponential decrease in average propagated delay time as the delay recovery percentage increases.

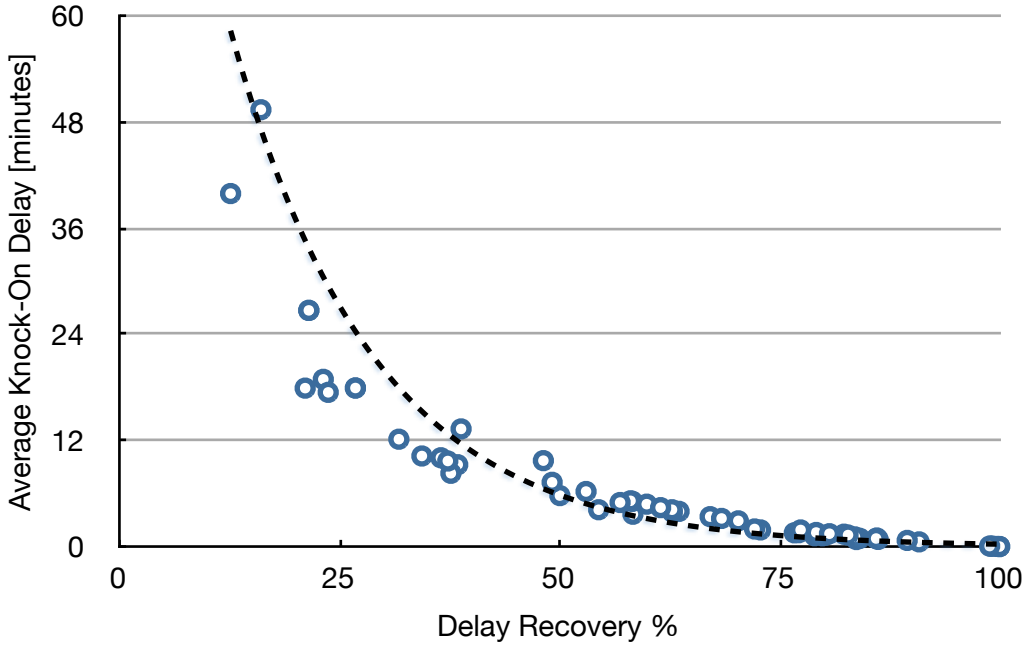


Figure 9-2: Scatter Plot of Average Propagated Delay Versus Delay Recovery

This is an expected result and shows that propagated delays are shorter in duration as the ability of the schedule to absorb delays increases.

Figure 9-3 shows the scatter plots of the average time difference between the scheduled and actual departure and arrival times against the density of the schedule. The dashed lines are the best fit exponential trend-line. There is an increase in delays as the schedule density increases. The scatter plots of the departure and arrival time differences are similar due to the obvious relationship between departure and arrival delays, this being that a flight which departs late will most likely also arrive late.

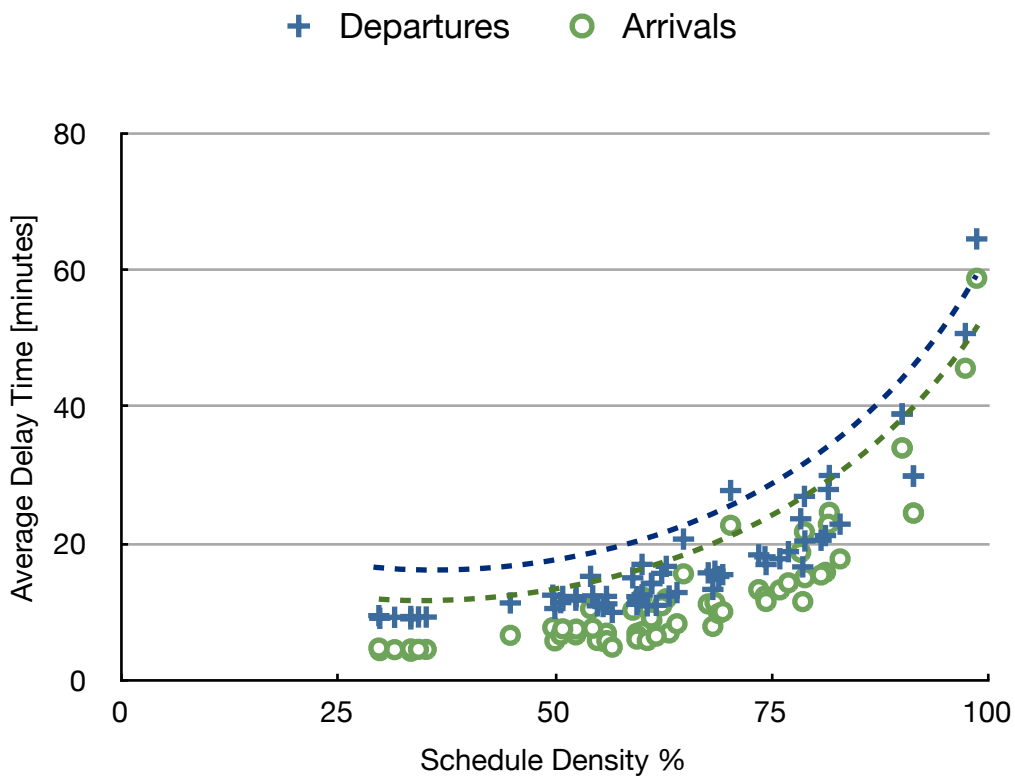


Figure 9-3: Scatter Plots of Average Delay Time Versus Schedule Density

It is apparent from the above figures that there is a relationship between schedule density and the effect which flight delays have on the schedule. The effect that the model parameters have on the schedule density, delay recovery percentage, and the profit of the generated schedules will be considered. Each figure (9-4, 9-5, 9-6) contains a plot for each test situation data set so that the relationships can be identified more clearly and the effect of the model parameters on the resultant performance measure for each situation can be identified.

Figure 9-4 shows schedule density versus the delay time setting used in the model parameters for each test situation. The dashed lines are the best fit polynomial trend-lines. The general trend of the plots is a decrease in schedule density as the delay time setting is increased. This is the expected behaviour and it allows us to adjust the delay time setting to influence the final schedule's density. It can be seen that the decrease in schedule density is more pronounced for certain test situations. This can be attributed to the characteristics of each test situation, namely the day length and the amount of aircraft available. The test situations which have large fleets of aircraft are able to spread the flights out more effectively and thus have a less dense schedule to begin with. The graph suggests that the use of stochastic time delay settings in the models is more useful and applicable to situations were the flight schedule is already dense with flights (more than 75%).

