Improving Aircraft Departure Time Predictability

by

William W. Vanderson

B.S. Computer Science and Electrical Engineering
Massachusetts Institute of Technology, 1999

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2000

©William W. Vanderson, 2000. All Rights Reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly
paper and electronic copies of this thesis document in whole or in part.

Author ....
Department of Electrical Engineering and Computer Science
June 21, 2000

Certified by ....
William D. Hall
Charles Stark Draper Laboratory
Thesis Supervisor

Certified by ....
John-Paul Clarke
Assistant Professor of Aeronautics and Astronautics
Thesis Supervisor

BARKER

Accepted by ....
Arthur C. Smith
Chairman, Department Committee on Graduate Students

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

JUL 27 2000

LIBRARIES

Improving Aircraft Departure Time Predictability

by

William W. Vanderson

Submitted to the Department of Electrical Engineering and Computer Science
on June 21, 2000, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

In this thesis, a forecasting model is described that improves departure predictions over those of Collaborative Decision Making (CDM), reducing error by up to 30% for a given day. This model propagates delay from incoming flights to outgoing flights by using minimum turn times calculated from Airline Service Quality Performance (ASQP) data. The model was run on data covering every day of March, April, and May of 1999, and produced departure predictions 6 hours, 4 hours, 2 hours, and 1 hour in advance of departure time. The greatest improvement on CDM predictions was achieved in the 4-hour predictions.

Thesis Supervisor: William D. Hall
Title: Senior Member of the Technical Staff, Charles Stark Draper Laboratory
Lecturer, MIT Department of Aeronautics and Astronautics

Thesis Supervisor: John-Paul Clarke
Title: Assistant Professor, MIT Department of Aeronautics and Astronautics

Acknowledgment

Thanks to my parents for sending me to MIT, and for all the knowledge and wisdom they have given me over my 23 years.

Thanks Bill Hall for being the most helpful and understanding supervisor possible, and for teaching me everything about airport and airline operations. He's a grammar whiz, too.

Thanks to John-Paul Clarke for supervising my thesis and giving me direction when I needed it.

Thanks to Gehrke for everything.

Thanks to Kari, my officemate, who made me waste so much time that I almost didn't finish.

Thanks to Mike, my partner in crime.

This thesis was prepared at The Charles Stark Draper Laboratory., Inc., under IR&D Project Number 15091, Collaborative Air Transportation System Throughput Enhancement.

Publication of this thesis does not constitute approval by Draper or the sponsoring agency of the findings or conclusions contained herein. It is published for the exchange and stimulation of ideas.

6/26/00

William W Vanderson — Date

Contents

Chapter 1

Introduction

This chapter contains the motivation and purpose of the work, the contributions, and a general background.

1.1 Motivation

The National Airspace System (NAS) has reached a crisis point. Demand for air travel in the United States is rapidly outstripping airport capacity. A special issue of Aviation Week titled *Air Travel in Crisis* predicts that at current capacity, the system will "descend into gridlock" by 2008[5]. The summer of 1999 saw 50% more delay than the same period in 1998, the beginning of a trend that continued throughout the year. Adding additional airport capacity is difficult, at best, and only a few airports have expansion plans for the near future. Aviation Week concludes that if changes are not implemented to use existing capacity more efficiently, delays will increase by 250% in the next five years.

One means to facilitate more efficient use of capacity is through improved data availability and accuracy. In order to use our existing capacity to accept increasing levels of air traffic, better and more informed decisions must be made by the FAA and the airlines. These decisions can only be improved by having accurate data reach the appropriate decision-makers at the right time. Collaborative Decision Making (CDM) is a step in the right direction.

The purpose of CDM is to produce better traffic flow management decisions by improving the level of cooperation between the FAA and the airlines [2]. The most extensive use of CDM to date is in the area of capacity allocation during times of reduced capacity. Through CDM, an electronic prediction of future demand for airport capacity called the Aggregate Demand List (ADL) is produced every five minutes. The ADL is a more accurate picture of the current and near-future demand than was previously available. However, CDM can be improved further. CDM relies on the parties involved to produce accurate data and to make that data available to the system. When one of the parties is delinquent in reporting the data it possesses, the quality of ADL predictions suffers.

CDM departure time data are often inaccurate. Especially in times of irregular operations, when many departures are delayed, departure demand differs significantly from CDM-predicted demand. CDM departure predictions are based on airline data feeds and the Enhanced Traffic Management System (ETMS). ETMS is myopic, as it has a wealth of information at its disposal concerning arrival times, but must rely on airlines to meet their schedules to be accurate when it comes to departure times. Unfortunately, airlines may not report their departure delays until the scheduled departure time or later.

The objective of this thesis is to design and test a tool that improves departure time predictability. Not only would implementation of such a tool enrich the CDM and ETMS data, it would also increase the effectiveness of ground delay programs, collaborative routing efforts, departure queue prediction, and airport arrival-departure capacity utilization.

1.1.1 Ground Delay Program (GDP)

Every airport has a certain maximum number of flights it can allow to arrive without compromising safety. The controllers' estimate of this limit is called the Airport Acceptance Rate (AAR), and is low enough to allow airport capacity to be used for departing aircraft as well as arriving aircraft. In fair-weather conditions, this limit may be much greater than demand, but in inclement weather, the AAR must be decreased so that safety is maintained. In the past, when arrival demand exceeded capacity, aircraft were delayed in airborne holding patterns, wasting jet fuel and increasing risk exposure.

Ground Delay Programs (GDPs) were created to keep the number of arriving flights at an airport at or below the AAR, while decreasing or eliminating airborne holding. When arrival demand exceeds capacity at an airport, a GDP is initiated. A GDP delays the departure times of aircraft bound for the capacity-limited airport so that they arrive at that airport at a rate no greater than the AAR. The FAA decreases the AAR for the troubled airport in order to delay flights sufficiently so that when they take off from their origin airports, they can proceed directly to the troubled airport with a minimum of airborne holding. GDPs effectively transform airborne delay at a flight's destination to ground delay at the flight's origin. This solution is more economical than airborne holding (by conserving fuel) and also avoids the safety implications of excessive airborne holding.

GDPs are performed by assigning "slots" to arriving flights. Because the decision to implement a GDP is made based on CDM data, cancellations or delays that have not been entered into the system may not be accounted for. In one hour on 1 July 1999, a GDP run at St. Louis Airport assigned 56 slots, only 32 of which were used. In the next hour, 52 slots were assigned and 30 were used [8]. In this case, bad weather on the East Coast forced airlines to delay many flights that would have landed at St. Louis, but the FAA did not account for these delays until hours after the aircraft had failed to appear in the St. Louis airspace. If the delays had been reported early enough, the GDP might not have been necessary.

A useful aspect of GDPs is substitution. Substitution allows airlines to change the assignment of flights to their GDP slots. Substitution provides for greater flexibility, as airlines can move up certain flights to take the place of cancelled or delayed ones. In addition, CDM allows the FAA to know of empty slots due to airline cancellations, so the FAA can parcel out those slots to other airlines if the original slot owner does not substitute. Unfortunately, empty slots can still occur, as seen in the St. Louis example, where demand was perceived to be much higher than it actually was.

Though an underutilization of the magnitude felt in St. Louis in July of 1999 is uncommon, the Ground Delay Program is not entirely effective. Taking updated departure information into account could greatly improve GDPs, and may demonstrate that they are unnecessary in some circumstances. In fact, in the St. Louis example, the GDP was cancelled after four hours because ATC realized that capacity far exceeded demand. Similarly, accurate departure data system-wide could reduce delays felt by all those affected by GDPs.

1.2 Purpose

This research investigates whether the quality of the departure time predictions used by Air Traffic Management (ATM) could be improved by modeling aircraft inventory. The result of the research is the design of a Departure Forecasting Model that takes delay on flights arriving at an airport and translates it into delay on flights departing the airport. This translation (propagation) of delay models aircraft inventory by not allowing flights to depart until the corresponding incoming leg has arrived and the aircraft has been turned.

1.3 Contributions

The contributions of this work are two-fold. First, a Departure Forecasting Model (DFM) is demonstrated that is able to improve the average absolute error of departure predictions by up to 30% on certain days. This improvement over CDM prediction is greatest when forecasting four hours in advance, in the presence of large departure delay (a prediction that is vital for Ground Delay Program (GDP) operation). Second, this work has shown that the accuracy of CDM departure data is degraded by large delays.

1.4 Background and Previous Work

1.4.1 Current State of Air Traffic Control (ATC)

Air traffic controllers make all major decisions concerning traffic management. Controllers do, however, have many information sources to help them in performing their duties, such as Enhanced Traffic Management System (ETMS). But final judgments concerning traffic control are still made by humans.

1.4.2 Collaborative Decision Making (CDM)

CDM can trace its beginnings to the FAA/Airline Data Exchange (FADE) experiments performed in 1993. These experiments established that timely airline submission of updated departure time estimates to the FAA would positively impact air traffic management decision-making [2]. CDM was officially launched in the spring of 1995, and relies on the sharing of information between airlines and the FAA. Using this more accurate data, better air traffic control decisions can be made to increase efficiency. In addition, CDM has proven useful for discovering errors or problems in both ATC and airline operations because more data are available concerning each flight. Not only has CDM produced enhancements in FAA programs such as GDPs, it has proven itself useful in that it allows the industry to interact with the FAA, leading to greater understanding between parties that should share similar goals.

Steps have been taken to improve departure delay information. CDM and substitution have already improved the efficiency of GDPs. However, CDM relies on airlines to provide timely, accurate data, which is not always in the airlines' best interest. For passenger satisfaction reasons, many airlines do not inform the FAA of delays until the passengers are informed; passengers will complain if other parties know of their delays before they do. Therefore, it is unlikely that the FAA will ever know all the information available to the airlines.

In addition, airlines may not advise passengers of delays the moment those delays are discovered. Since delays can be decreased by faster flight speeds and quick turns, the airline agents will sometimes wait until the scheduled boarding time to inform passengers of delay. Should a plane suddenly become available to service the flight, the passengers will be able to board since they have not left the gate area. Had they been advised of the erroneously predicted delay, they might not have been present to board the flight. This situation is one of many in which airlines would withhold departure delay information from passengers. Finally, the airlines' data may be incomplete. In each case, CDM accuracy suffers.

1.4.3 Departure Prediction

CDM and ETMS departure data can be inaccurate. Contributing to this lack of accuracy is the fact that ATC does not oversee ground operations. Once a plane has landed and arrived at the gate, controllers typically are unaware of the aircraft's status until the plane pushes back from the gate to taxi to the runway or makes a request to enter the taxiway system. There is seldom any way for ATC to know whether a flight departure will be delayed unless the airline reports the delay through CDM. Due to the complexity of ground operations, from refueling to catering to deplaning and boarding passengers, delays can be frequent and significant. In fact, there are 33 separate ground operations performed in an average turn (see Appendix A).

Additionally, ATC does not keep track of airplanes after they have parked at the gate. If no flights from a certain carrier have landed in the last two hours, ATC will still expect that carrier's outgoing flights to depart on schedule, even though the carrier may have no planes at the affected airport. Similarly, if all of a carrier's aircraft have landed in the last ten minutes and there are a number of flights scheduled to depart before those aircraft can be turned, the departures will be delayed. ATC would still expect the departure schedule to be followed, even though none of the planes could be turned around that quickly.

Due to ATC's inability to model ground operations, CDM must rely on the airlines to produce accurate departure information. In some cases this would be sufficient, but airlines usually post delays at a time near to the scheduled departure (as mentioned previously), which can be too late. If departure demand statistics are to be generated an hour or more in advance, airline updates also need to be produced at least an hour or more in advance.

1.4.4 Enhanced Traffic Management System

The Enhanced Traffic Management System (ETMS) provides much of the data used by CDM. ETMS derives predictions of airborne traffic from radar and flight plan data. ETMS outputs a projected flow of traffic into airports and sectors, which can be used by air traffic controllers to decide when to implement GDPs, airborne holding, or other flow management procedures. ETMS also runs the "Monitor Alert" system that displays an alert when demand is predicted to exceed capacity in a particular area.

1.4.5 Airline Service Quality Performance (ASQP)

The Federal Bureau of Transportation Statistics compiles ASQP data. ASQP data provide on-time departure statistics for all airlines with at least one percent of total domestic passenger revenues. At present, this list includes the ten largest carriers in the United States: Alaska Airlines, America West, American, Continental, Delta, Southwest, TWA, United, and USAir. Regional carriers are not represented, even if those carriers are partners with one of the ten largest. ASQP data contain many fields for each flight and those of the most concern to this work are gate arrival time, gate departure time, and aircraft tail number. The tail number is an alphanumeric string that uniquely identifies each registered airplane.

1.4.6 Previous Work

In his Ph.D. thesis titled "Dynamic Statistical Models for the Prediction of Aircraft Take-off Times", Robert Shumsky considers the problems of airport and runway congestion [6]. Shumsky's method predicts take-off times in order to alleviate congestion (caused by inaccurate prediction) down stream, either in the air or at arrival airports. He proposed a statistical model to account for gate delay that did not consider turn times. Shumsky concentrated mostly on runway congestion and departure queue delays, not on the gate departure delay. In normal operations, taxi delays tend to outweigh gate delays. The work described in this thesis is an extension of Shumsky's work, but to improve departure time predictions when an airport is under stress, in which case taxi delays tend to be much smaller than gate delays.

Kari Andersson models airline ground operations in a Master's thesis titled "Potential Benefits of Information Sharing During the Arrival Process at Hub Airports". Andersson uses this model to asses the potential benefit of sharing information between the airlines and ATC during the arrival process [1]. Though not explicitly concerned with predicting departures, Andersson's work schedules airline ground facilities, personnel, and aircraft in order to turn those aircraft, and produces a departure schedule. Since these schedules are formulated 3 or 4 hours in advance, airline departure predictions would be improved for those time periods. This improvement would likely lead to an improvement in CDM predictions as a whole, as long as airlines were able to report their predictions in a timely manner. Additionally, ATC could implement Andersson's model to obtain better predictions of departure times without airline involvement.

Chapter 2

Model Formulation

In this chapter, the design and function of the Departure Forecasting Model (DFM) will be discussed. The operation of the model will be described, followed by an explanation of the inputs and outputs. The DFM was implemented in a test environment in order to evaluate its operation. The environment itself and the differences between the test and a real-world implementation of the DFM are also described.

2.1 Model Function

The Departure Forecasting Model establishes a matching for each aircraft turn. Then, if the sum of incoming arrival time and minimum turn time for any outgoing flight is greater than its original departure time, a new departure time is generated. Next, the departure queue is updated. This process is repeated for every incoming ADL update.

In order to facilitate the explanation of the DFM, its functions will be described in the order that they are performed for a single flight. The following subsections expand on each of the descriptions of the steps below. The general functions that the DFM performs are:

1. Receive and read in flight information from CDM ADL updates

2. Match departing flights with arrivals (In the test environment, this matching is

accomplished by comparing ASQP tail numbers)

3. Propagate delay from arrivals to departures, accounting for turn time

4. Output predictions

5. Repeat Steps 3 and 4 for Each ADL Update

Two additional functions are performed in the test environment to evaluate the predictions of the DFM. The results of these functions are reported in Chapter 4:

1. Compare CDM predictions with DFM predictions

2. Output actual (ASQP) departure times

2.1.1 Flight Information Received

The first step the Departure Forecasting Model takes is to read in the data contained in an ADL update. The model obtains arriving flights by reading in ADL updates corresponding to the airport being modeled. Departing flights are read in from the ADLs corresponding to all of the other airports. These flights are entered into arrival or departure lists, respectively. Once all flights for an update are read in, departing flights are matched to arriving flights.

2.1.2 Flight Leg Matching

The DFM establishes a pairing between the incoming flights and their outgoing legs. First, the DFM matches departing flights to ASQP flight records containing the same carrier and flight number information. Then, an arriving flight is found in ASQP that matches the tail number of the departing ASQP flight. Third, the ASQP arrival is matched to a flight in the arrival list by carrier and flight number. Finally, the departing flight's data record is updated to contain a pointer to the matched, arrival flight. Every subsequent prediction update takes into account the arrival time of this matched flight.

If the DFM were to be implemented on-line, ASQP data would not be available to match arrivals to departures. In that case, a general matching could be accomplished using the Aircraft Type field of the ADL data. Though not as specific as Tail Numbers, the aircraft type field does differentiate between series of aircraft. For example, an Aircraft Type of B733 refers to a Boeing 737, series 300, while B732 refers to the 200 series. Appendix C contains a list of all ADL aircraft types used in the DFM. This general matching could affect the DFM's accuracy on a per-flight basis, as turns would not always be matched correctly. However, aggregate accuracy should not be affected too greatly, as aircraft inventory would still be taken into account.

Planes could be matched by placing all flights into FIFO queues, one queue for incoming flights, and one for outgoing flights. Matching would be accomplished by taking the first incoming plane with the correct Aircraft Type and pairing it with the plane departing the earliest in the future. If the departure time of a flight were to occur when there were no incoming aircraft to service it, the departure would be delayed.

2.1.3 Propagation of Delay

For each matched flight, the DFM makes a prediction based on the arrival time of the incoming leg. The average minimum turn time for the type of aircraft is found by table lookup, and added to the arrival time of the matched, incoming flight. The sum of these two times is the earliest that the flight can take off. If this time is later than the scheduled departure, the flight is predicted to be delayed.

After delay has been propagated to the departing flight, the CDM data value for departure time is compared with the sum of the incoming arrival time and the average turn time. The model chooses the later of the two values to be the new departure time.

After each ADL update, delay is propagated to every matched departure and the result is compared to the CDM data. This action ensures that any recent incoming delay or CDM updates have been taken into account.

2.1.4 Output Predictions

Finally, at the end of each 15-minute period, predictions are produced concerning departures 6 hours, 4 hours, 2 hours, and 1 hour in advance. In addition, a "forecast" is produced for the previous 15-minute time period, to find the baseline error between ASQP and CDM ADLs. For example, after the last ADL update received during the 0915 to 0930 period, predictions are generated for five departure periods: 0915-0930, 1015-1030, 1115-1130, 1315-1330, and 1515-1530.

Also at this point, departure and arrival queues are updated for all 15-minute periods.

2.2 Test Environment

A test environment was necessary to evaluate the Departure Forecasting Model because real-time CDM data were unavailable. In the test environment, the DFM was run on data from the Spring of 1999. The most significant difference between the test implementation of the DFM and a real-world implementation was the absence of the Cancellation and Swap Predictor.