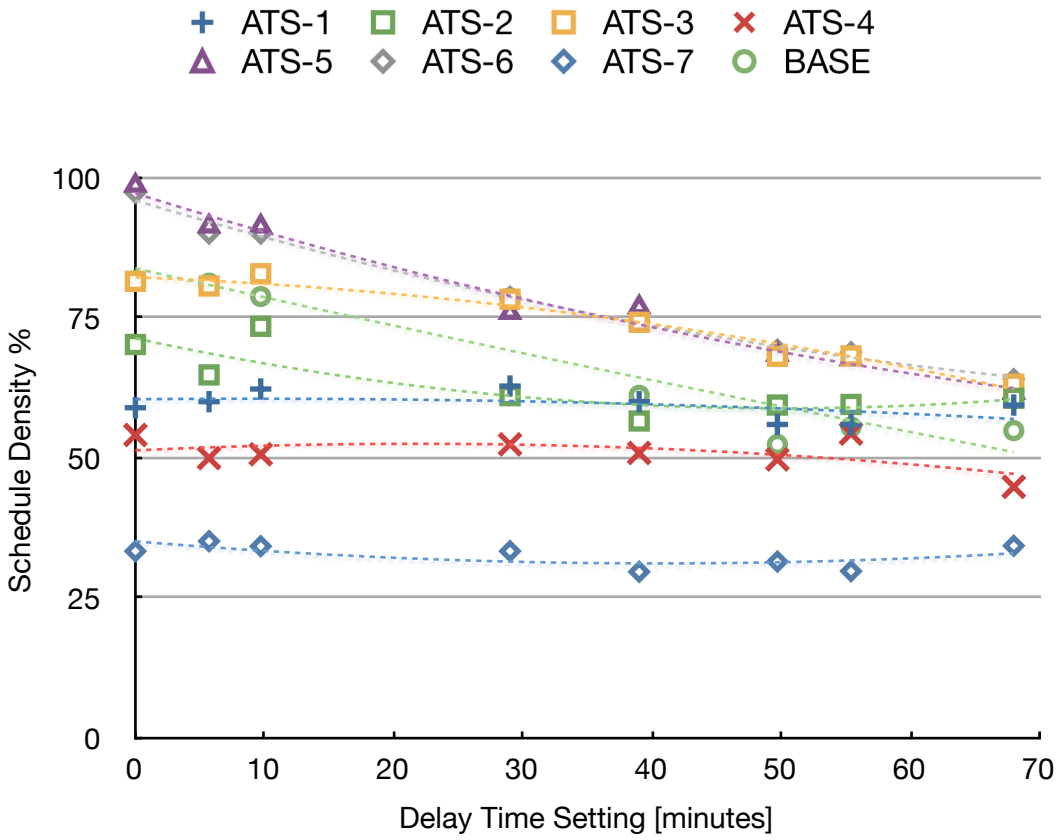


Figure 9-4: Plots of Schedule Density Versus Delay Time Setting For Each Test Situation

Figure 9-5 shows delay recovery percentage versus delay time setting for each test situation. The dashed lines are the best fit polynomial trend-lines. The plot shows a significant increase of delay recovery percentage as the delay time setting is increase for the majority of the test situations. This correlation validates

the approach taken in improving the flight schedules ability to handle delays by introducing the flight delay time setting into the models. The trend of some test situations gaining more benefit from the delay time setting in the model is repeated. The test situations which are not densely packed in the default naive model seem to be affected less by the delay time settings. This is the expected behaviour and suggests that airlines operating with dense flight schedules will gain more advantage from using stochastic programming techniques then those who do not.

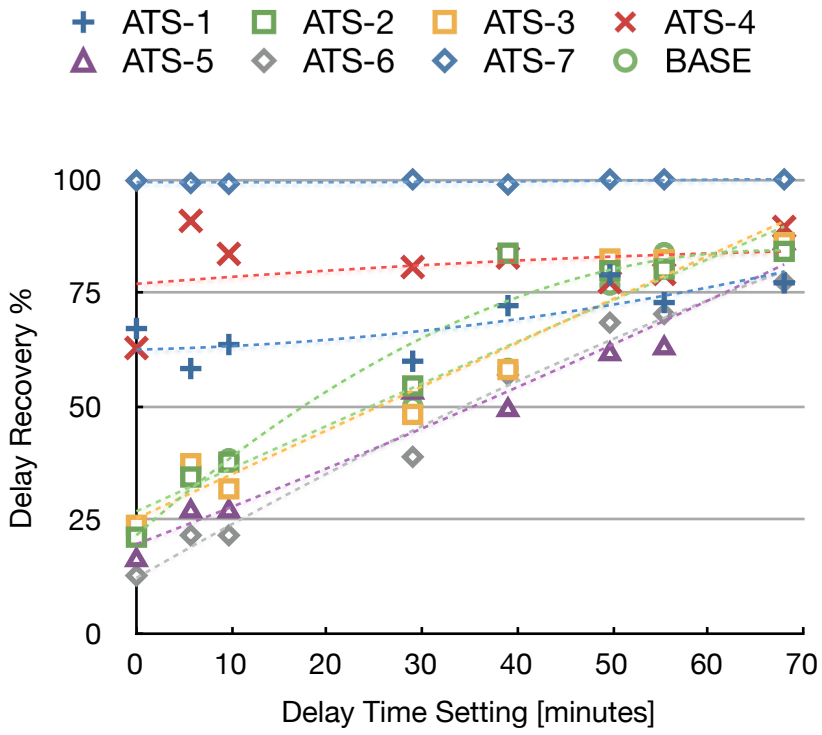


Figure 9-5: Plots of Delay Recovery Versus Delay Time Setting For Each Test Situation

It is important to identify the effect that the delay time setting has on the resultant schedule's profit. Figure 9-6 shows profit versus delay time setting for each of the test situations. The dashed lines are the best fit exponential trend-lines. The figure shows that for the test situations which didn't appear to gain any delay recovery performance as the delay time setting was increased (Figure 9-5) also didn't lose any profit by incorporating the delay time settings. Therefore airlines which are operating with flight schedules that are not dense are neither gaining any delay recovery performance nor losing any net profit. The delay time setting parameter has very little effect on these schedules. The delay time setting technique could be applied to these cases without consideration of the effect on the bottom line. However very little benefit will be gained with respect to delay

recovery. On the other hand, the test situations which gained the most improvement of their delay recovery performance also experienced the biggest decreases in schedule profit. Thus there exists a trade-off between the profit obtained from a schedule and the delay recovery performance of the schedule. The airline would need to resolve this trade-off by weighing the relative costs of each situation and deciding the best setting to use. The decision will include factors such as: profit, delay costs, and passenger goodwill.

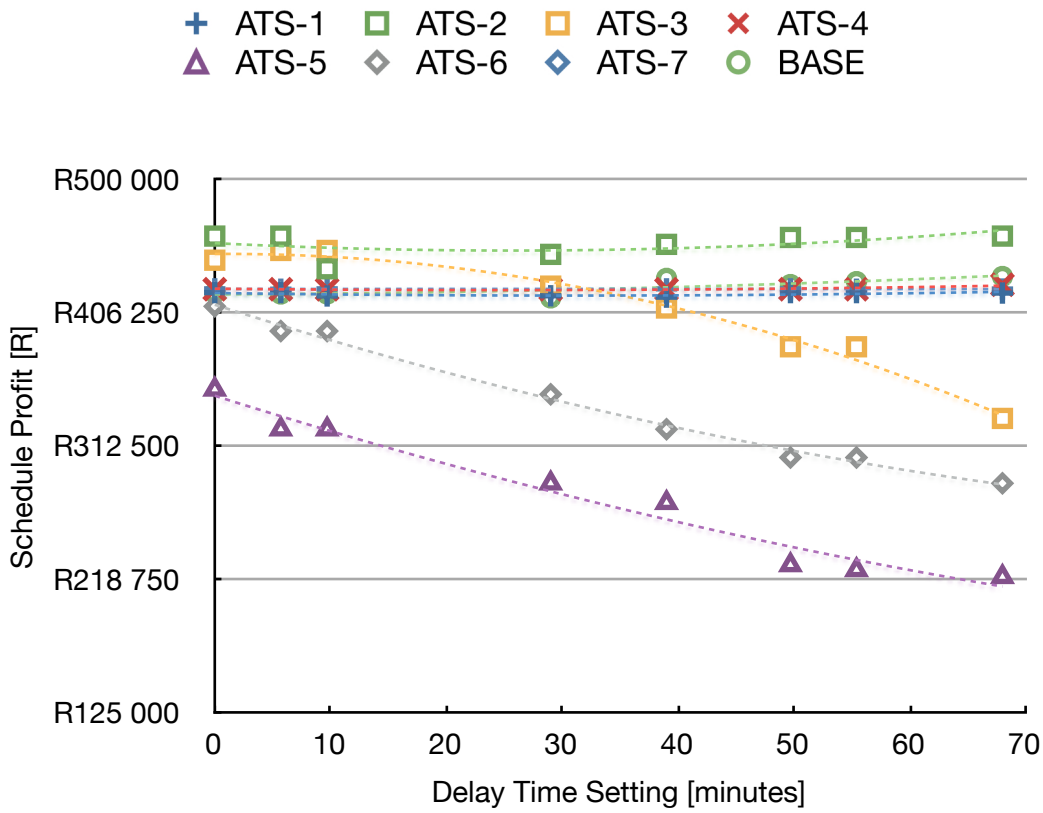


Figure 9-6: Plots of Schedule Profit Versus Delay Time Setting For Each Test Situation

The obvious question an airline would be interested in is 'which is the best delay time setting to use in the model?'. The answer is that it is dependent on the current situation at hand for which the schedule is being generated. Thus there is no single correct setting to use that will be the best for every case. To further illustrate the relationships shown in the figures above, the average effect of the delay time settings over the test situations is shown in Figures 9-7, 9-8, and 9-9. The bars in the charts are presented in order of increasing delay time setting.

Figure 9-7 shows a bar chart of the average profit of the schedules for each setting of the time delay parameter in the models. Average profit decreases as the delay time setting is increased.

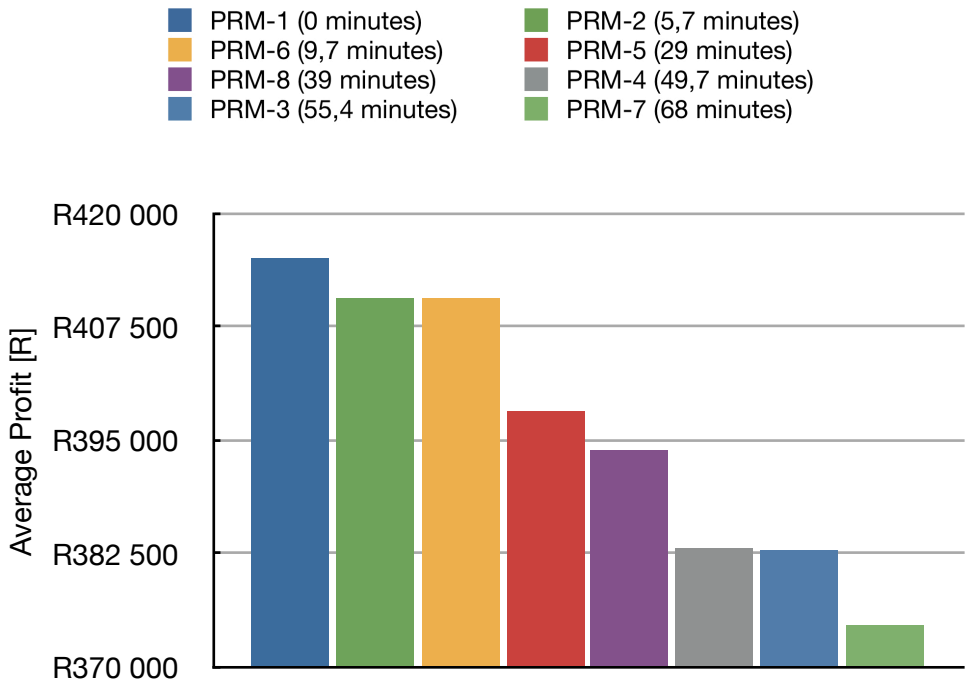


Figure 9-7: Bar Chart Of Average Profit Per Delay Time Setting

Figure 9-8 shows a bar chart of the average percentage of flights which experienced a propagated delay for each setting of the time delay parameter in the models. The percentage of propagated delays decreases as the delay time setting is increased.

Figure 9-9 shows a bar chart of the average delay recovery percentage for the schedules for each setting of the time delay parameter in the models. The delay recovery percentage increases with an increase in the delay time setting.

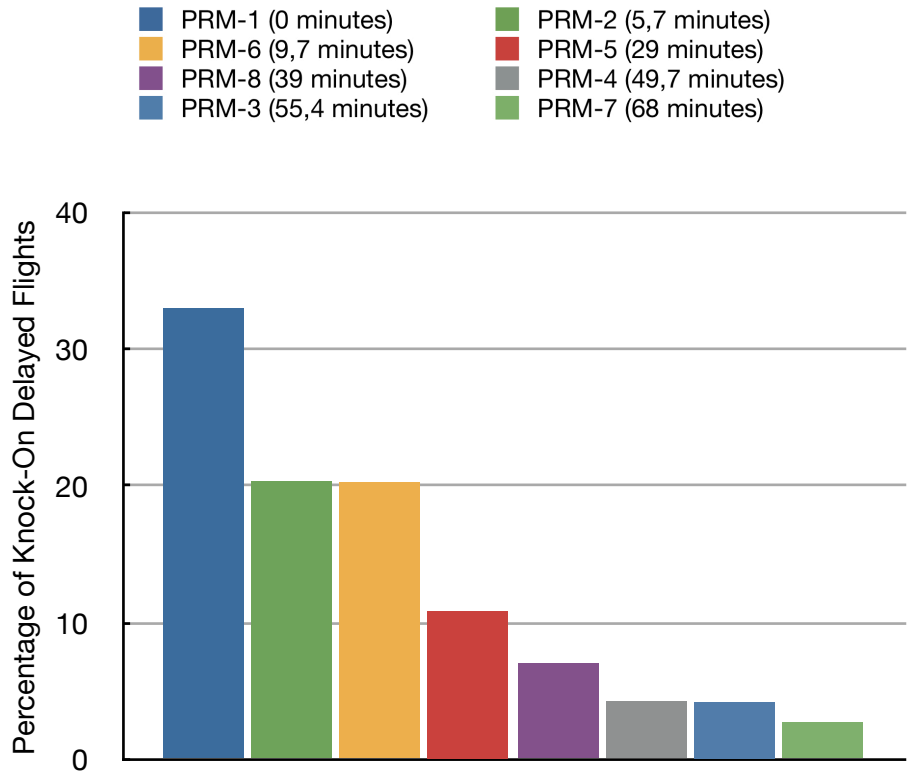


Figure 9-8: Bar Chart Of Average Percentage Of Propagated Delayed Flights Per Delay Time Setting

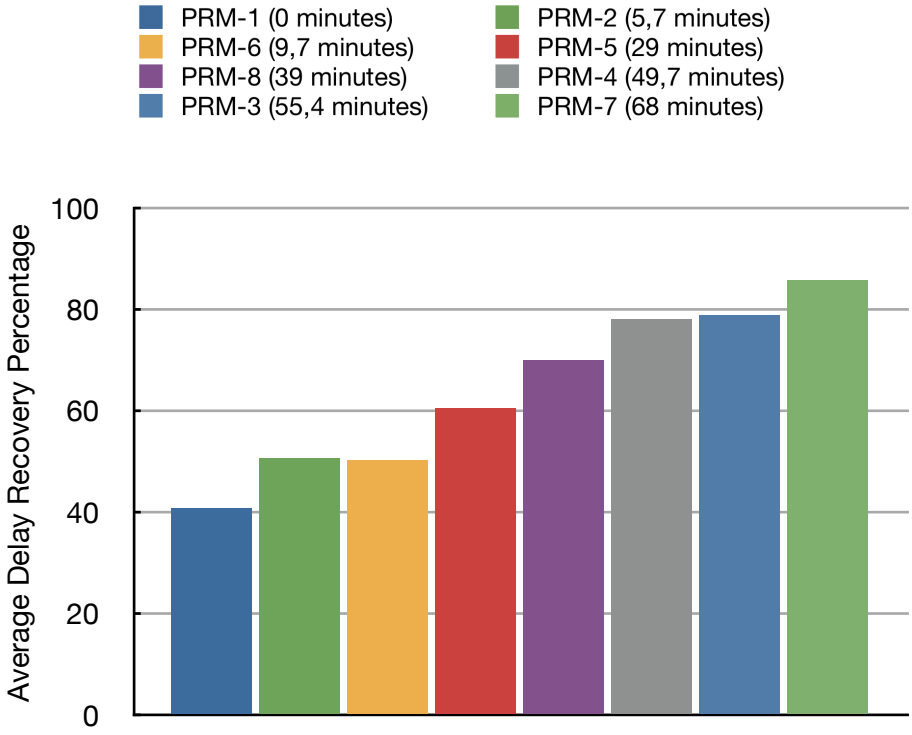


Figure 9-9: Bar Chart Of Average Delay Recovery Percentage Per Delay Time Setting

Assuming that the profit of the schedule and the delay recovery percentage of the schedule are equally important to the airline operator then it is possible to identify the best delay time setting for this case. Calculating the relative percentages of each measure for delay time setting (i.e largest profit gets 100% rating whilst smallest profit gets 0% rating for the profit portion), the sum of these measures can be used to get a measure of which setting offers the best compromise. Table 9-1 shows the relative percentage ratings for the measures, whilst Figure 9-10 shows the information in the form of a stacked bar chart, and Figure 9-11 is a bar chart of the compromise measure rating.