The design of the DFM includes a Swap and Cancellation Predictor as one of its inputs (see Figure 2-1). As modeled departures become more and more delayed, the Predictor would label some of these departures as swapped or cancelled. The Predictor was not implemented in the test environment because ASQP data were used to match departures to arrivals. Since ASQP data are generated after the fact, any swaps that may have occurred are already included in the data set. Therefore, it is already known exactly which aircraft flew each leg, and the DFM does not have to take equipment swaps into account. Additionally, cancellations reported in ASQP result in a largely empty data field. These cancelled flights have no actual arrival times (for obvious reasons) and have no tail number entered. Because the DFM in the test environment only predicts those flights that could be matched through their tail numbers, cancellations were ignored due to the lack of the tail number in ASQP.

Were the Departure Forecasting Model to be implemented on-line, however, a

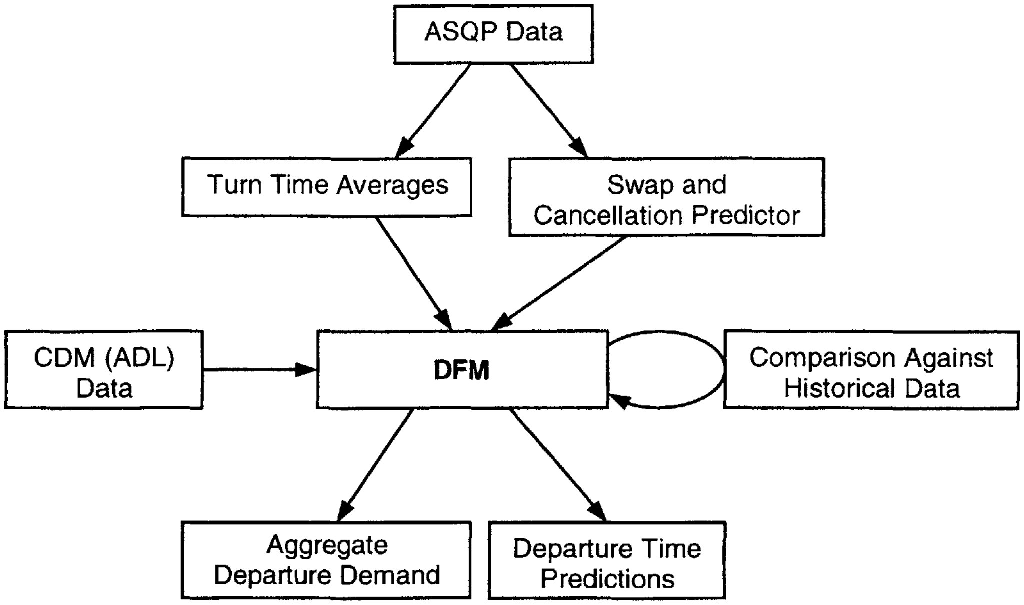

Figure 2-1: Departure Forecasting Model Diagram

Swap and Cancellation Predictor would be necessary. This Predictor would incorporate two conditional probabilities, one for equipment swaps and the other for cancellations. The probabilities would be applied each time a flight was updated with a greater delay than the previous ADL update. If a flight were to be swapped, no further swaps would be allowed, but cancellation probabilities would still be applied. If a flight were cancelled, on the other hand, neither probability would be applied.

If the Swap and Cancellation Predictor were added to the DFM in such a way as to mirror actual operations effectively, results similar to those reported in this thesis could be obtained without the use of ASQP data to determine swaps and cancellations.

2.3 Inputs and Outputs

As illustrated in Figure 2-1, the inputs to the DFM are the pre-calculated turn time averages, CDM data, and the Swap and Cancellation Predictor. The method by which average turn time data are obtained is discussed in detail in Chapter 3. Though the

Swap and Cancellation Predictor was not implemented in the test environment used in this thesis, the methods used to obtain its probabilities are also discussed in Chapter 3. Outputs of the DFM are the departure time estimates and departure demand queues for each time period.

2.3.1 CDM Data

CDM data are available through five-minute updates called Aggregate Demand Lists (ADLs). In every 24-hour period, starting at 0800$\mathcal{Z}$, there is an ADL update for each major domestic airport every five minutes. The first ADL of the day contains all the flights for a specific airport for the next 20 hours. Updates can consist of flights with changed information (cancellations and delays, for example) or new flights not in the original schedule. These new flights are called pop-ups because they appear suddenly in the CDM data when a new flight plan is filed. Though most pop-ups are the result of general aviation, airlines also submit flight plans when planes are diverted or unscheduled legs are flown.

ADL data contain information on all aircraft, from propeller planes to jets. The model described in this thesis only takes into account jet aircraft, because many props and turbo-props land or take off at regional airports that are not covered in the CDM data. To effectively model aircraft inventory, data must be available for both incoming and departing flights, which is not the case for many non-jet aircraft. Additionally, ASQP data cover only the top ten major carriers in the United States, exclusive of regional carriers. At a typical major airport, approximately 70% of all operations are preformed by jet aircraft (see Table 2.1).

In Table 2.2, the format of an ADL update is given. The format of a single flight record (such as one in an ADL update) is explained in CDM releases [7].

2.3.2 Updated Departure Times

As the DFM propagates delay, certain flights are given departure times different from those of the CDM data. The model outputs a new departure time for such flights

Airport	Percentage
Boston Logan	59
Chicago O'Hare	84
New York Laguardia	67
Philadelphia	66
Pittsburgh	54

Table 2.1: Percentage of operations performed by jet aircraft at five airports. Data compiled from CDM ADLs.

```
START_UPDATE 010057
START_ARRIVALS 3
UAL1687  DEN  ZDV  CMH  ZID LANDR 302239  B733 J L A302008 A302300 302008
302300 301959 302259 302012 302305 301959 302259 301959 302256 - - - - - - - -
- -  168 - - - - - - - - - - - - - - - - - C  0x00000000  Y N N
#
AWI130  DEN  ZDV  MTJ  ZDV POWDR 302240  D328 T S P302209 E302300 - - 302200
302255 302209 302300 302200 302255 302200 302245 - - - - - - - - - - -   45 - Y -
- - - - - - - - - - - - - - T  0x00000000  Y N N
#
UPS2804  DEN  ZDV  SDF  ZID DANDD 302244  B763 J H A302031 A302301 302031
302301 302031 302324 302033 302307 302031 302324 302016 302308 - - - - - - - -
- -  157 - - - - - - - - - - - - - - - - - A  0x00000000  N N N
#
END_ARRIVALS
```

Table 2.2: Arrivals section of an ADL update. The time of the update follows the header START_UPDATE. In this example, the date is the first of the month, and the time is 12:57AM (0057). The number of arrivals in the update follows the header START_ARRIVALS. In this case, there are 3 flights in the update. Each flight update contains estimated and actual departure times and arrival times, a unique call sign, origin and destination airports, equipment type, estimates from airlines, ETMS, or other agencies, GDP information, and other flight information [7].

End Time of 15-minute Period									Queue Output Time
930	945	1000	1015	1030	1045	1100	1115	1130	
3	1	5	6	5	4	4	6	1	7AM
3	1	5	6	5	4	4	6	1	9AM
3	2	5	5	5	5	5	5	1	11AM
3	2	5	5	5	5	1	5	3	1PM

Table 2.3: A two-hour portion of CDM departure predictions for the morning of May 31, 1999 at Chicago O'Hare Airport

in the format of an ADL update. This update complies with the ADL standards and has no missing fields, so it could be easily integrated into ADL updates in the real-world system. The only difference between the two would be an updated value in the "Estimated Time of Departure" field.

2.3.3 Departure Queue for CDM

The second output of the model is the CDM departure demand queue. As each ADL update is processed, the model outputs a departure demand forecast for every fifteen-minute period from the current time to the end of the 24-hour period. This forecast contains the number of flights that CDM predicts will take off in each fifteen-minute period. Table 2.3 shows a small portion of a CDM departure demand queue file.

2.3.4 Prediction Departure Queue

Another output of the DFM is the predicted departure queue. This is the number of flights the model has predicted will depart, after propagating delay and updating departure times. A side-by-side comparison of model-predicted and CDM departure queues is given in Chapter 4.

2.4 Implementation Details

The model was written in C++ on a PC in a Linux environment. The code for manipulating and averaging the ASQP data was also written in C++ and in the same environment. The model has a small number of options, which the user can change before running the model:

- Day of operations
- Airport of operations
- Carriers to include
- File names in which to save logged data

As the DFM is run, the departure and arrival queues are output to temporary files in spreadsheet format. Separate files are created by the model containing CDM, modeled, and actual (ASQP) departure times for each forecast.

At the present, the model has been run on the following airports: Chicago O'Hare International, Boston Logan International, Philadelphia International, Pittsburgh International, and New York La Guardia.

Chapter 3

Model Data

In this chapter, the data used in the Departure Forecasting Model and the methods and assumptions used in gathering that data are discussed. Two sets of values are necessary: (1) turn time averages, and (2) conditional probabilities for use in the Swap and Cancellation Predictor. Each is discussed in turn. Though the Swap and Cancellation Predictor was not added to the DFM in the test implementation, it was designed to play a significant role in forecasting. Formulation and discussion of the conditional probabilities used by the Swap and Cancellation Predictor are included not only as part of the design, but also to assist the implementation of future work.

3.1 Turn Times

Turn time averages are required for propagation of delay, which forms the backbone of the DFM, so they are described in detail. First, a description of an aircraft turn is given followed by assumptions and design choices made while collecting data. Additional constraints on data collection that were explored but not implemented are discussed in Appendix B.

An aircraft turn, as defined in this thesis, is the amount of time the aircraft is parked at the gate. The terms arrival and departure refer to parking at the gate and pushing back from the gate, respectively.

3.1.1 What Constitutes a Turn

There are three components of a normal turn. The nomenclature put forth by Jody Gittell [4] defines three quantifiers labeled Turn1, Turn2, and Turn3. Turn1 is the minimum time needed to prepare the aircraft for departure (deplane passengers and baggage, board new passengers and baggage, clean the aircraft, load food and fuel, etc.). Turn2 is the scheduled buffer to increase the likelihood of staying on schedule. Turn3 is extra time scheduled into the turn to allow for passenger connections. Airlines that do not operate with hubs, such as Southwest, do not have any Turn3 time in their turns, as they are not trying to make many connections for each flight. In addition, Southwest schedules very little Turn2 time, because it is trying to maximize flight time. Southwest also has the greatest on-time arrival percentage, so it does not have as much delay to make up in each turn [5].

ASQP data were collected to create realistic estimates for the minimum time needed to turn an aircraft, Turn1. The reason for focusing on the value for Turn1 is that when conditions deteriorate and flights are delayed, all the buffer in the schedule is taken up by delay, and airlines try to turn their planes as quickly as possible. In the following sections of the paper, the terms average minimum turn time and average turn time mean the same value: the average value for Turn1.

3.1.2 Assumptions

Two assumptions were made concerning the compilation of turn time data. The first is that individual Turn1 times do not vary significantly from year to year. The second is that an averaging of Turn1 times by carrier, airport, and type of aircraft is sufficiently granular to produce improvements in prediction accuracy. This assumption is reasonable because the same maintenance and gate crews will process the same type of aircraft, resulting in similar turn times between aircraft of the same type. Additionally, the major contributing factor to turn time is passenger unloading and loading, which is determined to a great degree by the number of passengers a type of aircraft can accommodate [4]. Distinguishing between carriers and airports is nec-

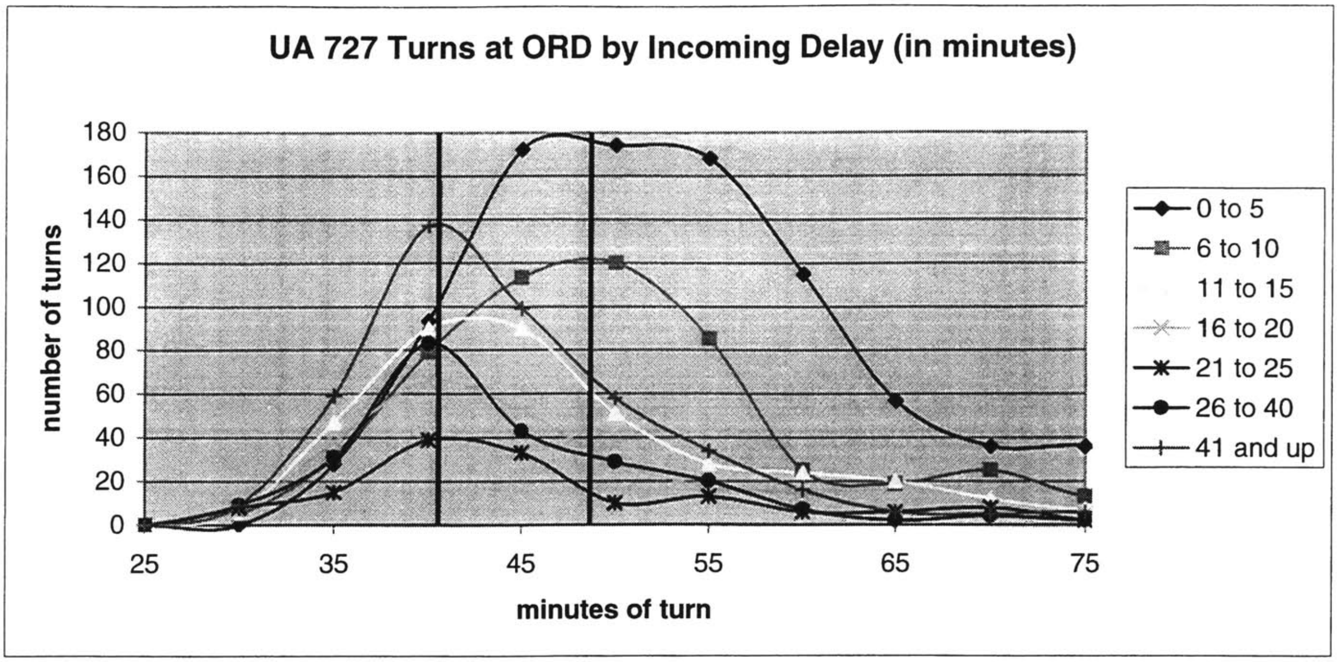

Figure 3-1: United Airlines 727 Turns at O'Hare Airport. Turns are graphed by the delay of the incoming leg of the turn. The left-most vertical line shows the most common turn time for those turns with incoming delays greater than 15 minutes. The vertical line on the right shows the most common turn time for those turns with incoming delays less than ten minutes. A saturation point at approximately 15 minutes of delay is evident. Data taken from ASQP; March, April, and May of 1999. All 727 turns are included that had incoming leg delays from 0 to 300 minutes.

essary because of differences in facilities, training, and motivation possessed by each airport and airline combination.

3.1.3 Effect of Incoming Delay on Turn Time

When a flight arrives either early or on schedule, the amount of time spent at the gate is greater than the minimum time it takes to turn the plane. As the incoming flight becomes more and more delayed, however, the time spent at the gate also decreases, until the incoming flight is so delayed that the turn cannot be completed before the scheduled departure time. These delayed turns are both the type of turn modeled in this thesis and the type of turn from which the data were collected.

As we can see from Figure 3-1, turn time decreases as incoming delay increases until approximately 15 minutes of incoming leg delay. This point at which incoming

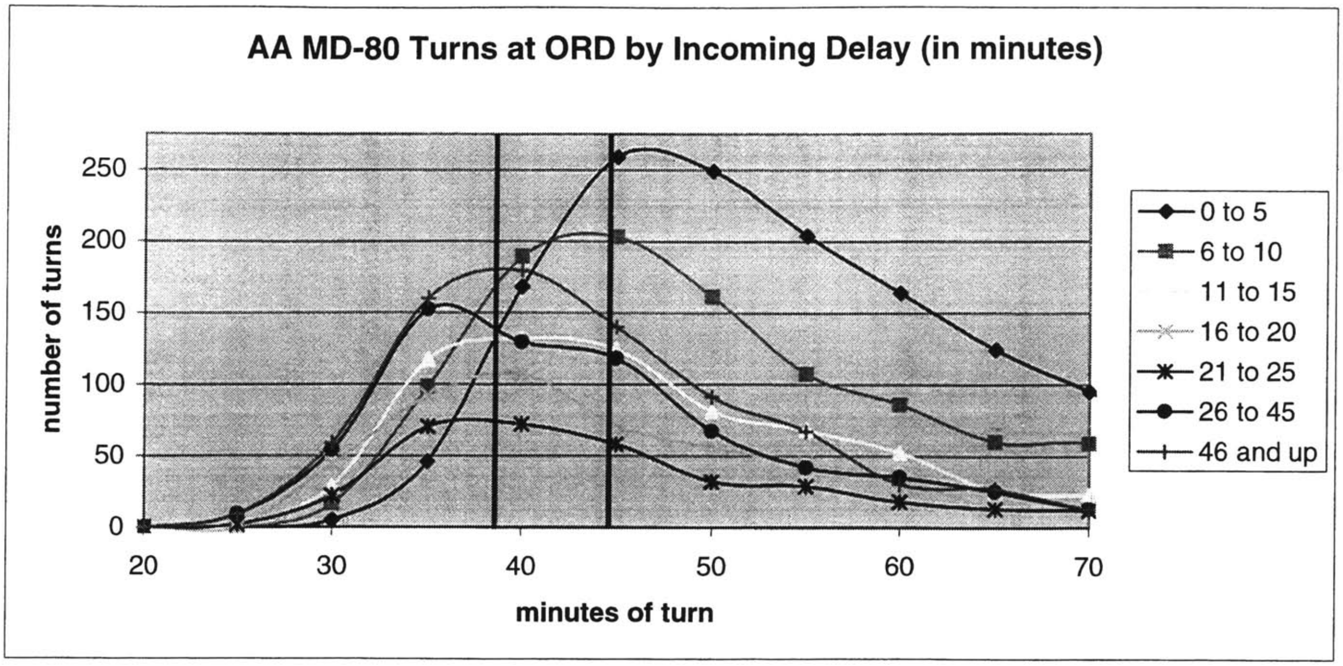

Figure 3-2: American Airlines MD-80 Turns at O'Hare Airport. Turns are graphed by the delay of the incoming leg of the turn. As in Figure 3-1, the leftmost line shows the minimum turn time and the right line shows an approximate average of the flights with an incoming delay of 0-10 minutes. Once again, a saturation point at approximately 15 minutes of delay is evident. Data taken from ASQP; March, April, and May of 1999. All MD-80 turns are included that had incoming leg delays from 0 to 300 minutes.

delay begins to translate directly into outgoing delay is very important. This point is the saturation point, where all padding in the scheduled turn has been exhausted. As delay increases from this point, turn time remains the same. Thus, averaging all turns made with an incoming delay sufficient to use up the scheduled padding gives a good estimate of the minimum time needed to turn a piece of equipment.