It can be seen in Figure 9-10 that the delay time setting of 39 minutes offers the best solution for the case when both the profit and delay recovery performance of the schedule is equally important to the airline. The 39 minute delay setting corresponds to the parameter specification that 90% of the flights should arrive on or before their scheduled arrival times.

Table 9-1: Relative Percentage Performance Measures

Delay Time Setting (minutes)	Relative Profit Percentage	Relative Delay Recovery Percentage	Net Performance Rating
0	100,00	0,00	100,00
5,7	88,87	21,80	110,67
9,7	89,00	21,10	110,10
29	58,14	44,08	102,21
39	47,70	64,88	112,59
49,7	20,81	83,00	103,81
55,4	20,29	84,75	105,04
68	0,00	100,00	100,00

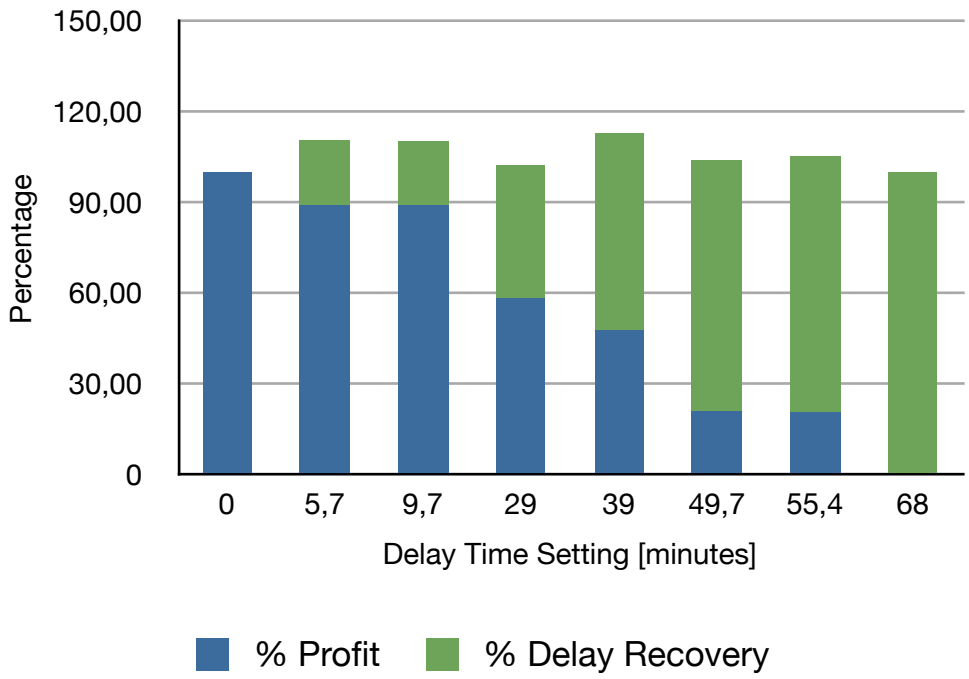


Figure 9-10: Stacked Bar Chart Of The Relative Percentages Of Profit And Delay Recovery For Each Delay Time Setting

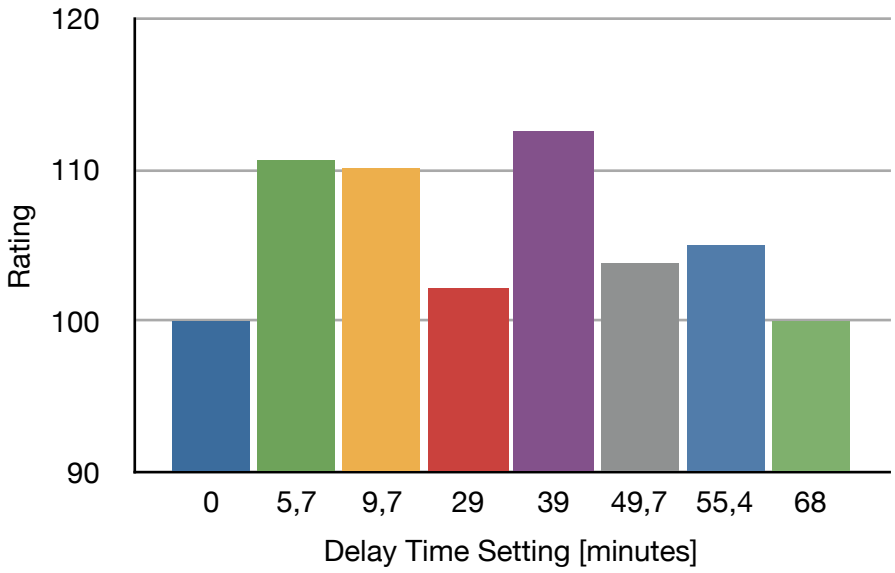


Figure 9-11: Bar Chart of Compromised Performance Rating For Each Delay Time Setting

10 SUMMARY CONCLUSIONS AND RECOMMENDATIONS

10.1 Summary Of Work Completed

An investigation into the appropriateness of stochastic programming techniques for reducing flight delays in airline schedules was carried out. The project commenced with a literature review of airline scheduling techniques, stochastic programming techniques, and the application of stochastic programming techniques in airline scheduling. Relevant real airline data required for the completion of the project was then obtained from various sources. The data allowed for the use of realistic situations to be modeled. A small test situation was created for the development and testing of the models to be created. Various models were created for use in this project, including both traditional mixed-integer models and stochastic models. A variety of test situations were created for the thorough testing and evaluation of the models. The models were solved using an external solver (Gurobi) and the solutions were used to generate flight schedules. These flight schedules were tested using Monte Carlo Simulation techniques to evaluate their performance in real-world conditions with regard to flight delays. Observations were recorded and the results analyzed to determine the success of the project.

10.2 Conclusions

1. The appropriateness of stochastic programming techniques for the airline scheduling problem was investigated. Stochastic programming techniques were able to improve the delay recovery performance of the schedule at the cost of decreasing the schedule's profit. Flight schedules which are more dense are affected more by the stochastic programming techniques.
2. The performance of various operations research techniques in formulating and solving the airline scheduling problem were evaluated. Observations were made of the computational time, optimality of the solution and the performance of the solution during simulation. The results show that the models were applicable for use by a local low-cost airline due to the fast solving times and quality of the schedules created.

10.3 Recommendations

- The use of stochastic programming techniques is recommended for cases where an airline's flight schedule has a high density of activity.

- The selection of parameters for stochastic programming models should be chosen using a well defined set of airline objectives such that the compromise between profit and flight delay recovery can be appropriately balanced to achieve the desired performance characteristics in the final schedule.

10.3.1 Recommendations for further research

- The S-CCP model could be modified to incorporate small independent time-windows for each flight duplicate within the scenario levels of the network structure. This would allow greater flexibility of the model at the expensive of model size and computational time.
- The scheduling model could be modified to incorporate additional steps of the scheduling process such as crew scheduling.
- Quantify the financial loses incurred due to flight delays. The financial lose could be used to select the most appropriate model parameters .

REFERENCES

AhmadBeygi, S., A. Cohn, et al. (2008). "Decreasing Airline Delay Propagation By Re-Allocating Scheduled Slack." Working Paper. University of Michigan.

Aouni, B. and D. La Torre (2010). "A generalized stochastic goal programming model." Applied Mathematics and Computation 215(12): 4347-4357.

Baoding, L. (1997). "Dependent-Chance Programming: A Class of Stochastic Optimization." Computers Math. Applic. 34(12): 89-104.

Barnhart, C., N. L. Boland, et al. (1998). "Flight String Models for Aircraft Fleeting and Routing." Transportation Science 32: 208-220.

Barros, C. P. and N. Peypoch (2009). "An evaluation of European airlines operational performance." International Journal of Production Economics 122(2): 525-533.

Benisch, M., A. Greenwald, et al. (2004). "A stochastic programming approach to scheduling in TAC SCM." Proceedings of the 5th ACM conference on Electronic commerce - EC '04: 152-152.

Bian, F., E. K. Burke, et al. (2005). Making airline schedules more robust. Berlin, Springer: 381-392.

Boland, N. L., L. W. Clarke, et al. (2000). "The Asymmetric Traveling Salesman Problem with Replishment Arcs." European Journal Of Operational Research 123: 408-427.

Bravo, M. and I. Gonzalez (2009). "Applying stochastic goal programming: A case study on water use planning." European Journal of Operational Research 196(3): 1123-1129.

Burke, E. K., Patrick, et al. (2010). "A multi-objective approach for robust airline scheduling." Computers & Operations Research 37(5): 822-832.

Charnes, A. and W. W. Cooper (1959). "Chance-constrained programming." Management Science 6(1): 73-79.

Chiraphadhanakul, V. (2010). Routing and Scheduling Models for Robust Allocation of Slack. Department of Civil and Environmental Engineering and the Sloan School of Management. Massachusetts, Massachusetts Institute of Technology.

Deb, K. (2001). Multi-objective optimization using evolutionary algorithms. New York, Wiley.

Deb, K. (2005). Multi-objective optimization. Berlin, Springer: 273-316.

Desaulniers G., Desrosiers J., Solomon M.M., and Soumis F (1994). Daily aircraft routing and scheduling. GERAD Technical Report, Montreal, Quebec, Canada.

Dobbyn, T. (2000). U.S. Prepares Plan to Cut Summer Flight Delays. Reuters Wire, Reuters.

Domenica, N. D., G. Mitra, et al. (2007). "Stochastic programming and scenario generation within a simulation framework : An information systems perspective." Decision Support Systems 42: 2197-2218.

Eurocontrol (2005). Delays to air transport in Europe. Digest Annual 2005.

Eurocontrol (2006). Eurocontrol medium-term forecast: IFR flight movements 2007-2013.

Floudas, C. A. and P. M. Pardalos (2005). "Stochastic scheduling." Encyclopedia of Optimization V: 367-372.

Grothklags, S. and U. Lorenz (2006). "Stochastic Airline Fleet Assignment is PSPACE-complete." Electronic Notes in Discrete Mathematics 25: 67-68.

Gu, Z., E. L. Johnson, et al. (1994). "Some Properties of the Fleet Assignment Problem." Operations Research Letters 15: 59-71.

Hane, C. A., C. Barnhart, et al. (1995). "The Fleet Assignment Problem: Solving a Large Integer Program." Mathematical Programming 70: 211-232.

Higle, J. and S. Wallace (2002). "Managing risk in the new power business - a sequel." IEEE Computer Applications in Power April: 12-19.

Jones, T. O. and W. E. Sasser (1995). "Why satisfied customers defect." Harvard Business Review 73(6): 88 - 99.

Kniker, T. (1998). Itererary-Based Airline Fleet Assignment, Massachusetts Institute of Technology. Ph.D.

Koch, T. (2004). Rapid Mathematical Programming. Technische Berlin. Berlin. PhD.

Lan, S., J. Clarke, et al. (2006). "Planning for robust airline operations: optimizing aircraft routings and flight departure times to minimize passenger disruptions." Transportation Science 40: 15-28.

Landa-Silva, J., E. K. Burke, et al. (2004). An introduction to multiobjective metaheuristics for scheduling and timetabling. Lecture notes in economics and mathematical systems. Berlin, Springer. 535: 91-129.

Lettovsky, L. (1997). Airline Operations Recovery: An Optimization Approach. Atlanta, Georgia Institute of Technology. Ph.D.

Levin A. (1971). Scheduling and Fleet Routing Models for Transportation Systems. Transportation Science 5: 232-255

Man, K. W. (2008). "A Multi-Stage Stochastic Linear Programming Model for Apparel Production Planning."

McDaniel, D. and M. Devine (1977). "A modified Benders' partitioning algorithm for mixed integer programming." Management Science 24(2): 312-319.

Mercier, A., J.-F. Cordeau, et al. (2005). "Computational study of Benders decompostion for the integrated aircraft routing and crew scheduling." Computers & Operations Research 32(6): 1451-1476.

Mirabi, M., S. M. T., et al. (2010). "Efficient stochastic hybrid heuristics for the multi-depot vehicle routing problem." Robotics and Computer-Integrated Manufacturing.

Möller, A., W. Römisch, et al. (2007). "Airline network revenue management by multistage stochastic programming." Computational Management Science 5(4): 355-377.

Ostrowski, P. L., T. V. O'Brien, et al. (1993). "Service quality and customer loyalty in the commerical airline industry." Journal of Travel Research 32(2): 16 - 24.

Papadakos, N. (2009). "Integrated airline scheduling." Computers & Operations Research 36: 176-195.

Parasuraman, A., V. A. Zeithaml, et al. (1988). "SERVQUAL: a multi-item scale for measuring consumer perceptions of the service quality." Journal of Retailing 64(1): 12 - 40.

Petkov, S. B. (1997). "Multi-period planning and scheduling of multi-product batch plants under demand uncertainty." Industry and Engineering Chemistry Research 36(11): 4864-4881.

Rosenberger, J. M., A. J. Schaefer, et al. (2002). "A Stochastic Model of Airline Operations." Transportation Science 36(4): 357-377.

Ruszczynski, A. and A. Shapiro (2003). "Stochastic Programming Models." Handbooks in OR & MS 10: 1-64.

Sastry, K., D. Goldberg, et al. (2005). Genetic Algorithms. Berlin, Springer: 97-126.

Schaefer, A. J., E. L. Johnson, et al. (2005). "Airline Crew Scheduling Under Uncertainty." Transportation Science 39(3): 340-348.

Shapiro, A. and A. Philpott (2007). "A Tutorial on Stochastic Programming." INTERNET. http://www.stoprog.org/SPTutorial/SPTutorial.html.

Shaprio, A. and S. Ahmed (2002). "The sample average approximation method for stochastic programs with integer recourse."

Sherali, H., E. Bish, et al. (2006). "Airline fleet assignment concepts, models, and algorithms." European Journal of Operational Research 172(1): 1-30.

Stojkovic, G. (2002). "An optimization model for a real-time flight scheduling problem." Transportation Research Part A: Policy and Practice 36(9): 779-788.

Wen, Y. and C. Hsu (2006). "Interactive multiobjective programming in airline network design for international airline code-share alliance." European Journal of Operational Research 174(1): 404-426.

Wu, C. L. (2006). "Improving Airline Robustness and Operational Reliability by Sequential Optimisation Algorithms." Network and Spatial Economics(6): 235-251.

Yan, S., C.-H. Tang, et al. (2008). "An airline scheduling model and solution algorithms under stochastic demands." European Journal Of Operational Research 190: 22-39.

Yen, J. W. and J. R. Birge (2006). "A Stochastic Programming Approach to the Airline Crew Scheduling Problem." 1-31.

(2000). Researchers Warn Airline Service Declining. Reuters Wire, Reuters.