Figures 3-1 and 3-2 show that turn times are relatively constant after a certain amount of incoming delay has been exceeded. This relationship held for all carrier, airport, and aircraft sets examined and led to the decision to average all turns with incoming delay greater than 15 minutes to determine a value for Turn1. In the case of larger aircraft, such as the 747 and DC-10, cutoff points were 25 and 20 minutes of incoming delay, respectively. Larger planes tend to have larger overall turn times, larger minimum turn times, and more padding scheduled into their turns.

Aircraft Type	Cutoff (minutes)
B-727	15
B-737-1/2	15
B-737-3	15
B-737-5	15
B-757	15
FOKKER	15
DC-9	15
MD-80	15
DC-10	20
B-747	25
A-300	20

Table 3.1: Cutoff for each equipment type for Chicago's O'Hare airport. Values were obtained by examining the plot of turn time versus incoming delay, as in Figure 3-1. Cutoffs were found to be the same for each equipment type regardless of carrier. Cutoffs for other airports matched the values for Chicago.

3.1.4 Method for Obtaining Turn1 Averages

For every airport examined, incoming and outgoing flights were paired by tail number. Because the same aircraft can arrive and depart from the same airport many times in the same day, only pairings when the plane was scheduled to be on the ground for less than two hours were used. When all pairings were found for a specific day, those whose incoming legs were not sufficiently delayed were thrown out. The remaining pairings had their turn times calculated, and their equipment type looked up in an equipment database. All sufficiently delayed flights were averaged together for three months (March, April, and May of 1999) to arrive at a value for each carrier and equipment type.

3.1.5 Minimum Average Turn Times

The following is a list by carrier of the average minimum turn time (Turn1) for each type of aircraft used in this thesis (Tables 3.2 and 3.3). The average turn values for Chicago O'Hare and Boston Logan are displayed here, and the turn averages for the other three airports are included in appendix D.

	Airline						
Aircraft Type	**AA**	**UA**	**US**	**DL**	**NW**	**CO**	**TW**
B-727	-	45.5	-	41.3	41.3	-	-
B-737-1/2	-	-	38.0	39.8	-	-	-
B-737-3	50.6	43.4	43.0	-	-	33.6	-
B-737-5	-	52.9	46.2	-	-	-	-
B-757	56.4	53	56.7	56.9	44.8	-	-
FOKKER	42.5	-	37.6	-	-	-	-
DC-9	-	-	39.1	-	35.4	-	-
MD-80	43.6	-	41.9	53.2	-	59.8	41.8

Table 3.2: Most common equipment types and corresponding average minimum turn times for the 7 major carriers operating out of Boston Logan airport. Empty cells signify equipment that is not flown by the specific carrier. All turns with an incoming leg delay of 15 or more minutes were averaged from March to May of 1999. The averages for the 737 aircraft are broken up due to the difference between configurations. 737-1/2 refers to the F12Y96 configuration, 737-3 to the Y119 configuration, and 737-5 to the F16Y132 configuration (F = first class, Y = coach class).

	Airline						
Aircraft Type	**AA**	**UA**	**US**	**DL**	**NW**	**CO**	**TW**
B-727	48.6	47.2	-	49.6	32.7	45.3	42.1
B-737-1/2	-	39.3	36.6	-	-	-	-
B-737-3	-	45.7	41.7	48.7	-	44.0	-
B-737-5	-	40.2	46.8	-	-	52.8	-
B-757	60.8	52.9	-	56.3	48.9	-	-
DC-9	-	-	38.4	-	33.7	-	37.4
MD-80	46.8	-	42.7	45.5	-	41.6	40.8
A320-1/2	-	46.5	-	-	39.1	-	-

Table 3.3: Most common equipment types and corresponding average minimum turn times for the 7 major carriers operating out of O'Hare airport. All turns with an incoming leg delay of 15 or more minutes were averaged from March to May of 1999. The averages for the 737 aircraft are broken up just as in the previous table.

3.2 Swap and Cancellation Predictor

The Predictor incorporates two conditional probabilities that can be applied to each modeled aircraft turn. They are a probability of equipment swap and a probability of cancellation. The motivation behind modeling swaps and cancellations and the methods for obtaining the conditional probabilities are discussed in turn.

3.2.1 Equipment Swaps

Many factors can influence the decision to swap equipment, including delay, mechanical problems, gate assignment mistakes, and cancellations. However, delay and cancellation data are more readily available than that of the other factors and are more likely to affect equipment swaps in a predictable manner. Therefore, the swap portion of the Predictor discussed in this thesis will depend on cancellation and delay data.

3.2.2 Detecting Swapped Flights

Airlines schedule their planes to be on the ground only for the time necessary to turn the plane, permit connections to be made, and allow for a small amount of delay to be absorbed. Therefore, especially with smaller aircraft (DC-9 and 727, for example), scheduled turns are very rarely more than 90 minutes in length. Regardless of the actual time spent on the ground, scheduled turns in the ASQP data are nearly always less than two hours.

In almost every case, swapped turns appear to have a *scheduled* turn time greater than 90 minutes in the ASQP data. This phenomenon occurs because swaps are often used when the incoming leg delay is greater than the time needed to turn the plane. When the airline determines that the original turn cannot be made, it re-schedules the incoming aircraft for a different outgoing leg, making a swap. When the substituted leg is flown, ASQP does not record the original equipment connection. However, in the ASQP data, the incoming flight still has its badly missed scheduled arrival. The difference between the scheduled arrival and the scheduled departure time of the swapped leg has been observed to be greater than 90 minutes with a probability of 90%. This statistic was confirmed by comparing flights presumed to be swapped with the same flights on different days in the ASQP data. If, on most days, an incoming and departing flight pair shared the same aircraft, those days on which they did not likely involved a swap.

Swapped turns can therefore be located in the ASQP data by selecting those flights whose scheduled turn was greater than 90 minutes. This action leads to an accurate count of swaps for each day/route combination.

3.2.3 Swap Probability

When an incoming flight is severely delayed or cancelled, its aircraft will not be able to serve the corresponding outgoing leg in a timely manner. In this case, an aircraft originally scheduled to serve a different departure may be switched to the affected outgoing leg. This equipment swap minimizes the impact of the incoming delay and/or cancellation on that outgoing leg. Though specific exchange techniques of carriers are too detailed to study here, a swap function can be introduced into the Swap and Cancellation Predictor to model swap effects in aggregate. A swap normally affects real-world operations by reducing the delay experienced by any one flight. Therefore, it should have the same effect on a modeled flight when forecasted by the Predictor.

A swap probability can be determined by the number of swaps per equipment type, and is conditioned on incoming delay. In the Predictor, as an incoming leg becomes more and more delayed, the chance of a swap increases as long as there is an aircraft on the ground to be switched. Swaps' effect on the DFM, therefore, is to temper the delay felt by any one flight, assuming that an aircraft is present to provide service.

As with turn time, swaps are not heavily correlated with the overall delay felt at an airport. Therefore, only incoming delay affects the swap probability.

3.3 Cancellations

Cancellations affect departure times especially when the departing flight itself is cancelled. But cancellations at origin airports can significantly delay operations at destination airports as well. In order to model the cancellation process, a cancellation probability can be generated for use in the Swap and Cancellation Predictor. The cancellation probability is produced by an amalgamation of two factors: incoming delay and route.

Though lightly booked flights are more likely to be cancelled, data concerning booking is not available to the public or to air traffic control. Intuitively, flights that service the same route should have similar booking levels from day to day and from month to month. If cancellation does affect lightly booked flights, the route flown by these flights should produce higher probabilities of cancellation.

Cancellation probability is approximated by the sum of two conditional probabilities, one conditioned on the route being flown and another conditioned on the amount of incoming delay. Route and delay effects are explained in the next two subsections.

3.3.1 Conditioning on Route

The decision to cancel a flight is most affected by factors other than incoming delay. A greater percentage of the flights (almost twice as many) to heavily traveled destinations is cancelled than the percentage of flights to less heavily traveled destinations. This bias is likely due to the ease of accommodating passengers on heavily traveled routes. If 5 flights fly the same route every day, canceling one of them still allows passengers to travel on one of the other four. But in the case of a single flight per day for a certain route, cancellation not only creates accommodation issues, but also produces much more passenger delay (24 as opposed to 3 or 4 hours).

Another airline method for dealing with heavily traveled routes is to schedule "wing tip" flights. These flights fly the same route but depart only a few minutes from each other. Wing tip flights allow airlines to accommodate passengers on both low and high booking days. When booking levels are high, both flights are operated.

When booking levels are low, however, the first flight of the two is cancelled and only the second flight is flown. These wing tip flights and their greater chance of cancellation contribute to the cancellation percentage observed on heavily traveled routes.

It was found that flights on heavily traveled routes (those with 4 or more flights to the same destination each day) are approximately twice as likely to be cancelled as those flying more sparsely traveled routes. These observations were collected from ASQP data from June 1998, concerning hubs at Chicago O'Hare, Dallas-Ft. Worth, and St. Louis airports.

The route probability is conditioned on the origin-destination and equipment pair. The probability is a value between 0 and 1 determined by the actual number of cancelled flights divided by the total number of flights to a destination in a single month. This value was not calculated in this thesis because the Swap and Cancellation Predictor was not used in the test environment for the DFM.

3.3.2 Conditioning on Incoming Delay

As with swap probability, cancellation probability increases with incoming delay. It should reach a value of 1 after approximately 12 hours (at which point the flight probably has been cancelled but has not been reported as such). The rate at which the probability of cancellation increases with delay should be determined from ASQP data from the same month of the preceding year. The value was not calculated in this thesis as the Swap and Cancellation Predictor was not included in the test environment for the DFM.

Chapter 4

Results

The Departure Forecasting Model was run on ADL data from March, April, and May of 1999. DFM predictions were generated and CDM predictions accessed for each departure flight matched with an incoming flight (CDM predictions are merely the CDM predicted departure times when the forecast is made). The ASQP wheels-off times were also recorded to determine the accuracy of the predictions by comparison.

Representative results are discussed in this chapter for each of the five studied airports: Chicago O'Hare International, Boston Logan International, Philadelphia International, Pittsburgh International, and New York La Guardia. More detailed results can be found in Appendix E. Results are presented and discussed by airport.

4.1 Predictions

Predictions are evaluated in this chapter using two metrics. They are mean and standard deviation of absolute error, comparing predictions to actual wheels-off times in ASQP data. These predictions are presented by month and day.

Five predictions were made every 15 minutes: a current flight list and a prediction of the flights leaving in one hour, two hours, four hours, and six hours in the future. Each prediction spans all the modeled flights in a fifteen-minute period. For example, when the ADL update for 1000 (10AM) is received, the model predicts departure times for four periods: 945-1000, 1045-1100, 1145-1200, 1345-1400, and 1545-1600. All times are in local time for their specific airports. The Departure Forecasting Model updates departure times for each ADL update it receives. When these departure times fall within the forecasting periods, they are included in the forecasts.

Airports were chosen to maximize domestic departures and to include the hubs of major airlines. Boston Logan was added due to its proximity to MIT.

4.1.1 Model Coverage

Predictions and their error with respect to ASQP data are only included for those flights whose turns could be matched. Due to the limited number of ADL-included airports for the study period, a subset of the total flights for an airport was examined. As discussed in Chapter 2, ADLs contain arrival information only. Many departures are not found in the ADL data set because their arrival airports are not included.

To use operations at O'Hare as an example, approximately 800 jet aircraft land and take off in a 24-hour period, for a total of 1600 flights. A usual day of ADL arrival data will contain approximately 25000 updates concerning jet aircraft, and the list of departures compiled from the other airport arrival lists usually contains about 10000 updates. Approximately 500 of the 800 departing flights are matched from the departure list to the arrival list via ASQP, and the reasons for 300 missing flights are two-fold. Only certain major airlines and airports are contained in the ASQP data, and only certain airports are contained in the ADL arrival lists. Departures operated by regional carriers and those that serve international cities are two sets of departures that are not contained in the modeled data.

Only those flights that could be matched were examined and input to the model. Though only a subset of flights was examined, the subset contains well over 50% of the jet operations for each modeled airport. Similar results should be attainable if the model were applied to all jet flights.

Though the Departure Forecasting Model does not take all flights into account, the results for the flights it does model are valid. Additional flights could cause congestion on the ground and in the air, but CDM updates reflect this congestion in the delay felt by modeled flights. In addition, as discussed in Chapter 3, aggregate delay does not significantly affect turn times. Therefore, the operations of those flights not included in the DFM should not significantly affect its results.

4.2 Correlation of Improvement and Delay

As will be seen in the following sections, the greatest improvement in departure time prediction was achieved on heavily delayed days. As the average delay of incoming flights increased, so did the accuracy of the DFM predictions when compared to CDM data. Figure 4-1 illustrates this correlation by showing the average delay by day for March, 1999 at Chicago O'Hare. When comparing this figure with the predictions in the next section (Figure 4-3), it is observed that the three most heavily delayed days, March 6th, 9th, and 17th, are also those days in which the DFM most improved on CDM predictions.

The same trend was found throughout the results, but is not reported here for space reasons. The DFM outperformed CDM by the greatest degree on those days exhibiting the most significant delay.

4.3 Prediction Distribution

In the following figure (Figure 4-2), the distribution of predictions for both the DFM and CDM are presented for March 6th, 1999, at Chicago O'Hare. This day was chosen because the DFM produced much better predictions than CDM for this day. The one noticeable trend in the figure is that most of the improvement achieved by the DFM occurs in the afternoon. This trend was also seen on the other days in which the DFM produced a marked improvement. Delay usually increases throughout the day, as early delays are propagated to subsequent flights. Thus, it is not surprising that the DFM provides the most improvement late in the day.

Afternoon improvement was found to a varying degree in all results, in much the same manner as in Figure 4-2.

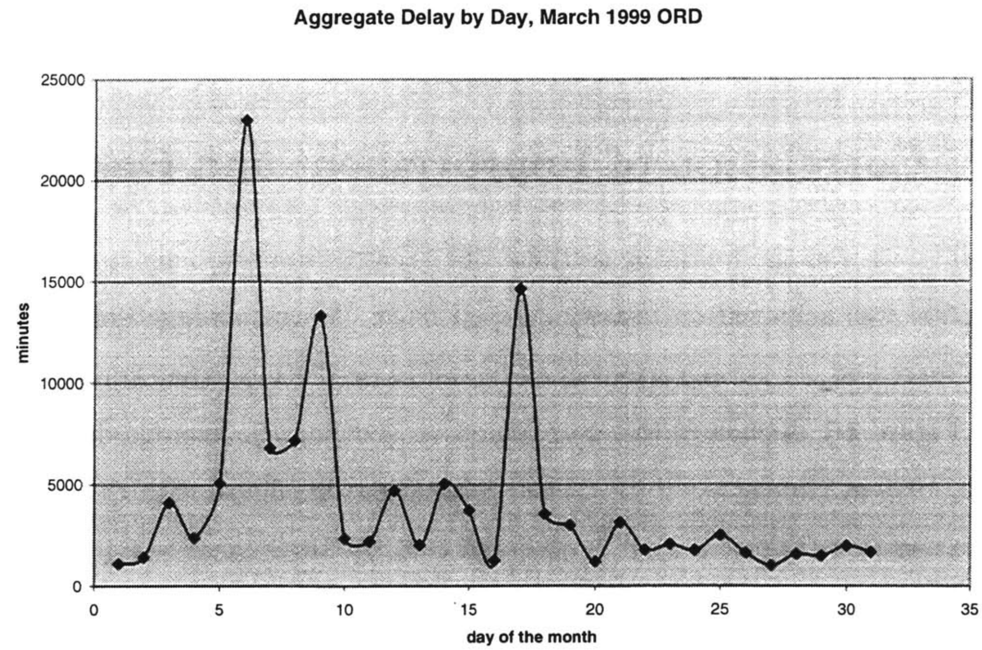

Figure 4-1: ORD Average Delay by day for March, 1999

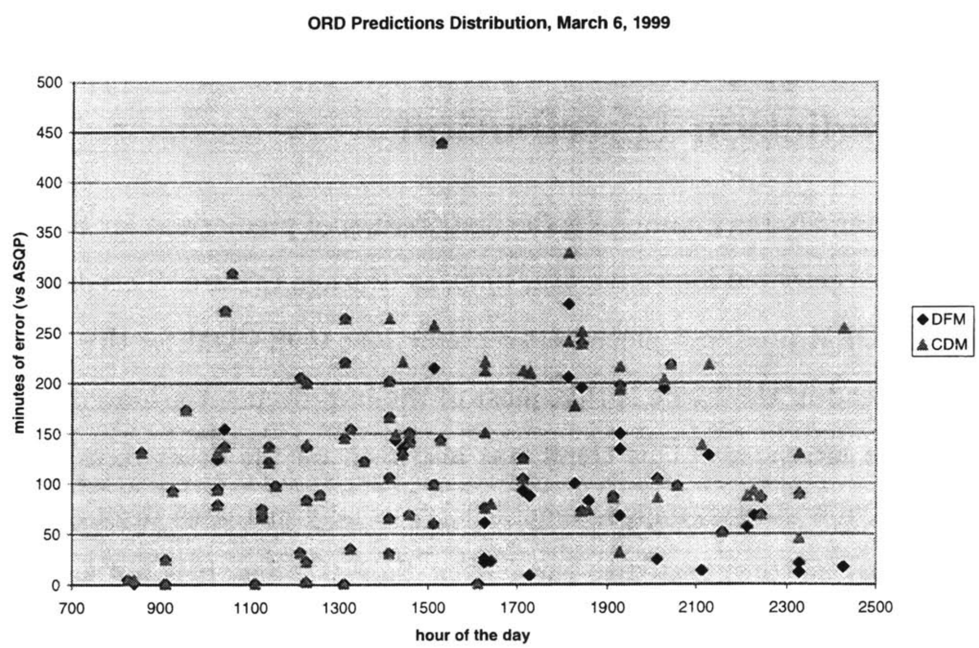

Figure 4-2: Distribution of Absolute Prediction Error for 4 Hour Forecast for March, 1999. Prediction error for the DFM and CDM is displayed for each flight by the time of day.