APPENDIX A COMPLETE OBSERVATIONS

Table A-1: Complete Solver Observations

Model Code	Solving Time (s)	Original			Pre-solved			Optimality Gap (%)
		Rows	Columns	NonZero	Rows	Columns	NonZero	
BASE PRM-1	66,83	1563	5445	14907	1211	5096	14286	0
BASE PRM-2	1200	1563	5425	14847	1193	5058	14135	0,5587
BASE PRM-3	213,89	1563	5209	14199	1062	4709	13174	0,0042
BASE PRM-4	40,52	1563	5229	14259	1070	4737	13227	0
BASE PRM-5	65,69	1563	5317	14523	1139	4895	13623	0,0041
BASE PRM-6	1200	1563	5415	14817	1191	5046	14103	0,6752
BASE PRM-7	115,18	1563	5153	14031	1020	4610	12876	0
BASE PRM-8	276,49	1563	5283	14421	1106	4828	13446	0,0094
ATS-1 PRM-1	11,87	1755	6165	16875	1379	5792	16206	0
ATS-1 PRM-2	14,86	1755	6145	16815	1361	5754	16055	0
ATS-1 PRM-3	30,09	1755	5929	16167	1230	5405	15094	0
ATS-1 PRM-4	64,73	1755	5949	16227	1238	5433	15147	0,0054
ATS-1 PRM-5	123,22	1755	6037	16491	1309	5593	15547	0,0057
ATS-1 PRM-6	12,34	1755	6135	16785	1359	5742	16023	0
ATS-1 PRM-7	15,32	1755	5873	15999	1188	5306	14796	0
ATS-1 PRM-8	74,51	1755	6003	16389	1275	5525	15368	0
ATS-2 PRM-1	1200	1371	4725	12939	1043	4400	12366	0,1052
ATS-2 PRM-2	66,06	1371	4705	12879	1025	4362	12215	0
ATS-2 PRM-3	13,23	1371	4493	12243	894	4015	11237	0
ATS-2 PRM-4	22,47	1371	4509	12291	902	4039	11299	0,0061
ATS-2 PRM-5	1200	1371	4597	12555	970	4197	11697	0,1749
ATS-2 PRM-6	42,07	1371	4695	12849	1023	4350	12183	0
ATS-2 PRM-7	5,26	1371	4433	12063	852	3911	10944	0
ATS-2 PRM-8	1200	1371	4563	12453	937	4130	11520	0,1017

Model Code	Solving Time (s)	Original			Pre-solved			Optimality Gap (%)
		Rows	Columns	NonZero	Rows	Columns	NonZero	
ATS-3 PRM-1	16,18	1179	4005	10971	875	3704	10446	0
ATS-3 PRM-2	19,52	1179	3985	10911	857	3666	10295	0
ATS-3 PRM-3	12,86	1179	3785	10311	728	3333	9356	0
ATS-3 PRM-4	7,44	1179	3789	10323	734	3343	9379	0
ATS-3 PRM-5	17,66	1179	3877	10587	802	3501	9777	0
ATS-3 PRM-6	18,99	1179	3975	10881	885	3654	10263	0
ATS-3 PRM-7	54,30	1179	3713	10095	684	3215	9024	0
ATS-3 PRM-8	9,42	1179	3843	10485	769	3434	9600	0
ATS-4 PRM-1	5,14	1947	6885	18843	1547	6488	18126	0
ATS-4 PRM-2	11,48	1947	6865	18783	1529	6450	17975	0
ATS-4 PRM-3	8,68	1947	6653	18147	1398	6107	17055	0
ATS-4 PRM-4	10,03	1947	6669	18195	1406	6131	17075	0,0023
ATS-4 PRM-5	71,40	1947	6757	18459	1476	6289	17468	0
ATS-4 PRM-6	5,79	1947	6855	18753	1527	6483	17943	0
ATS-4 PRM-7	74,16	1947	6593	17967	1356	6005	16728	0
ATS-4 PRM-8	6,72	1947	6723	18357	1443	6222	17292	0
ATS-5 PRM-1	23,94	1050	3642	9958	818	3412	9559	0
ATS-5 PRM-2	79,53	1050	3626	9910	804	3382	9436	0
ATS-5 PRM-3	15,72	1050	3488	9496	716	3156	8790	0
ATS-5 PRM-4	15,70	1050	3496	9520	724	3172	8828	0
ATS-5 PRM-5	6,82	1050	3558	9706	770	3278	9102	0
ATS-5 PRM-6	100,52	1050	3622	9898	802	3376	9422	0
ATS-5 PRM-7	3,26	1050	3446	9370	688	3086	8608	0
ATS-5 PRM-8	1,96	1050	3534	9634	744	3230	8984	0,0021
ATS-6 PRM-1	317,76	1306	4602	12582	1042	4340	12119	0,0002
ATS-6 PRM-2	45,31	1306	4586	12534	1028	4310	11996	0

Model Code	Solving Time (s)	Original			Pre-solved			Optimality Gap (%)
		Rows	Columns	NonZero	Rows	Columns	NonZero	
ATS-6 PRM-3	5,87	1306	4448	12120	940	4084	11350	0
ATS-6 PRM-4	5,30	1306	4456	12144	948	4100	11388	0
ATS-6 PRM-5	35,01	1306	4518	12330	998	4212	11678	0
ATS-6 PRM-6	34,68	1306	4582	12522	1026	4304	11982	0
ATS-6 PRM-7	4,36	1306	4406	11994	912	4014	11168	0
ATS-6 PRM-8	40,09	1306	4494	12258	970	4160	11548	0
ATS-7 PRM-1	9,77	2331	8325	22779	1883	7880	21996	0
ATS-7 PRM-2	7,14	2331	8305	22719	1865	7842	21815	0
ATS-7 PRM-3	5,20	2331	8093	22083	1734	7499	20895	0
ATS-7 PRM-4	11,26	2331	8109	22131	1742	7523	20915	0
ATS-7 PRM-5	7,06	2331	8197	22395	1812	7681	21309	0
ATS-7 PRM-6	6,61	2331	8295	22689	1863	7830	21783	0
ATS-7 PRM-7	4,24	2331	8033	21903	1692	7397	20568	0
ATS-7 PRM-8	4,58	2331	8163	22293	1779	7614	21132	0

Table A-2: Complete Schedule Observations

Model Code	Aircraft Used	Unused Aircraft	Flights	Schedule Density (%)	Profit (R)	Load Factor
BASE PRM-1	3	4	26	81,66	R419 766,63	0,99692
BASE PRM-2	3	4	26	81,22	R419 086,43	0,996898
BASE PRM-3	7	0	24	55,56	R427 816,07	1
BASE PRM-4	5	2	24	52,42	R425 757,64	1
BASE PRM-5	5	2	24	78,60	R416 349,17	1
BASE PRM-6	3	4	26	78,84	R419 086,43	0,996898
BASE PRM-7	7	0	24	54,90	R431 418,48	1
BASE PRM-8	6	1	26	61,19	R430 108,76	0,996856
ATS-1 PRM-1	4	4	28	58,99	R417 730,13	0,974925

Model Code	Aircraft Used	Unused Aircraft	Flights	Schedule Density (%)	Profit (R)	Load Factor
ATS-1 PRM-2	4	4	28	60,06	R417 985,63	0,974925
ATS-1 PRM-3	4	4	26	55,96	R420 073,51	0,99693
ATS-1 PRM-4	4	4	26	56,00	R420 073,51	0,99693
ATS-1 PRM-5	4	4	28	62,85	R417 857,54	0,974925
ATS-1 PRM-6	4	4	28	62,33	R417 730,13	0,974925
ATS-1 PRM-7	4	4	26	59,46	R419 766,63	0,99692
ATS-1 PRM-8	4	4	26	60,07	R416 887,58	0,994179
ATS-2 PRM-1	5	0	24	70,28	R433 290,59	0,999045
ATS-2 PRM-2	5	0	24	64,84	R436 431,31	0,999041
ATS-2 PRM-3	5	0	24	59,57	R458 858,90	1
ATS-2 PRM-4	5	0	24	59,47	R458 858,90	1
ATS-2 PRM-5	5	0	24	61,14	R446 957,71	0,99696
ATS-2 PRM-6	5	0	24	73,52	R436 834,78	0,999041
ATS-2 PRM-7	5	0	24	60,66	R459 814,47	1
ATS-2 PRM-8	5	0	24	56,62	R454 002,25	0,999029
ATS-3 PRM-1	4	0	24	81,55	R442 863,46	1
ATS-3 PRM-2	4	0	24	80,71	R449 720,31	1
ATS-3 PRM-3	4	0	16	68,23	R382 006,69	1
ATS-3 PRM-4	4	0	16	68,17	R382 006,69	1
ATS-3 PRM-5	4	0	21	78,36	R424 631,50	1
ATS-3 PRM-6	4	0	24	82,90	R449 720,31	1
ATS-3 PRM-7	4	0	16	63,19	R331 642,07	1
ATS-3 PRM-8	4	0	19	74,24	R409 268,32	1
ATS-4 PRM-1	5	4	28	54,12	R422 458,00	0,980025
ATS-4 PRM-2	5	4	28	50,00	R422 458,00	0,980025
ATS-4 PRM-3	5	4	28	54,37	R422 458,00	0,980025