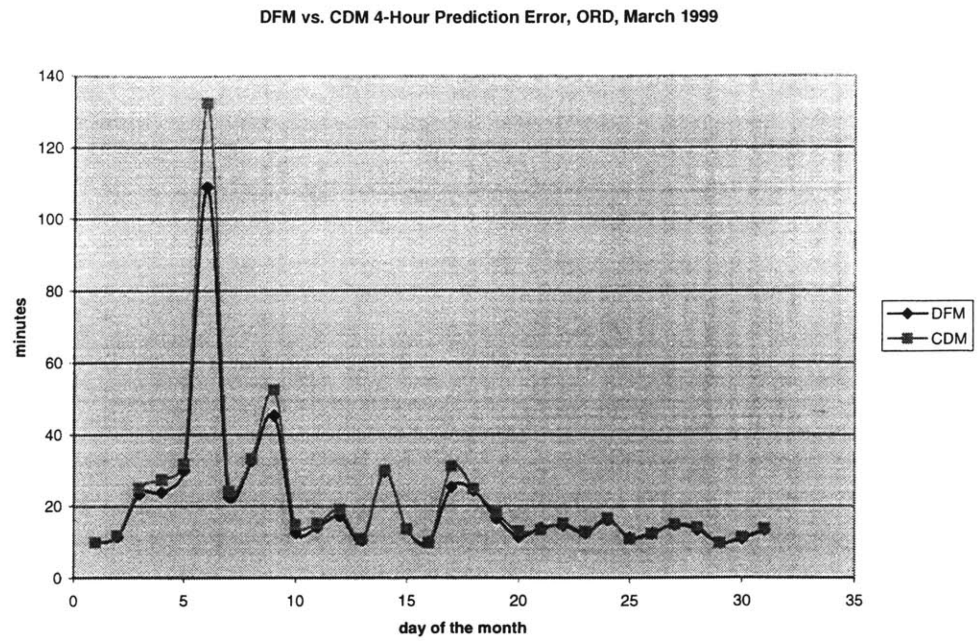

Figure 4-3: ORD Average Absolute Error for 4 Hour Forecast by day for March, 1999

4.4 Chicago O'Hare International (ORD)

At Chicago O'Hare, the greatest improvement over CDM predictions was achieved with the 4-hour forecast. Average absolute error by day is displayed in the next three figures (4-3, 4-4, 4-5), one for each month of the study.

As can easily be seen from the figures, the DFM follows CDM predictions closely, except for a few days each month. On those days, however, the model achieves an average absolute error reduction of 10-30% when compared to the CDM values.

Additionally, the DFM is less accurate than the CDM predictions on only one or two days each month, and on those days is only slightly less accurate. The model also, as can be seen in the next three figures, rarely has a greater standard deviation of error than the CDM predictions.

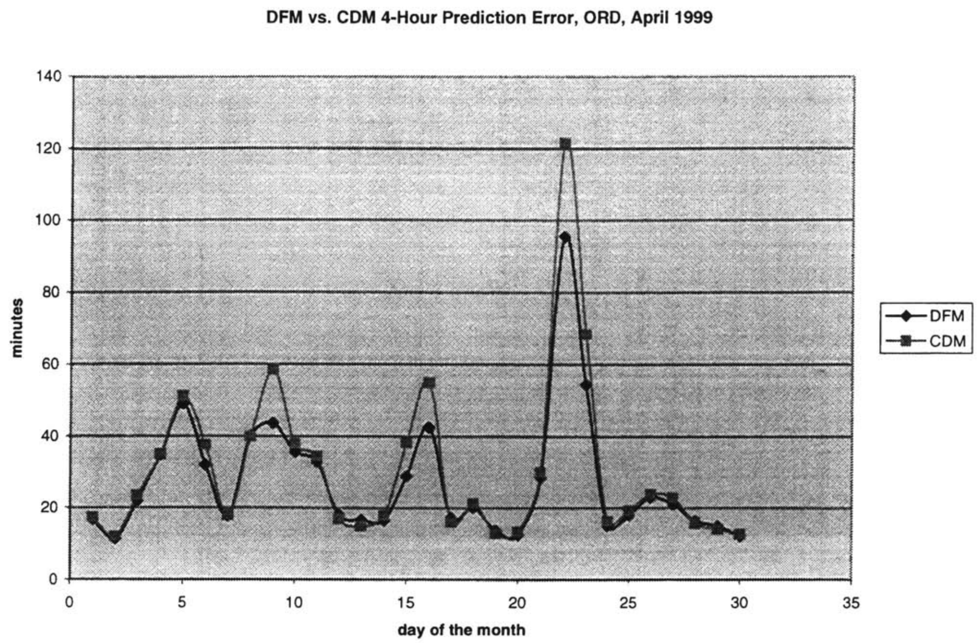

Figure 4-4: ORD Average Absolute Error for 4 Hour Forecast by day for April, 1999

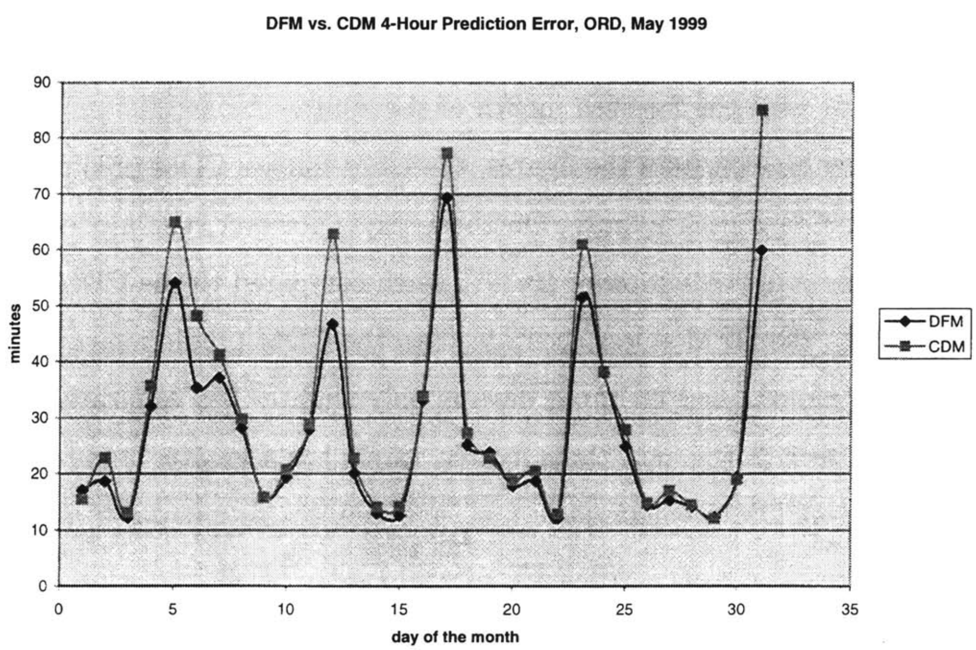

Figure 4-5: ORD Average Absolute Error for 4 Hour Forecast by day for May, 1999

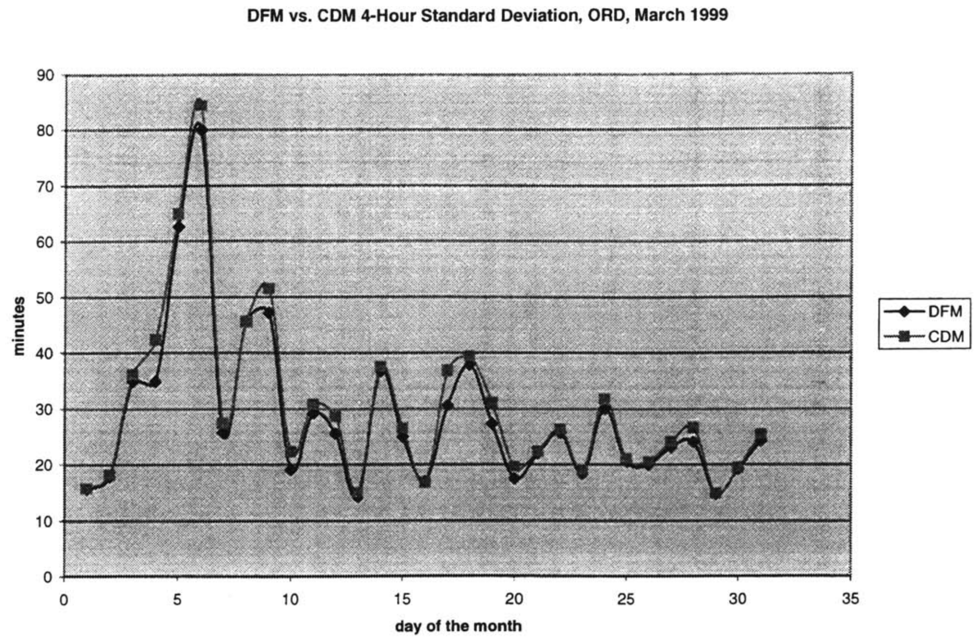

Figure 4-6: ORD Standard Deviation of Error for 4 Hour Forecast by day for March, 1999

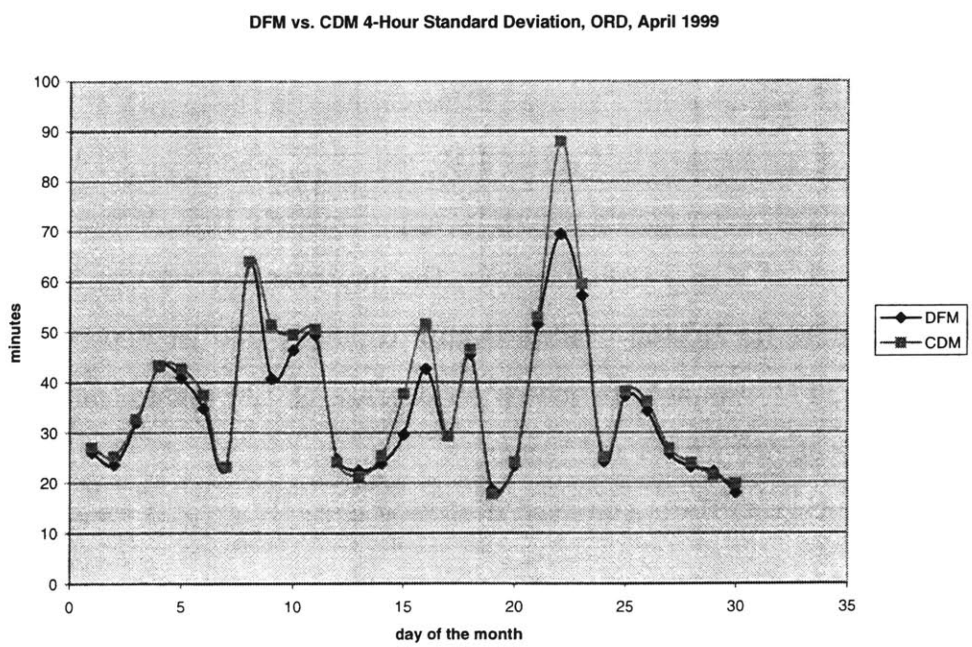

Figure 4-7: ORD Standard Deviation of Error for 4 Hour Forecast by day for April, 1999

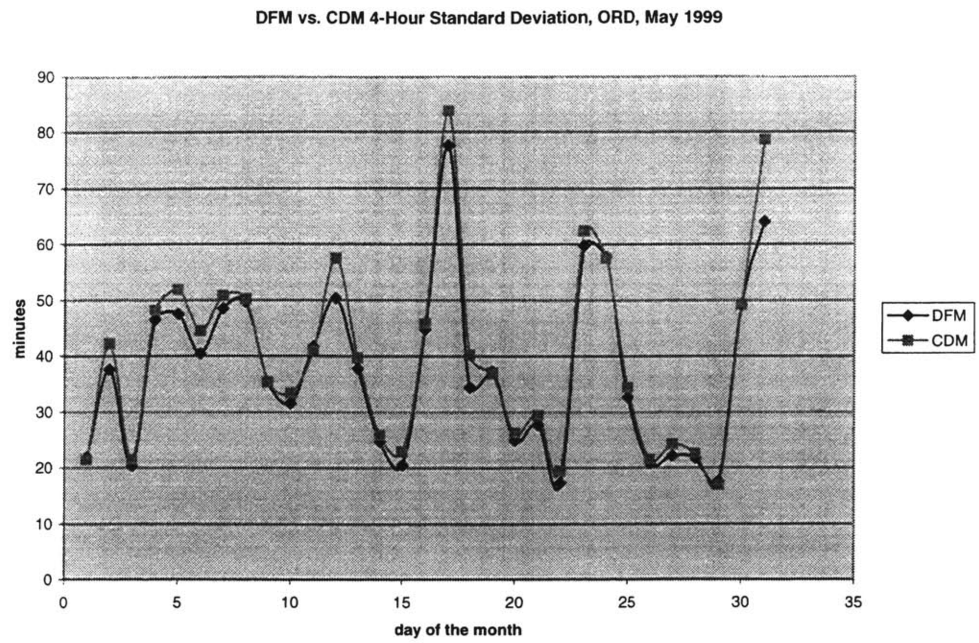

Figure 4-8: ORD Standard Deviation of Error for 4 Hour Forecast by day for May, 1999

4.5 Boston Logan International (BOS)

The results obtained from the Departure Forecasting Model are not as promising for Boston Logan as they were for Chicago O'Hare, except in the month of March (Figure 4-9). This trend can be seen in Figures 4-10, and 4-11, where the DFM outperforms CDM predictions on a few days a month by approximately 10%. When the DFM was run on the month of March 1999, however, the improvement reached 35%.

The ADL data for Boston contained approximately 230 jet departures per day. Of these, the DFM was able to match an average of 120 flights, on par with the percentage of matched flights at Chicago.

As with Chicago O'Hare, standard deviation is rarely greater for the DFM than it is for CDM.

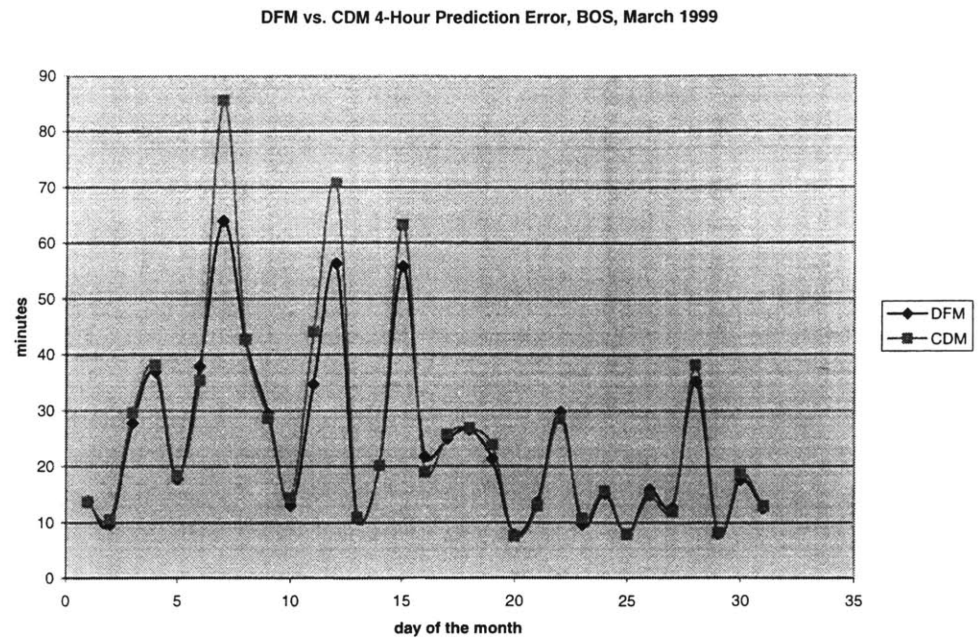

Figure 4-9: BOS Average Absolute Error for 4 Hour Forecast by day for March, 1999

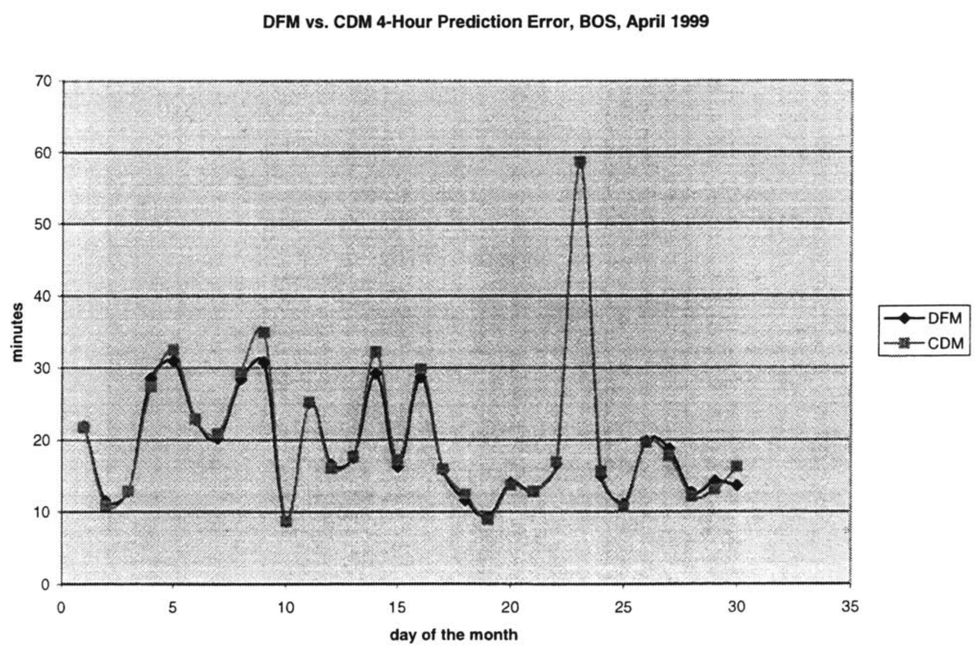

Figure 4-10: BOS Average Absolute Error for 4 Hour Forecast by day for April, 1999

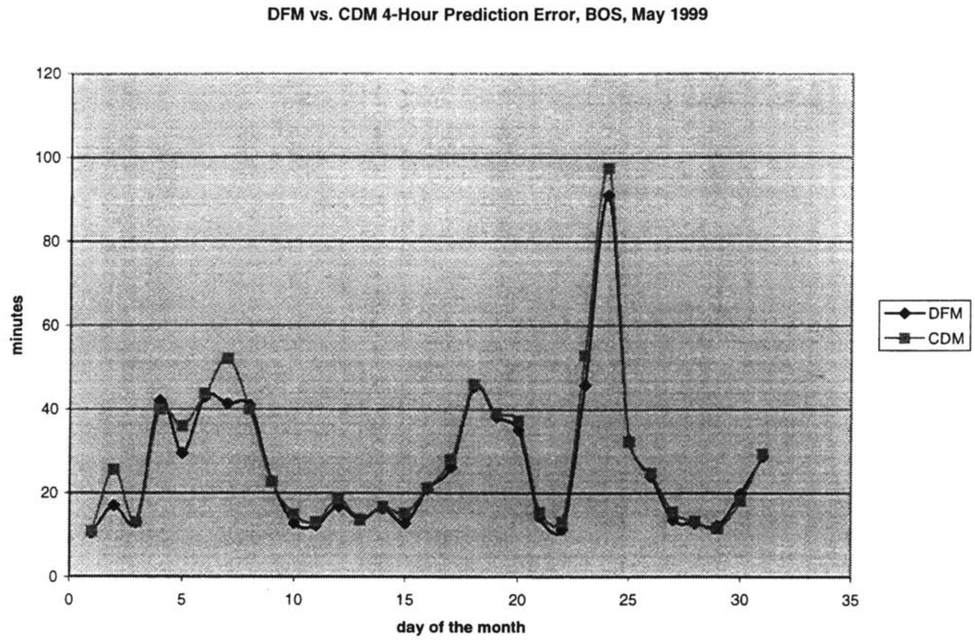

Figure 4-11: BOS Average Absolute Error for 4 Hour Forecast by day for May, 1999

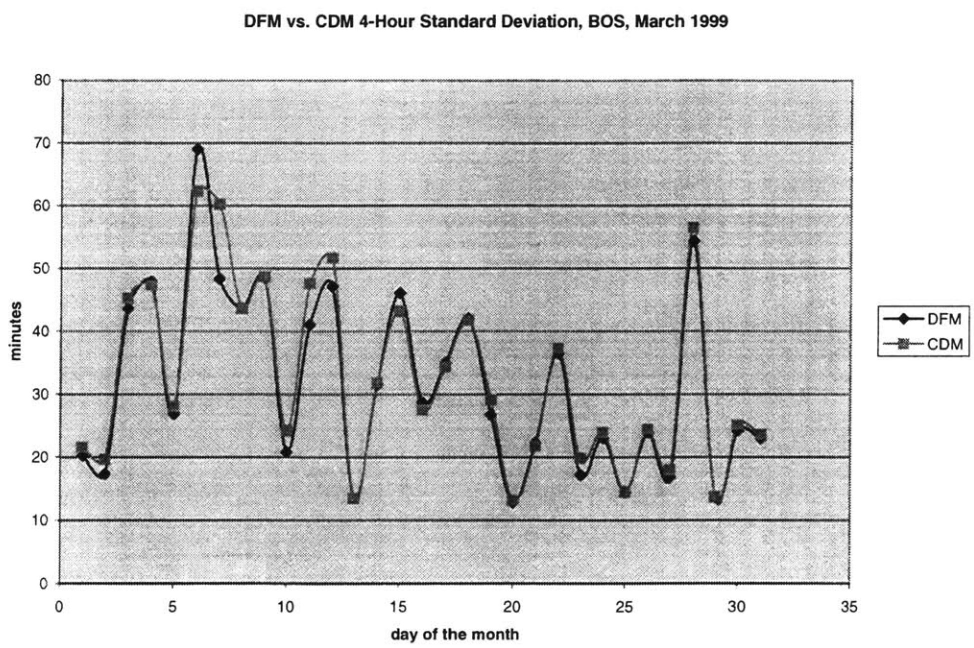

Figure 4-12: BOS Standard Deviation of Error for 4 Hour Forecast by day for March, 1999

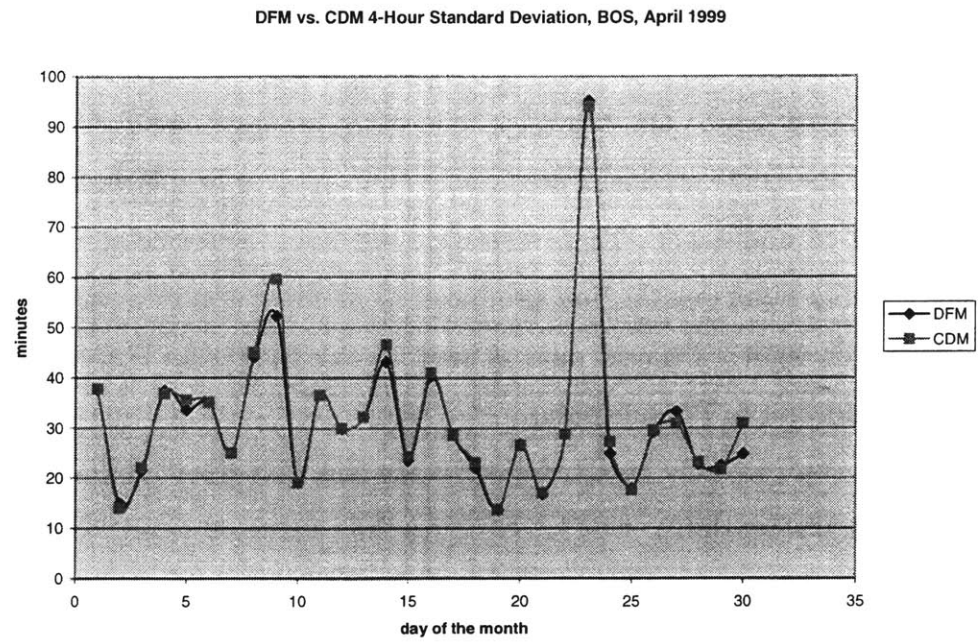

Figure 4-13: BOS Standard Deviation of Error for 4 Hour Forecast by day for April, 1999

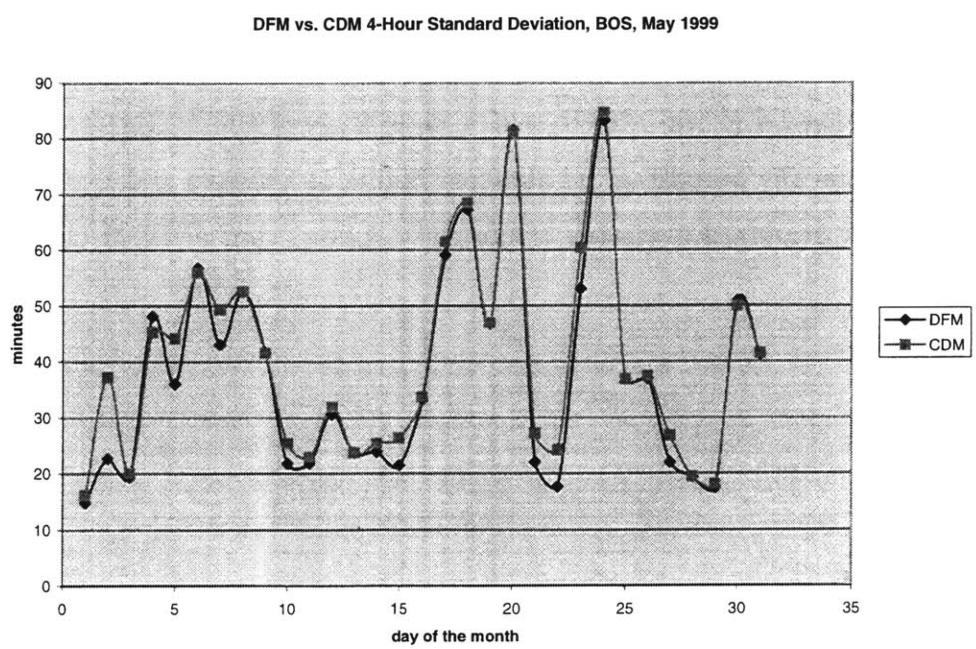

Figure 4-14: BOS Standard Deviation of Error for 4 Hour Forecast by day for May, 1999

4.6 Philadelphia International (PHL)

Philadelphia showed similar results to Chicago O'Hare, though not as pronounced on heavily delayed days. On two days, one early in April and the other early in May, however, the DFM produced predictions that were up to 30% better than CDM (see Figures 4-16 and 4-17). These two days did not experience as much delay as other days in the same months, but produced much more accurate predictions. This improvement on lightly delayed days is seen in the results for Pittsburgh and La Guardia, in addition to Philadelphia.

This effect can probably be attributed to the fact that there are fewer operations for Pittsburgh, Philadelphia, and La Guardia than for Boston and Chicago. Pittsburgh, Philadelphia, and La Guardia operate fewer flights, and therefore have fewer flights in their ADL updates. In addition, these three airports had a smaller percentage of flights matched through ASQP. Both effects lead to fewer modeled flights overall. If a single departure were to be predicted more accurately, the effect on the average error would be more considerable than at Boston or Chicago, where there are more data points.

The ADL data for Philadelphia contained approximately 200 jet departures per day. Of these, the DFM was able to match an average of 100 flights, for a percentage slightly less than the percentage of matched flights at Chicago and Boston.

Once again, standard deviation of the DFM follows that of CDM closely.

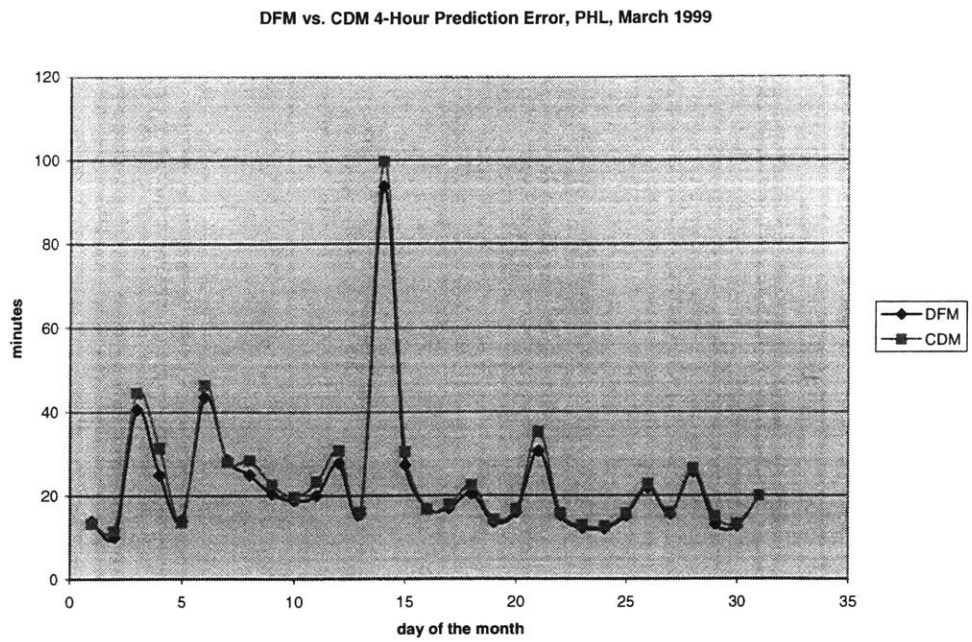

Figure 4-15: PHL Average Absolute Error for 4 Hour Forecast by day for March, 1999

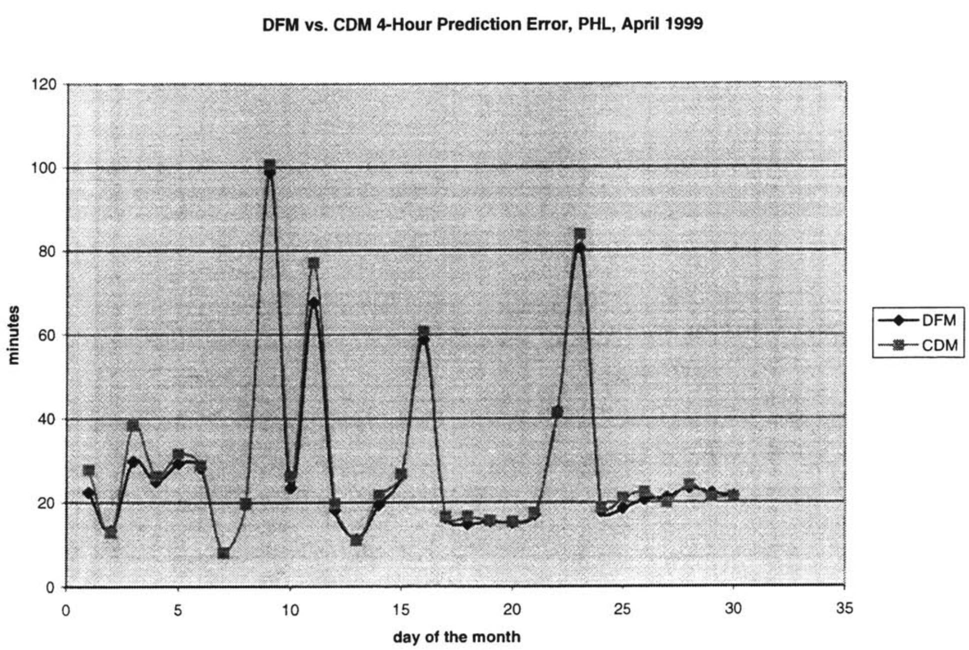

Figure 4-16: PHL Average Absolute Error for 4 Hour Forecast by day for April, 1999

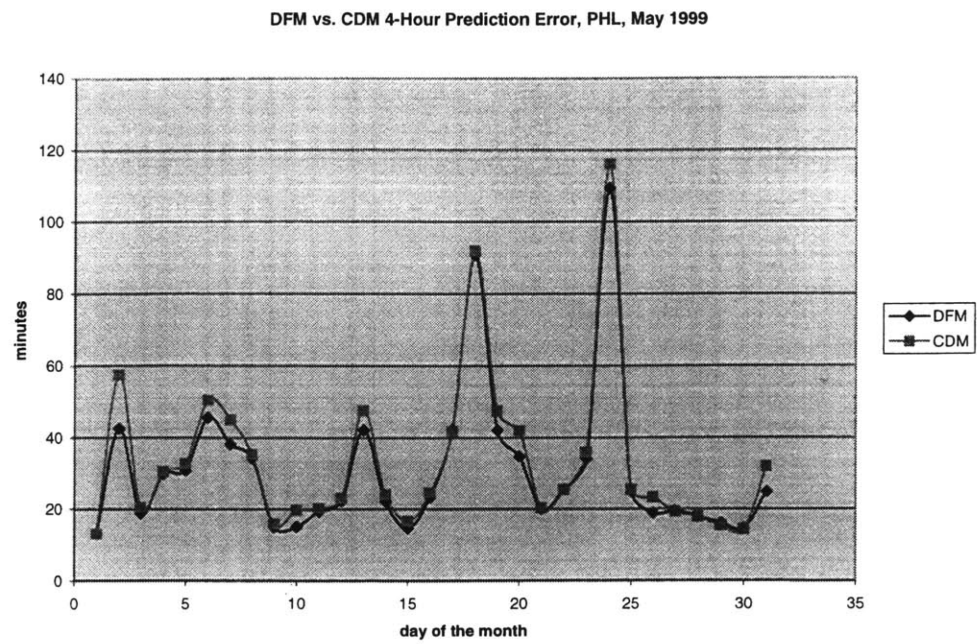

Figure 4-17: PHL Average Absolute Error for 4 Hour Forecast by day for May, 1999

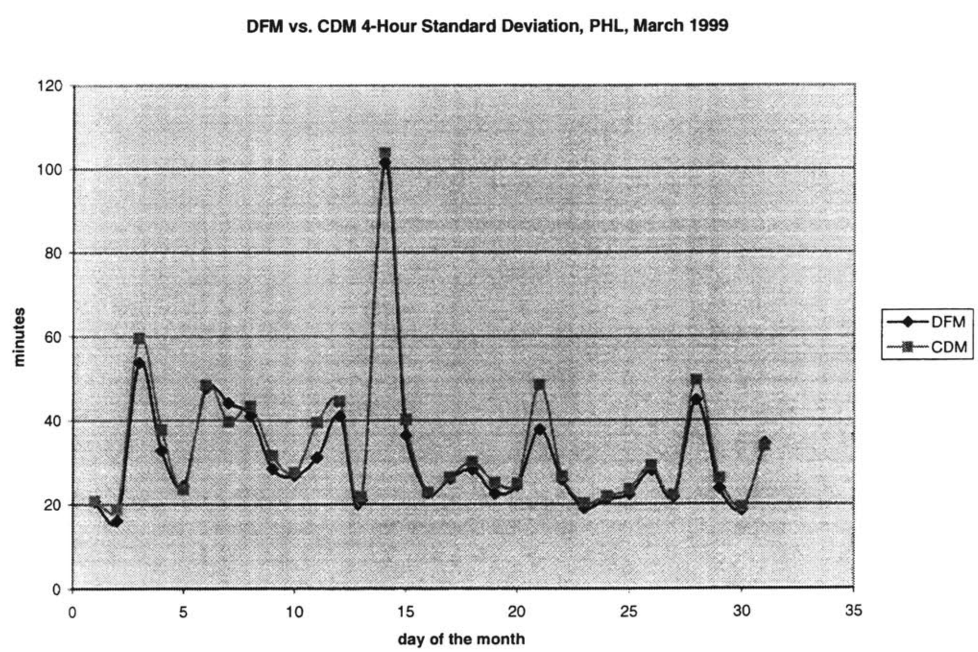

Figure 4-18: PHL Standard Deviation of Error for 4 Hour Forecast by day for March, 1999

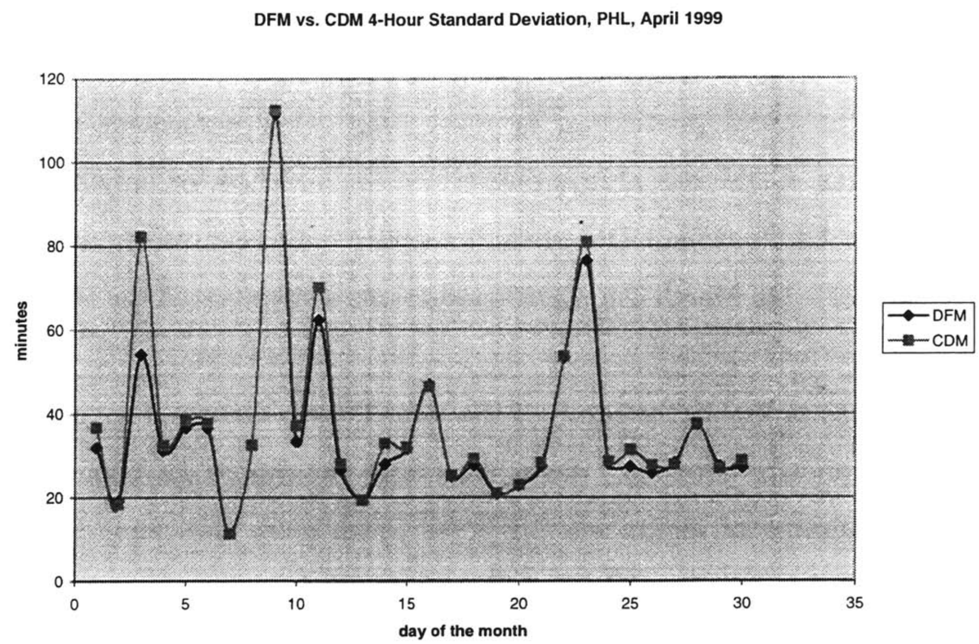

Figure 4-19: PHL Standard Deviation of Error for 4 Hour Forecast by day for April, 1999