Model Code	Aircraft Used	Unused Aircraft	Flights	Schedule Density (%)	Profit (R)	Load Factor
ATS-4 PRM-4	5	4	28	49,76	R422 458,00	0,980025
ATS-4 PRM-5	5	4	28	52,45	R422 458,00	0,980025
ATS-4 PRM-6	5	4	28	50,68	R422 458,00	0,980025
ATS-4 PRM-7	5	4	28	44,86	R425 643,93	0,982573
ATS-4 PRM-8	5	4	28	50,91	R422 458,00	0,980025
ATS-5 PRM-1	2	0	18	98,67	R352 030,41	1
ATS-5 PRM-2	2	0	16	91,37	R324 041,99	1
ATS-5 PRM-3	2	0	12	67,70	R225 130,05	1
ATS-5 PRM-4	2	0	12	68,47	R228 603,00	1
ATS-5 PRM-5	2	0	12	75,95	R286 046,41	1
ATS-5 PRM-6	2	0	16	91,37	R324 041,99	1
ATS-5 PRM-7	2	0	10	61,69	R220 287,10	1
ATS-5 PRM-8	2	0	12	76,94	R272 332,71	1
ATS-6 PRM-1	2	0	24	97,36	R410 449,95	1
ATS-6 PRM-2	2	0	22	90,07	R393 053,00	0,998468
ATS-6 PRM-3	2	0	14	68,91	R304 175,62	1
ATS-6 PRM-4	2	0	14	69,37	R304 175,62	1
ATS-6 PRM-5	2	0	17	78,80	R348 610,89	1
ATS-6 PRM-6	2	0	22	90,07	R393 053,00	0,998468
ATS-6 PRM-7	2	0	12	64,09	R286 046,41	1
ATS-6 PRM-8	2	0	16	74,37	R324 041,99	1
ATS-7 PRM-1	10	7	28	33,38	R422 713,50	0,980025
ATS-7 PRM-2	10	7	28	35,16	R422 458,00	0,980025
ATS-7 PRM-3	10	7	28	29,81	R422 458,00	0,980025
ATS-7 PRM-4	10	7	28	31,52	R422 713,50	0,980025
ATS-7 PRM-5	10	7	28	33,38	R422 713,50	0,980025

Model Code	Aircraft Used	Unused Aircraft	Flights	Schedule Density (%)	Profit (R)	Load Factor
ATS-7 PRM-6	10	7	28	34,27	R422 713,50	0,980025
ATS-7 PRM-7	10	7	28	34,34	R422 586,09	0,980025
ATS-7 PRM-8	10	7	28	29,70	R422 713,50	0,980025

Table A-3: Complete Monte Carlo Simulation Observations

Model Code	Percentage of Propagated Delayed Flights	Delay Recovery %	Average Propagated Delay	Total Average Propagated Delay	Average Time Difference Departures	Average Time Difference Arrivals
BASE PRM-1	36,54	23,16	51,73	18,901	29,98	24,60
BASE PRM-2	23,46	36,52	42,78	10,037	21,19	15,81
BASE PRM-3	2,50	83,87	36,93	0,923	10,81	5,73
BASE PRM-4	3,54	76,68	43,86	1,553	11,74	6,66
BASE PRM-5	12,50	50,04	46,07	5,759	16,61	11,53
BASE PRM-6	20,96	38,43	44,13	9,250	20,40	15,03
BASE PRM-7	2,08	86,20	38,62	0,804	10,91	5,83
BASE PRM-8	9,62	58,34	37,99	3,653	14,14	8,77
ATS-1 PRM-1	7,86	67,11	43,00	3,379	15,01	10,27
ATS-1 PRM-2	11,07	58,29	46,13	5,107	16,98	12,24
ATS-1 PRM-3	5,00	72,84	37,56	1,878	12,29	6,91
ATS-1 PRM-4	3,85	78,97	31,84	1,225	11,20	5,82
ATS-1 PRM-5	10,89	59,91	43,98	4,790	16,68	11,95
ATS-1 PRM-6	8,93	63,58	44,55	3,977	15,66	10,92
ATS-1 PRM-7	4,42	77,19	35,48	1,569	12,26	6,88
ATS-1 PRM-8	5,58	72,16	35,64	1,988	12,52	7,14
ATS-2 PRM-1	37,29	21,10	48,01	17,905	27,78	22,69
ATS-2 PRM-2	21,25	34,37	48,16	10,234	20,67	15,59
ATS-2 PRM-3	3,75	80,22	34,18	1,282	11,56	6,48

Model Code	Percentage of Propagated Delayed Flights	Delay Recovery %	Average Propagated Delay	Total Average Propagated Delay	Average Time Difference Departures	Average Time Difference Arrivals
ATS-2 PRM-4	3,54	79,87	34,35	1,216	11,12	6,04
ATS-2 PRM-5	7,92	54,45	52,55	4,160	14,21	9,13
ATS-2 PRM-6	19,17	37,66	43,20	8,281	18,36	13,28
ATS-2 PRM-7	2,92	84,12	31,75	0,926	10,91	5,83
ATS-2 PRM-8	2,92	83,73	27,13	0,791	9,95	4,86
ATS-3 PRM-1	38,96	23,72	44,74	17,429	27,93	22,85
ATS-3 PRM-2	22,92	37,31	42,19	9,669	20,50	15,42
ATS-3 PRM-3	3,75	82,32	37,28	1,398	13,31	7,90
ATS-3 PRM-4	3,44	82,45	40,28	1,385	13,30	7,89
ATS-3 PRM-5	17,38	48,17	55,87	9,711	23,64	18,73
ATS-3 PRM-6	28,75	31,75	42,20	12,133	22,86	17,78
ATS-3 PRM-7	1,88	86,05	51,59	0,967	12,35	6,94
ATS-3 PRM-8	8,95	58,07	57,79	5,171	18,08	12,33
ATS-4 PRM-1	7,68	62,82	50,78	3,899	15,22	10,49
ATS-4 PRM-2	2,14	90,86	24,67	0,529	10,52	5,78
ATS-4 PRM-3	4,11	79,16	38,72	1,590	12,37	7,63
ATS-4 PRM-4	4,82	77,35	36,32	1,751	12,47	7,73
ATS-4 PRM-5	3,21	80,67	45,16	1,452	12,24	7,51
ATS-4 PRM-6	2,86	83,57	38,18	1,091	11,38	6,64
ATS-4 PRM-7	2,50	89,50	27,71	0,693	11,33	6,60
ATS-4 PRM-8	3,39	82,78	38,31	1,300	12,28	7,55
ATS-5 PRM-1	71,94	16,08	68,63	49,376	64,59	58,83
ATS-5 PRM-2	37,19	26,81	48,19	17,921	29,90	24,49
ATS-5 PRM-3	10,00	62,78	41,62	4,162	15,74	11,18
ATS-5 PRM-4	9,58	61,49	45,89	4,398	15,97	11,42