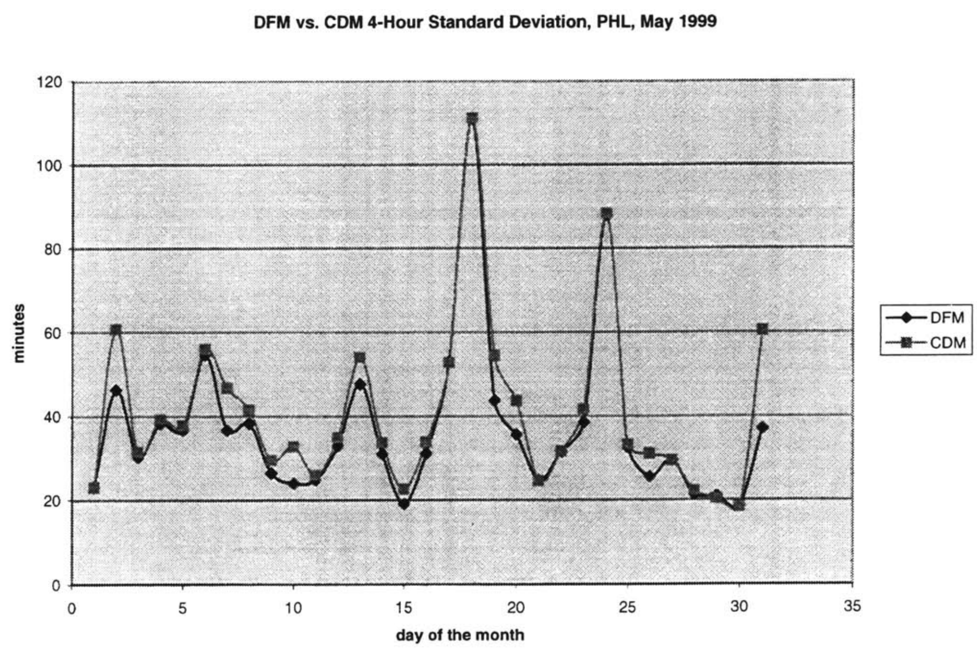

Figure 4-20: PHL Standard Deviation of Error for 4 Hour Forecast by day for May, 1999

4.7 New York La Guardia (LGA)

The results for La Guardia stand out from the other airports because the DFM was outperformed by CDM predictions by more than a few percentage points. In all other DFM runs, CDM was always outperformed, matched, or only slightly underperformed. For La Guardia, CDM predictions were more accurate on one day: March 4th (Figure 4-21). On March 4th, CDM predictions outperformed the DFM by almost 25%.

This day was probably the result of the same factors discussed in the last section, but with an opposite outcome. When the DFM was run on La Guardia on March 4th, a small number of flights was predicted much more accurately by CDM than the DFM. Since La Guardia had fewer flights overall that were predicted, a relatively small number of flights could affect the error averages considerably.

The ADL data for La Guardia contained approximately 200 jet departures per day. Of these, the DFM was able to match an average of 100 flights, for a percentage slightly less than the percentage of matched flights at Chicago and Boston.

Again, standard deviations of the DFM predictions are very close to those of CDM predictions.

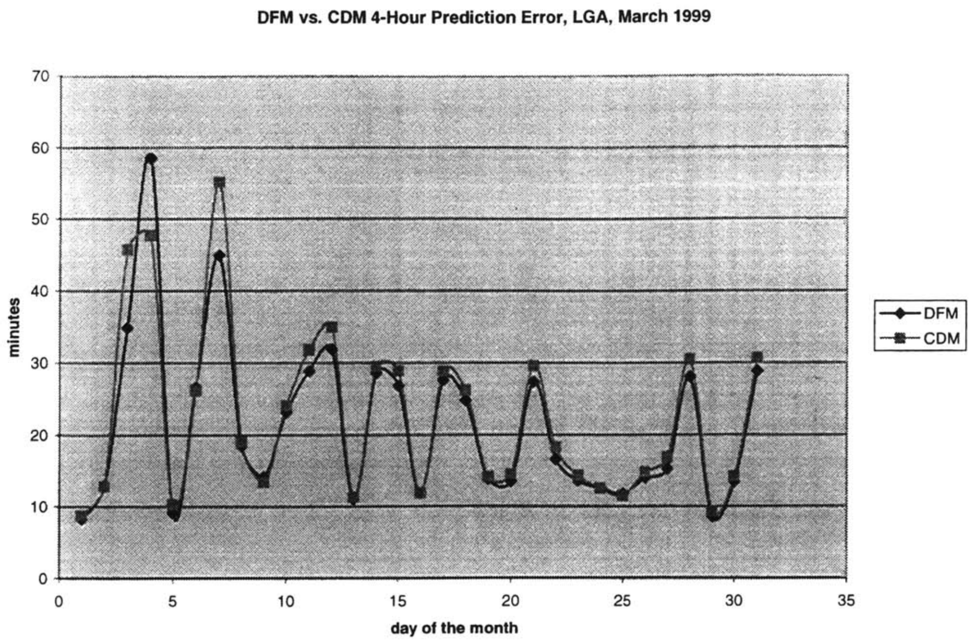

Figure 4-21: LGA Average Absolute Error for 4 Hour Forecast by day for March, 1999

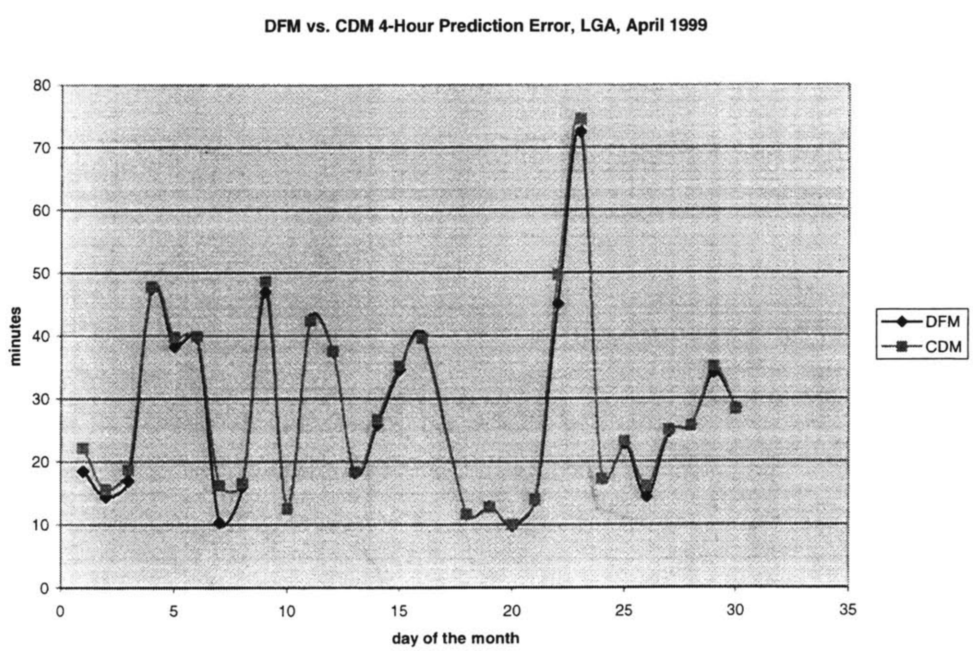

Figure 4-22: LGA Average Absolute Error for 4 Hour Forecast by day for April, 1999

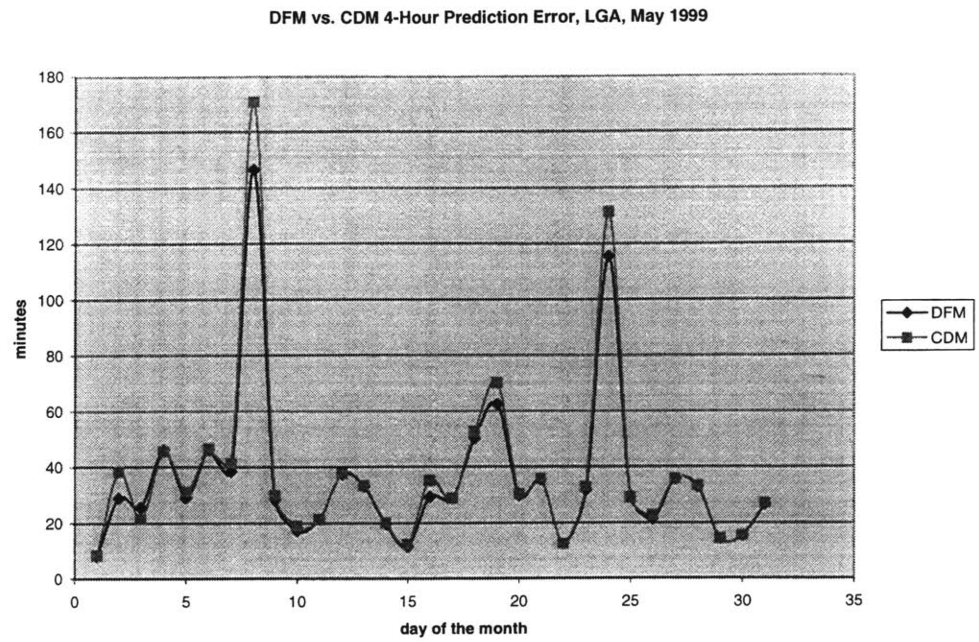

Figure 4-23: LGA Average Absolute Error for 4 Hour Forecast by day for May, 1999

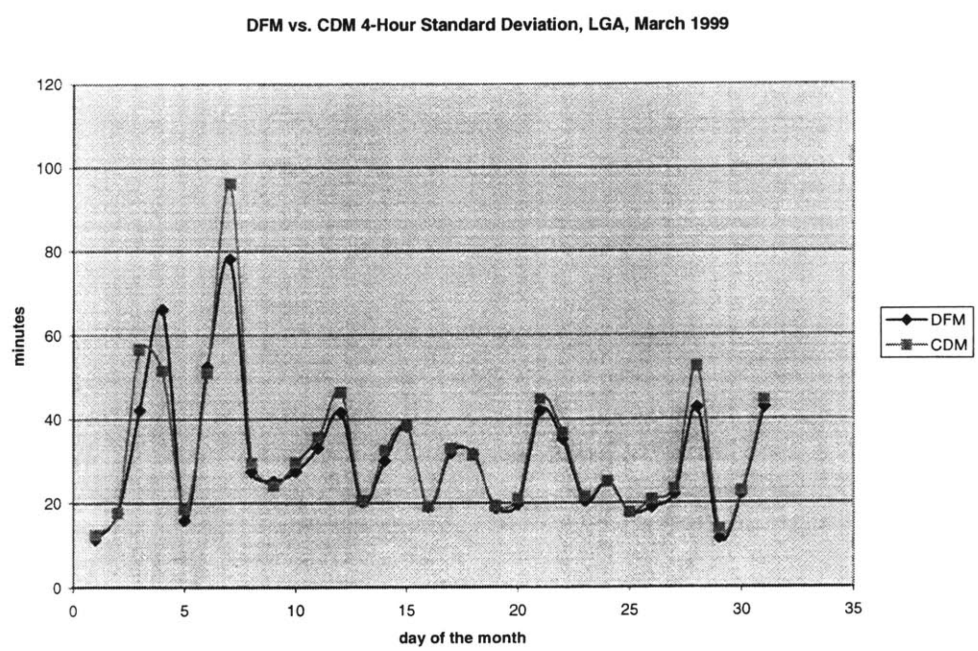

Figure 4-24: LGA Standard Deviation of Error for 4 Hour Forecast by day for March, 1999

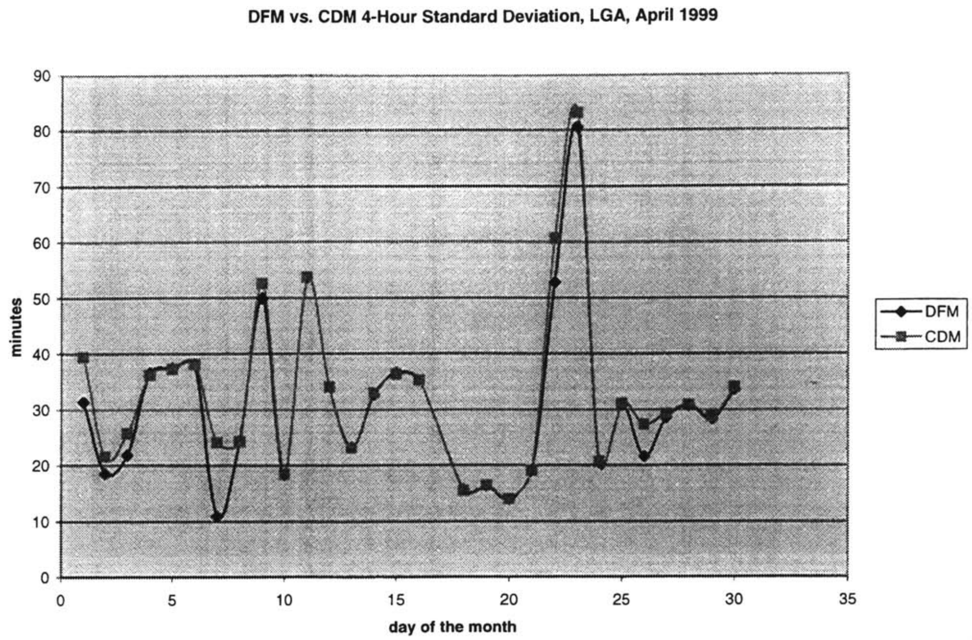

Figure 4-25: LGA Standard Deviation of Error for 4 Hour Forecast by day for April, 1999

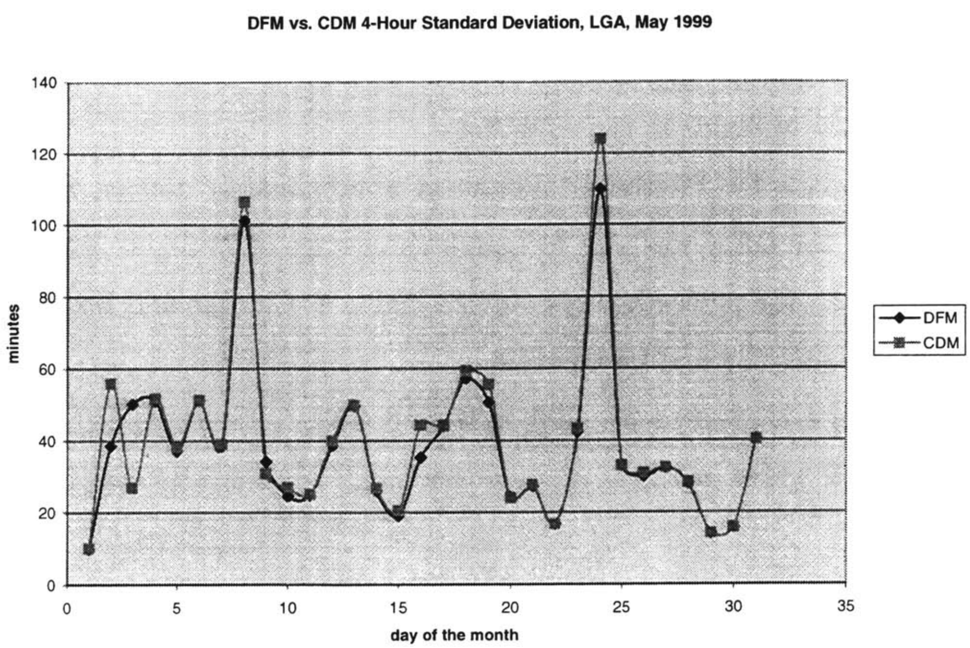

Figure 4-26: LGA Standard Deviation of Error for 4 Hour Forecast by day for May, 1999

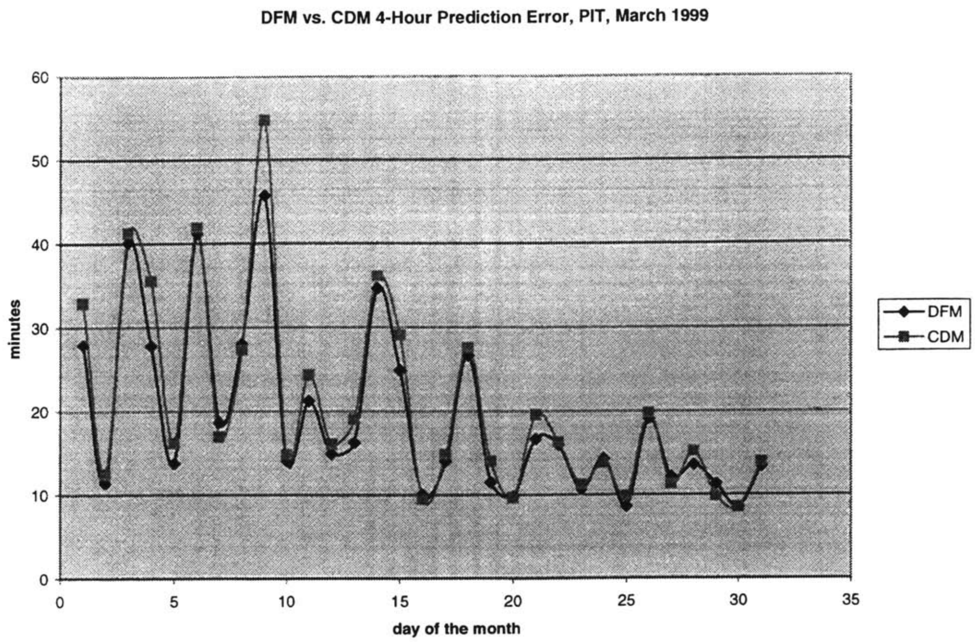

Figure 4-27: PIT Average Absolute Error for 4 Hour Forecast by day for March, 1999

4.8 Pittsburgh International Airport (PIT)

The DFM run on Pittsburgh produced very little difference from CDM predictions, except for a very small number of days. Pittsburgh was also the airport with the smallest number of flights in the ADL data.

The ADL data for Pittsburgh contained approximately 150 jet departures per day. Of these, the DFM was able to match an average of 75 flights, for a percentage slightly less than the percentage of matched flights at Chicago and Boston, but equal to that of Philadelphia and La Guardia.

See Figures 4-27, 4-28, 4-29, 4-30, 4-31, and 4-32.

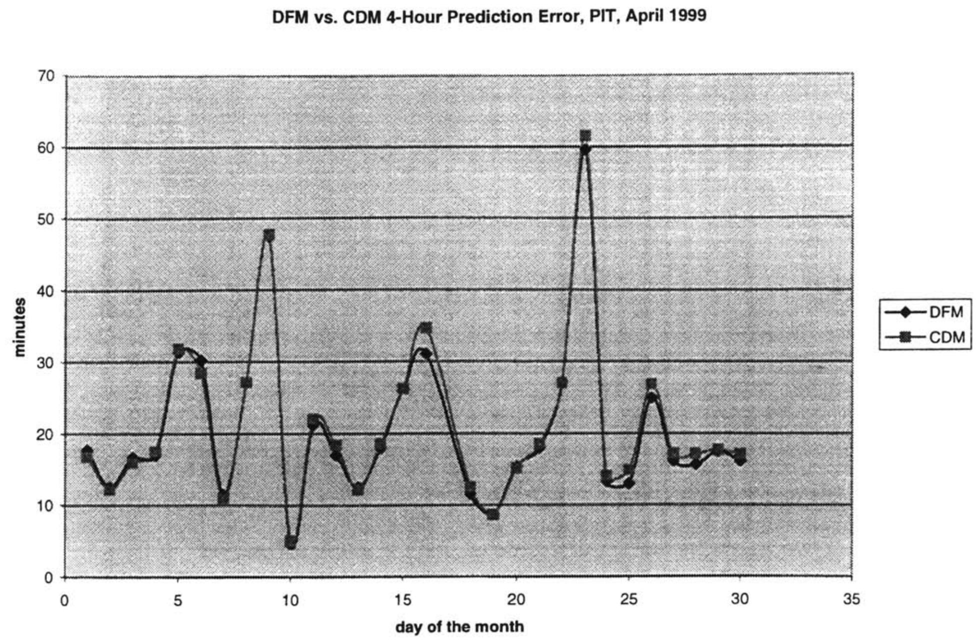

Figure 4-28: PIT Average Absolute Error for 4 Hour Forecast by day for April, 1999

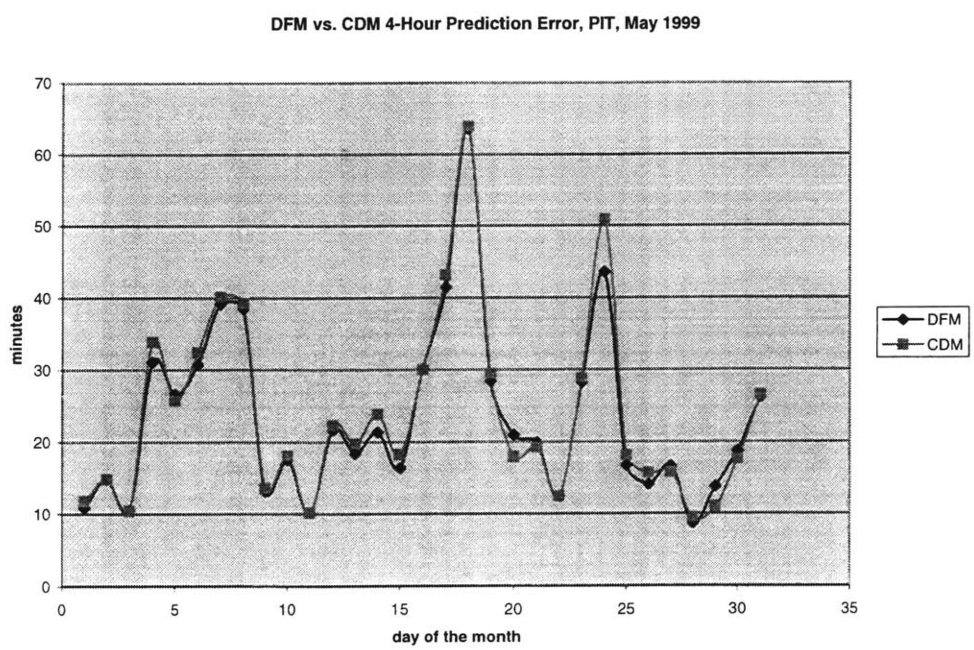

Figure 4-29: PIT Average Absolute Error for 4 Hour Forecast by day for May, 1999

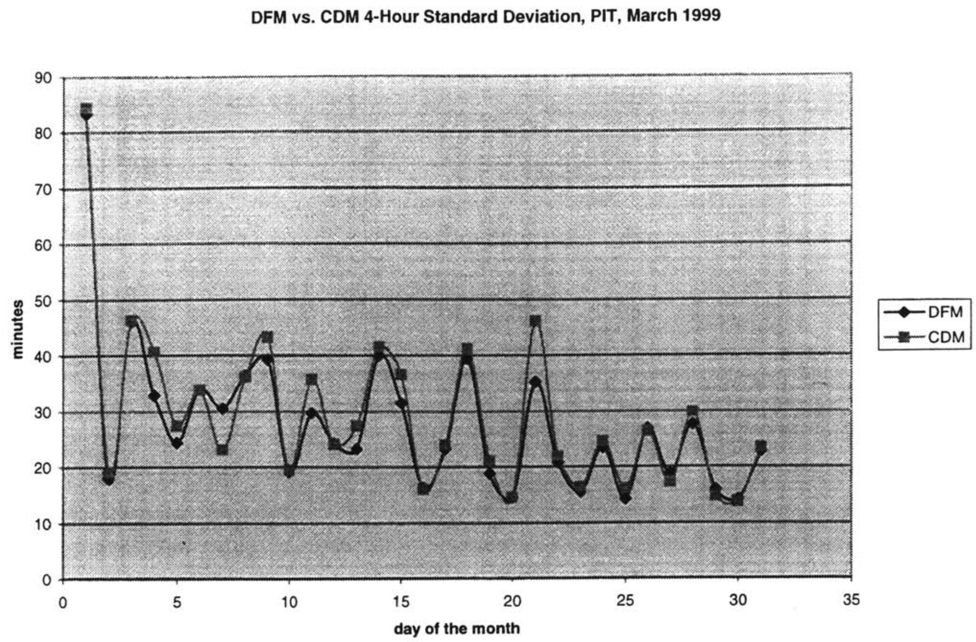

Figure 4-30: PIT Standard Deviation of Error for 4 Hour Forecast by day for March, 1999

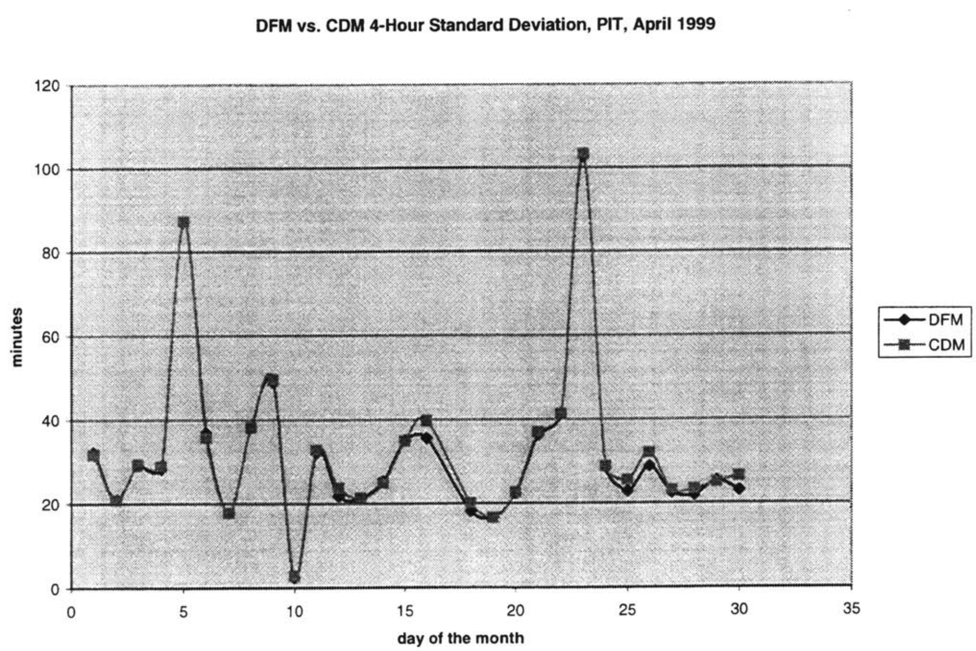

Figure 4-31: PIT Standard Deviation of Error for 4 Hour Forecast by day for April, 1999

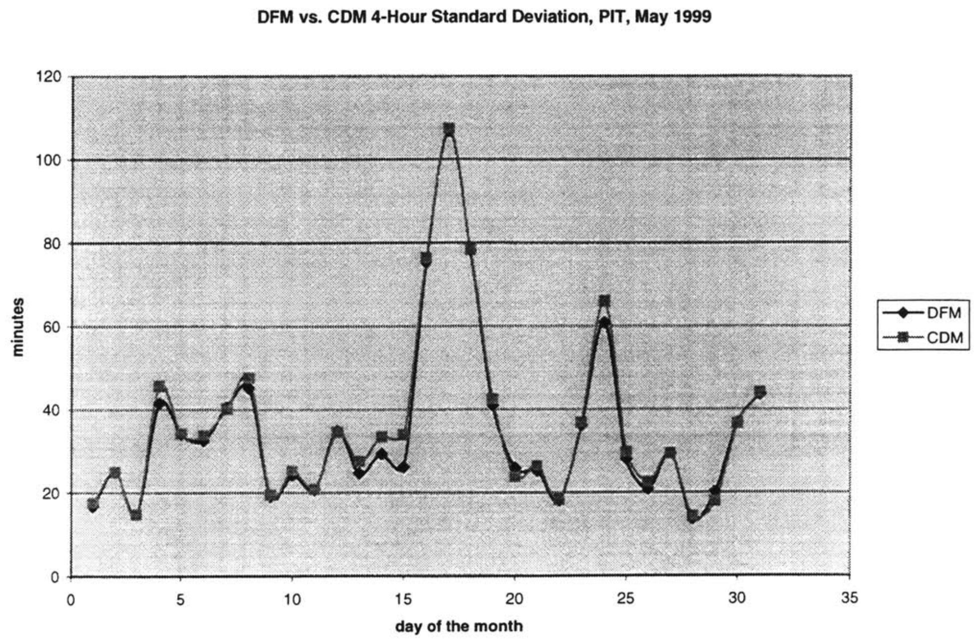

Figure 4-32: PIT Standard Deviation of Error for 4 Hour Forecast by day for May, 1999

Chapter 5

Contributions and Future Work

5.1 Contributions

The principal contribution of this thesis is a method of departure prediction which, by modeling aircraft inventory, is able to reduce average absolute error by up to 30% when compared to current CDM techniques. This reduction of error is an aggregate reduction, averaged over all the flight predictions on one day. This improvement is achieved while an airport is experiencing large amounts of delay, precisely the situation in which current predictions are inaccurate and an increase in accuracy may have the greatest operational benefit.

The Departure Forecasting Model is significantly more accurate than CDM for only a few days each month. On these days, airports experienced system-wide delays, where the reporting method of CDM seems to have broken down. Unfortunately, as the number of flights in our nation's airspace increases, those days where delays are felt across the system will become increasingly frequent. However, those days where delays are rampant are the days when the DFM is of the most help. Therefore, it is likely that the DFM presented in this thesis, or other models like it, will become more and more useful in the coming years.

As discussed in Chapter 4, the DFM predicts approximately 60% of all jet operations in its present state. Unfortunately, at some airports, such as Pittsburgh International, jet aircraft only make up 54% of total operations, resulting in an approximate model coverage of 30% of all operations. However, the DFM design can be expanded to include all operations for a given airport, if all operations are contained in ADL updates. DFM was not expanded to include all operations because ADL updates did not yet contain departure information, which would be necessary for a more complete implementation.

It is assumed that a certain amount of accuracy will be lost when the Departure Forecasting Model is run without using ASQP data for matching departures to arrivals. Some flights will be mis-matched by the DFM, leading to inaccuracy on a per-flight basis. However, it is believed that error averages would still be significantly lower using the model than using CDM alone. In addition, even if accuracy on a per-flight basis is somewhat reduced (compared to actual departure times) by an expanded implementation of the DFM, forecast departure demand queues may still be more accurate than CDM-predicted queues. In both cases, though per-flight accuracy may suffer, aggregate values should yet provide an improvement.

Since ASQP data were used in the test environment to match arriving aircraft to departing aircraft, part of the design of the DFM, the Swap and Cancellation Predictor, was not implemented. If the DFM design were to be realized in the current ATC system, ASQP data would not be available, and the Swap and Cancellation Predictor would be a necessary module of the DFM. An additional contribution of this thesis is a top-level design of the Predictor and suggestions for its function.

5.2 Future Work

A promising step to take in expanding this thesis work is to generalize the Departure Forecasting Model test environment to run on all jet operations, not just those turns contained in ASQP data. Though ASQP data should still be used to determine minimum turn times, the DFM implemented in the test environment will never be implementable in a real-time situation for which ASQP data do not yet exist. Other matching techniques could be used, such as a matching by the Aircraft Type field in the ADL data or a matching by FIFO queues (as mentioned in Chapter 2).

The second step is to further generalize the DFM to include all operations. Determining minimum turn times for these non-jet aircraft will be more difficult (as ASQP data do not exist for them), but since these aircraft are usually small with a minimal crew complement, turn times could be approximated to 10 or 15 minutes. Additionally, many non-jet aircraft do not follow the tight schedule constraints of jet aircraft.

Regardless of the next step, ADL update information, as of the publication of this thesis, has been modified to include departure information. One of the areas of information which could have improved the DFM coverage was ADL departure information. As mention in Chapter 4, ADL arrival updates outnumbered departure updates 2.5:1. With the ADL modification, a one-to-one matching of departure flights to arrival flights is possible.

Another continuation of this work could be the application of its principles to airline operations, as opposed to ATC operations. Airlines could use delay propagation in order to improve the timeliness of their announcements to both the FAA and passengers. In addition, airlines could observe their minimum turn times reported in this work and decide whether they wish to decrease the amount of time necessary to turn their aircraft. Though incoming delay affects departing delay considerably, smaller turn times serve to reduce delay throughout the system.

Finally, a grandiose picture of the future has the Departure Forecasting Model working at each major airport in the United States. In this case, delay could be propagated, not just though one turn, but through multiple turns at multiple airports. Departure prediction, in this case, would overlap into arrival prediction, as delay propagates throughout the system.

Appendix A

Typical Flight Departure Process

The following list enumerates the activities that usually occur when a plane is turned at the gate. The airline employee responsible for each action is listed in parentheses. Notice the complexity, not only of actions, but also concerning the number of different employees that must coordinate their actions.

The following list appears in Jody Gittell's 1995 PHD thesis [4]:

1. Connect skybridge to aircraft, open door, help passengers disembark (gate agent)
2. Assist passengers with special needs (gate agent)
3. Clean the plane, do magazines, headsets (ramp workers; gate agent, maybe flight attendant)
4. Load food onto plane (caterers; ramp workers)
5. Empty bathrooms (ramp workers)
6. Check weather, flight conditions, determine fuel needed (ops agent; central dispatch)
7. Refuel (ramp workers; ops agent)
8. Maintenance check (mechanics; captain, upline/downline mechanics, ops agent)

9. Equipment check, cockpit (captain, co-pilot)
10. Assist arriving passengers with their bags (Skycaps)
11. Direct passengers to ticket counter, gates (Skycaps, ticket agents)
12. Sell tickets, check baggage, maybe reroute passengers (ticket agents; baggage room)
13. Reorganize passenger lines to check in those who are leaving soon (ticket agents)
14. Send baggage downstairs (ticket agents; baggage room)
15. Accept priority mail (ticket agents, lead agents; baggage room)
16. Accept freight from outside vendors (freight handlers; ramp area, ops agent)
17. Sort and load baggage (baggage handlers; freight agents, ops agent)
18. Sort and load mail and freight (freight handlers; baggage handlers, ops agent)
19. Compute weight and balance of freight, baggage, and mail. Give to ops agent (Ramp agents; freight and baggage agents, ops agent)
20. Check people at entrance to gate area (security)
21. Check passenger seating, decide whether to add passengers or freight (gate agents; ticket, ops agents, flight attendants)
22. Call in number of passengers to ops agent (gate agents; ops agents, flight attendants)
23. Compute weight, balance of freight, mail, passengers, fuel (ops agent; everybody)
24. Assist passengers (gate agents; ticket agents, gate and ticket agents at next destination)

25. Load passengers, check carryon luggage, if necessary (gate agents, flight attendants)
26. Seat passengers, stow carryon luggage (flight attendants)
27. Serve refreshments to first class (flight attendants)
28. Prepare passengers for takeoff, make destination announcement (flight attendant, captain)
29. Give weather, load, and fuel information to captain (gate agent; ops agent, captain)
30. Close plane doors, dispatch flight (gate agents; ramp agents, captain)
31. Signal readiness to tower (captain)
32. Turn on engines (captain)
33. Pushback from gate (ramp workers; captain, control tower)

Appendix B

Turn1 Model Justifications

A few issues concerning turn time collection and averaging to obtain values for Turn1 were examined but not used in the data collection. The factors and the reasons for excluding them appear below.

B.1 Turn Time and Aggregate Airport Delay

One of the issues examined was that turn times were affected by overall airport delay. It was assumed that operations on those days with a large number of delayed flights would be slower than normal operations. However, a careful examination of the data revealed that average delay and the number of delayed flights have very little correlation. See Figure B-1.

The lack of correlation makes sense except in one case, when delays are such that too many planes are on the ground at one time. At this point, the demand for airline employees could be too great to turn all the planes at once. But in most circumstances, the conditions that cause delays would tend to lead to a decrease on the number of turns being performed. These circumstances do not lead to a decrease in turn time because only one crew is assigned to each turn, even if other crews are not busy at the time.