Model Code	Percentage of Propagated Delayed Flights	Delay Recovery %	Average Propagated Delay	Total Average Propagated Delay	Average Time Difference Departures	Average Time Difference Arrivals
ATS-5 PRM-5	12,50	53,02	49,78	6,222	17,80	13,24
ATS-5 PRM-6	37,19	26,81	48,19	17,921	29,90	24,49
ATS-5 PRM-7	3,00	85,97	30,30	0,909	10,97	6,48
ATS-5 PRM-8	14,17	49,16	51,25	7,261	18,84	14,28
ATS-6 PRM-1	63,54	12,63	62,79	39,900	50,75	45,67
ATS-6 PRM-2	44,09	21,52	60,54	26,694	38,97	34,01
ATS-6 PRM-3	4,64	70,31	62,58	2,906	15,20	9,79
ATS-6 PRM-4	5,36	68,41	59,32	3,178	15,47	10,06
ATS-6 PRM-5	22,65	38,82	58,71	13,295	26,92	21,73
ATS-6 PRM-6	44,09	21,52	60,54	26,694	38,97	34,01
ATS-6 PRM-7	5,00	77,39	37,52	1,876	12,85	8,30
ATS-6 PRM-8	11,56	56,88	43,04	4,977	16,95	11,54
ATS-7 PRM-1	0,18	99,73	6,43	0,011	9,00	4,27
ATS-7 PRM-2	0,54	99,19	6,82	0,037	9,27	4,53
ATS-7 PRM-3	0,00	100,00	0,00	0,000	9,09	4,35
ATS-7 PRM-4	0,00	100,00	0,00	0,000	9,23	4,50
ATS-7 PRM-5	0,00	100,00	0,00	0,000	9,35	4,61
ATS-7 PRM-6	0,18	99,02	25,04	0,045	9,28	4,54
ATS-7 PRM-7	0,00	100,00	0,00	0,000	9,46	4,72
ATS-7 PRM-8	0,18	98,87	30,04	0,054	9,47	4,74

APPENDIX B ZIMPL MODEL INDEX FILE

```
# Airline scheduling model
# Mark Silverwood (0512705N)
# University of the Witwatersrand
# 2011

set V := { read "profits.dat" as "<1n>" comment "#"};
# of Variables
set VG      := { 1 to 2304 };# of ground variables
set R       := { read "routes.dat" as "<1n>" comment "#"};
# of routes
set RW      := { 1 to 276 }; # width of routes parameter
set A := { read "fleet.dat" as "<1n>" comment "#"}; # of fleets
set ND      := { read "takeoffs.dat" as "<1n>" comment "#"};
# of nodes
set TKW     := { 1 to 7}; #width of takeoffs parameter
set LDW     := { 1 to 7}; #width of landings parameter
set TGW     := { 1 to 1}; #width of takeoffs ground parameter
set LGW     := { 1 to 1}; #width of landings ground parameter
set CFW := { 1 to 44};
set CGW := { 1 to 8};

param p[V]  := read "profits.dat" as "<1n> 2n" comment "#";
# cost per flight variable
param rnum[R]     := read "routes.dat" as "<1n> 2n" comment "#";
# number of flight variables per route
param air[A]      := read "fleet.dat" as "<1n> 2n" comment "#";
# of aircraft per fleet type
param tkn[ND] := read "takeoffs.dat" as "<1n> 2n" comment "#";
# number of flights taking off at specific node
param ldn[ND] := read "landings.dat" as "<1n> 2n" comment "#";
# number of flights landing at specific node
param tgn[ND] := read "takeoffsgnd.dat" as "<1n> 2n" comment "#";
# number of flights taking off at specific ground node
```

```
param lgn[ND] := read "landingsgnd.dat" as "<1n> 2n" comment "#";
# number of flights landing at specific ground node
param cf[A] := read "countf.dat" as "<1n> 2n" comment "#";
# count flights air
param cg[A] := read "countg.dat" as "<1n> 2n" comment "#";
# count flights ground
param pr[R] := read "prices.dat" as "<1n> 2n" comment "#";
# seat prices
param dm[R] := read "demand.dat" as "<1n> 2n" comment "#";
# route demands

include "routes.txt";
# Parameter table containing the variable numbers for each route
include "takeoffs.txt";
# Parameter table containing flow out of each node in network
include "landings.txt";
# Parameter table containing flow into each node in network
include "takeoffsground.txt";
# Parameter table containing ground arc flow out of each node in
network
include "landingsground.txt";
# Parameter table containing ground arc flow into each node in
network
include "countf.txt";
# Flight arcs at counting point
include "countg.txt";
# Ground arcs at counting point
include "fleetcaps.txt";
# Seating Capacities of each Fleet Type

var x[V]    binary;
var y[VG]   integer;
var pen[R]  integer;

maximize    profit: sum <i> in V: p[i] * x[i]
```

```
                - sum <j> in R: pen[j]*pr[j];

subto       pen: forall <i> in R do
                sum   <j>   in   {   1   to   rnum[i]   }   :
x[routes[i,j]]*fleetcaps[i,j]
                - pen[i]
                <= dm[i];
 subto          flow: forall <i> in ND do
                sum <j> in { 1 to tkn[i] } : x[takeoffs[i,j]]
                +   sum   <k>   in   {   1   to   tgn[i]   }   :
y[takeoffsground[i,k]]
                - sum <w> in { 1 to ldn[i] } : x[landings[i,w]]
                -   sum   <q>   in   {   1   to   lgn[i]   }   :
y[landingsground[i,q]] == 0;
subto       avail:      forall <i> in A do
                sum <j> in { 1 to cf[i] } : x[countf[i,j]]
                + sum <k> in { 1 to cg[i] } : y[countg[i,k]] <=
air[i];
```

APPENDIX C GENERATED FLIGHT SCHEDULE FOR BASE PRM-1

MODEL

Boeing 737-800

Departure Time	Origin	Arrival Time	Destination
06:00	CPT	07:30	DUR
08:00	DUR	09:30	CPT
10:00	CPT	11:30	DUR
12:00	DUR	12:33	BFN
13:15	BFN	14:21	CPT
15:00	CPT	16:06	BFN
16:45	BFN	17:12	JNB
17:45	JNB	19:17	CPT
19:45	CPT	21:17	JNB

Boeing 737-800

Departure Time	Origin	Arrival Time	Destination
06:00	JNB	07:32	CPT
08:00	CPT	09:32	JNB
10:00	JNB	11:06	PLZ
11:45	PLZ	12:51	JNB
13:30	JNB	13:57	BFN
14:30	BFN	15:03	DUR
15:45	DUR	17:15	CPT
17:45	CPT	19:17	JNB
19:45	JNB	21:17	CPT

McDonnell Douglass MD-81

Departure Time	Origin	Arrival Time	Destination
06:00	DUR	06:37	JNB
08:15	JNB	08:35	NLP
10:15	NLP	10:35	JNB
12:15	JNB	12:43	BFN
14:15	BFN	14:43	JNB
16:15	JNB	16:52	DUR
18:30	DUR	19:07	JNB
20:45	JNB	21:22	DUR

APPENDIX D REAL-WORLD TEST CASE (JETBLUE AIRWAYS)

JetBlue Airways is an American low-cost airline, and is currently placed seventh in the North American Domestic market with a market share of 4.4%. It operates a hub-and-spoke type flight network with the main base hub of the network being JFK Airport in New York City. It operates a fleet of 164 aircraft, consisting of 117 Airbus A320-200 aircrafts and 47 Embraer E190 aircraft. JetBlue serves 50 destinations worldwide and 24 destinations with the borders of the United States of America (airports served is more since several destinations has multiple local airports).

Presented in Table D-1 are the observations of the JetBlue optimization model.

Table D-1: JetBlue Model Observations

Delay Time Setting	39 minutes
Airports	24
Fleet Types	2
Total Aircraft	164
Solving Time	8,8 seconds
Model Size	
Rows	3908
Columns	12466
NonZeros	33738
Aircraft Used	33
Unused Aircraft	131
Flights	80
Schedule Density %	65,8
Delay Recovery %	98,2198
Total Average Propagated Delay	0,0704 minutes