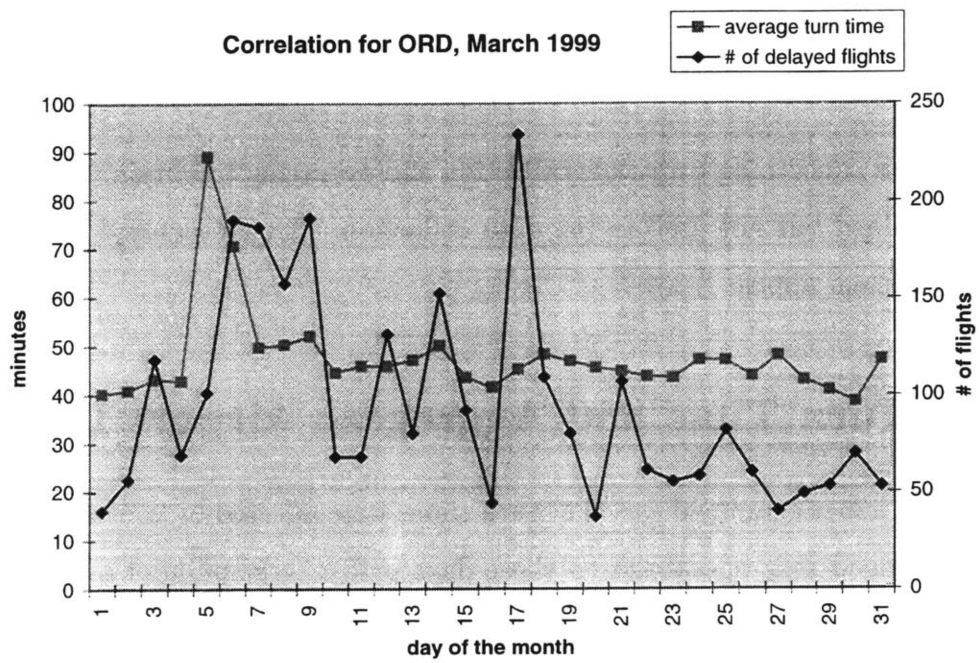

Figure B-1: Correlation between number of delayed flights per day and average turn time. Data taken from March 1999 ASQP for Chicago O'Hare. There is no noticeable positive or negative correlation between number of delayed flights and the average turn time. The correlation coefficient for the data points plotted above was .36.

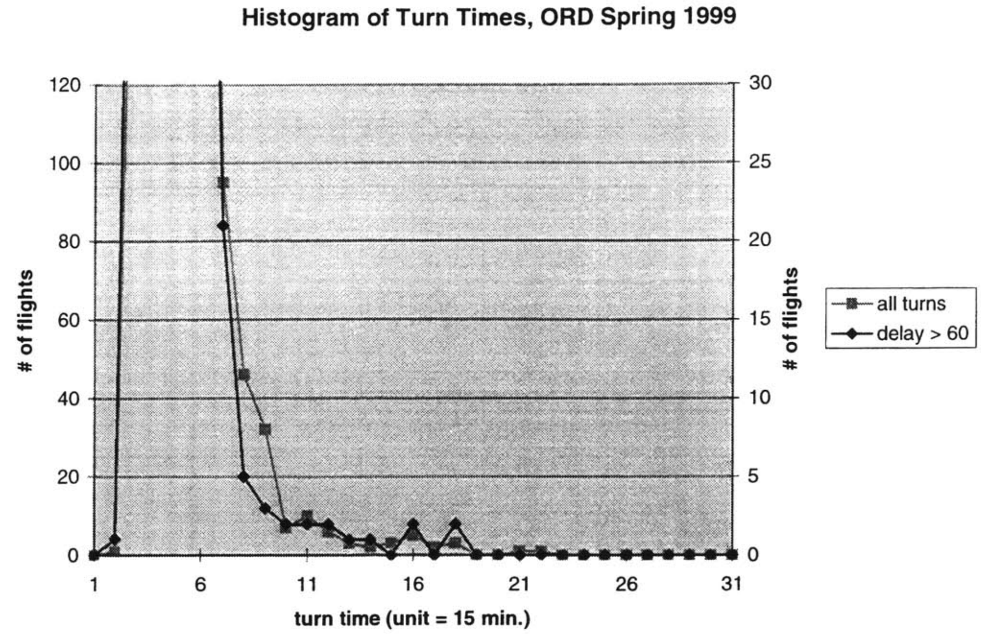

Figure B-2: Very late flights still have the same turn times as partially delayed flights. In the graph above, turn times of all flights are compared with those flights whose incoming delay is greater than one hour. In both cases, a vast majority of turns are completed in 3-8 time units (45-120 minutes). The turns completed in greater than 10 time unites (150 minutes) are statistically insignificant.

B.2 Turn Time and Large Individual Delay

In Chapter 2, it was illustrated that only delayed flights could be examined if a minimum turn time was to be found. On the flip side is the question, what if the delay is very great, i.e., larger than an hour, larger than two hours? In Figure B-2 we see that as an incoming leg is more and more delayed, as long as it is not swapped, the turn is still made as quickly as possible. This conclusion is reasonable from an operations standpoint: if the flight is very late but not cancelled, the airlines will still struggle to get the plane out as quickly as possible. In light of this argument, turns that had very late incoming legs were included in the average.

Appendix C

ADL Aircraft Type explanation

A list of the equipment types present in ADL updates is given in Table C.1:

Equipment Code	Meaning
AB3	A300-600
A320	A320-1/2
A319	A319
MD80	DC/MD-80
F100, FK10	FOKKER-100
B722	B-727-200
B732	B-737-1/2
B733	B-737-300
B734	B-737-400
B735	B-737-500
B73J, B73S	B-737-700
B752	B-757-200
B762	B-676-200
B772	B-777-200
MD80	DC/MD-80
DC9	DC-9-30, -40, -50
DC10	DC-10

Table C.1: Equipment types present in the CDM ADL updates. Not all equipment types are shown. Equipment definitions taken from Official Airline Guide (OAG) [3].

Appendix D

Values for Turn1 for LaGuardia, Philadelphia, and Pittsburgh airports.

This appendix holds average minimum turn times for the other three airports, whose turn times were not included in Chapter 3.

	Airline						
Aircraft Type	**AA**	**UA**	**US**	**DL**	**NW**	**CO**	**TW**
B-727	46	50	-	41	45	-	49
B-737-1/2	-	-	42	-	-	-	-
B-737-3	-	47	42	-	-	37	-
B-737-5	-	34	51	-	-	36	-
B-757	59	51	-	59	55	-	-
DC-9	-	-	35	-	40	-	-
MD-80	46	-	52	54	51	-	52
A320-1/2	-	42	-	-	-	-	-

Table D.1: Most common equipment types and corresponding average minimum turn times for the 7 major carriers operating out of LaGuardia airport. All turns with an incoming leg delay of 15 or more minutes were averaged from March to May of 1999. The averages for the 737 aircraft are broken up just as in the table in Chapter 3.

	Airline						
Aircraft Type	**AA**	**UA**	**US**	**DL**	**NW**	**CO**	**TW**
B-727	48	39	-	57	40	-	50
B-737-1/2	-	46	47	-	-	-	-
B-737-3	-	43	45	49	-	44	-
B-737-5	-	37	47	-	-	32	-
B-757	41	-	66	49	49	-	-
DC-9	-	-	42	-	40	28	37
MD-80	35	-	54	50	39	34	57
A320-1/2	-	41	-	-	44	-	-

Table D.2: Most common equipment types and corresponding average minimum turn times for the 7 major carriers operating out of Philadelphia airport. All turns with an incoming leg delay of 15 or more minutes were averaged from March to May of 1999. The averages for the 737 aircraft are broken up just as in the previous table.

	Airline						
Aircraft Type	**AA**	**UA**	**US**	**DL**	**NW**	**CO**	**TW**
B-727	-	47	-	-	-	-	-
B-737-1/2	-	46	42	-	-	-	-
B-737-3	-	44	48	-	-	-	-
B-737-5	-	46	47	-	-	34	-
B-757	-	-	59	-	-	-	-
DC-9	-	-	40	-	30	28	48
MD-80	45	-	49	39	-	-	-
A320-1/2	-	37	-	-	-	-	-

Table D.3: Most common equipment types and corresponding average minimum turn times for the 7 major carriers operating out of Pittsburgh airport. All turns with an incoming leg delay of 15 or more minutes were averaged from March to May of 1999. The averages for the 737 aircraft are broken up just as in the previous table. Airlines other than United and USAir have very few flights servicing Pittsburg, leading to the lack of turn time averages for certain types of equipment seen above.

Appendix E

Additional Results

This appendix lists exhaustive results for all airports.

E.1 Chicago O'Hare Intenational Airport (ORD)

Figures E-1, E-2, E-3, E-4, and E-5.

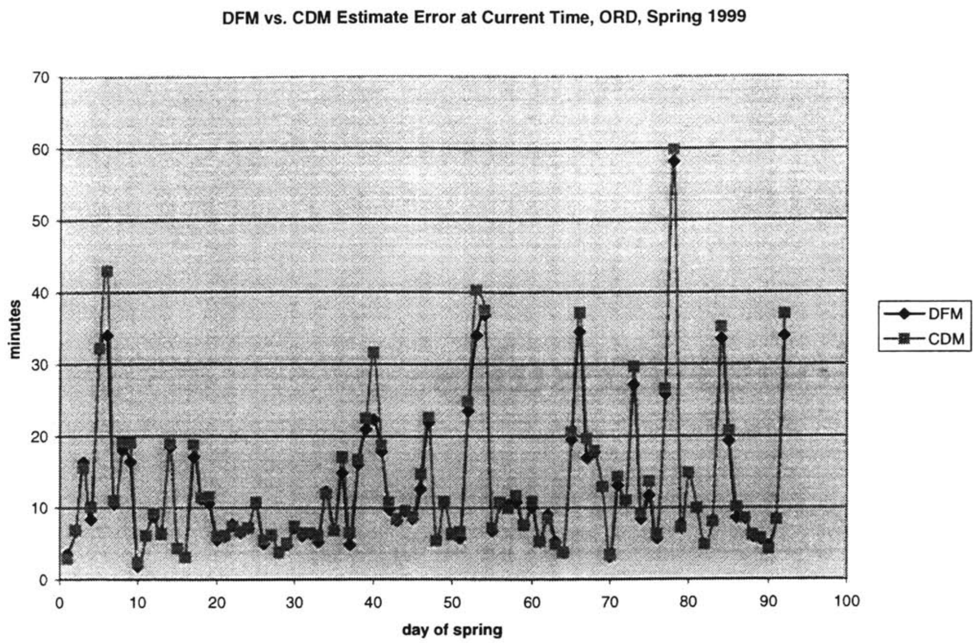

Figure E-1: ORD Error for Current Time, Spring 1999

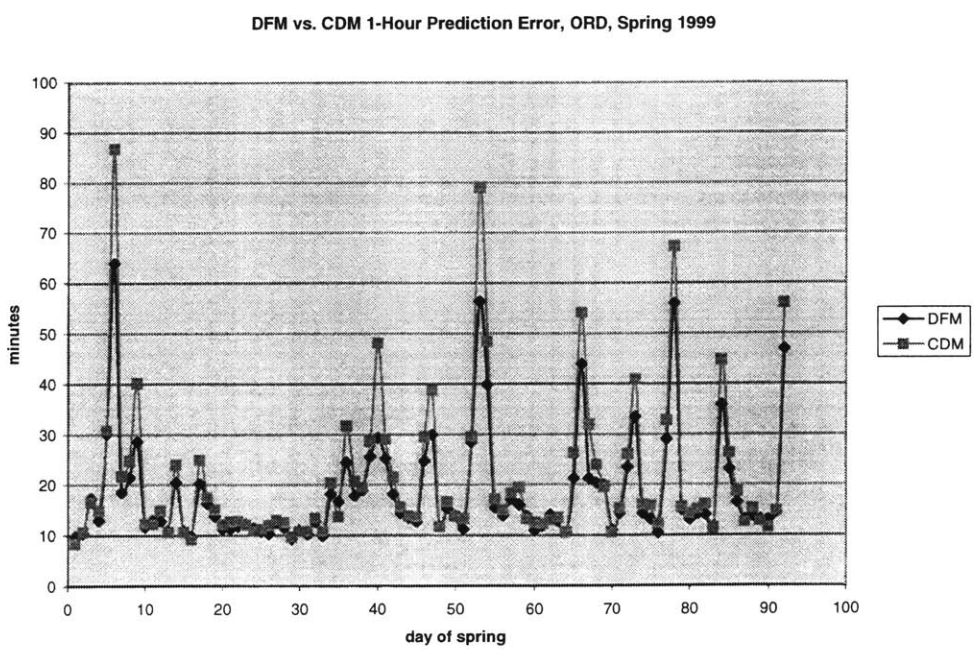

Figure E-2: ORD Error for 1-Hour Prediction, Spring 1999

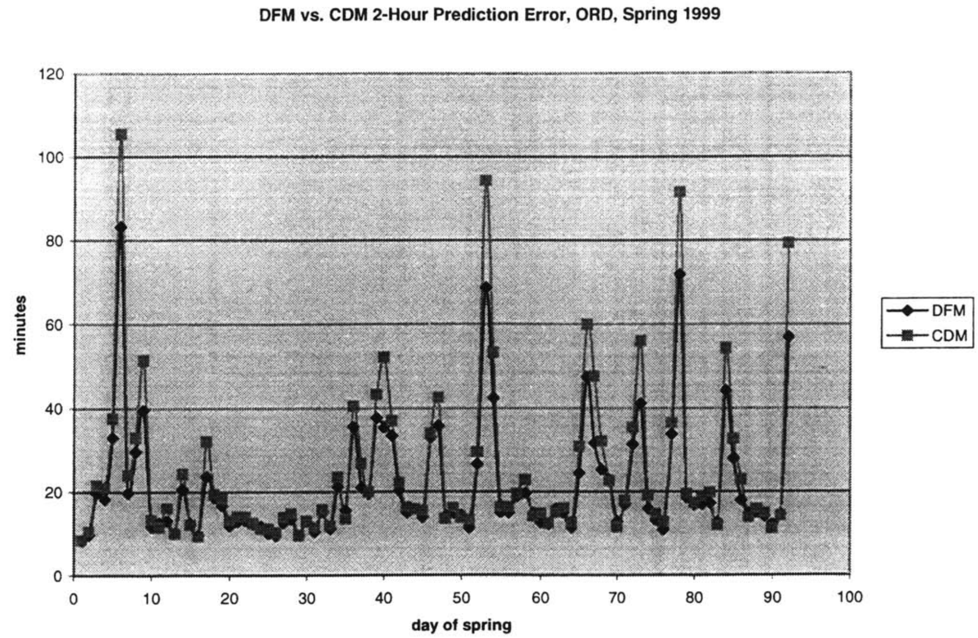

Figure E-3: ORD Error for 2-Hour Prediction, Spring 1999

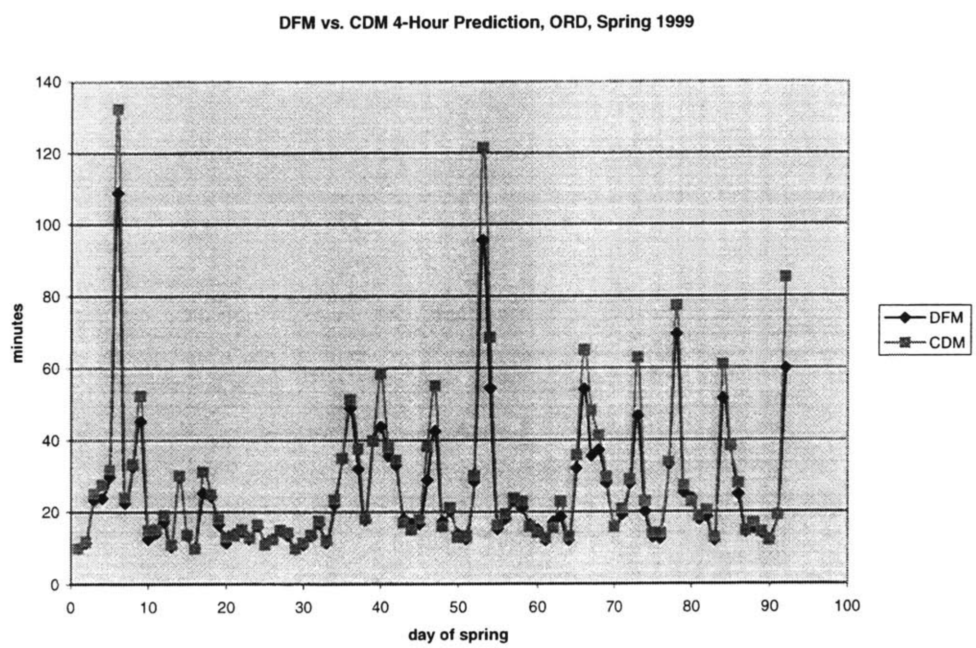

Figure E-4: ORD Error for 4-Hour Prediction, Spring 1999

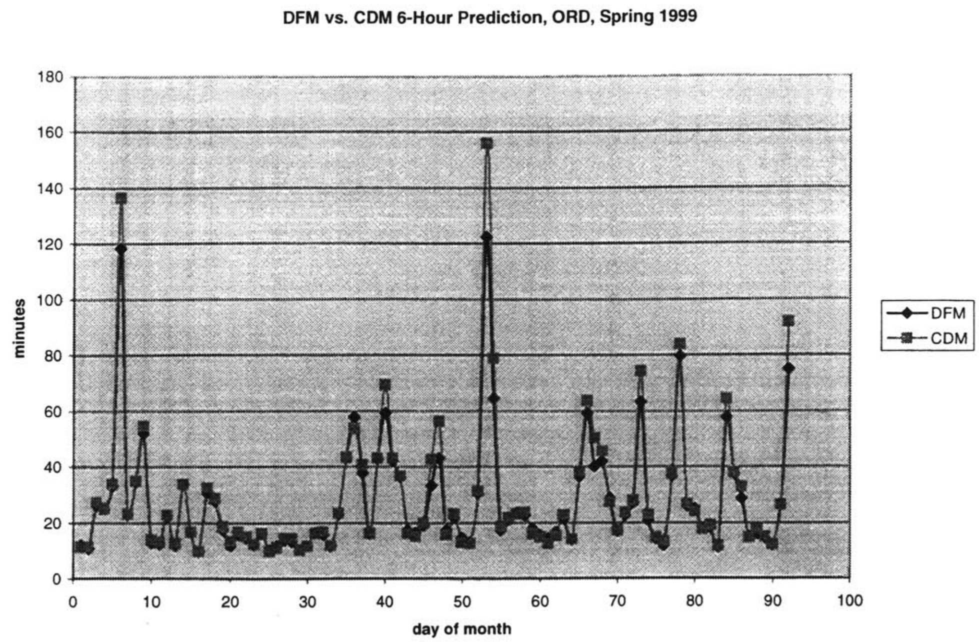

Figure E-5: ORD Error for 6-Hour Prediction, Spring 1999

E.2 Boston Logan Intenational Airport (BOS)

Figures E-6, E-7, E-8, E-9, and E-10.

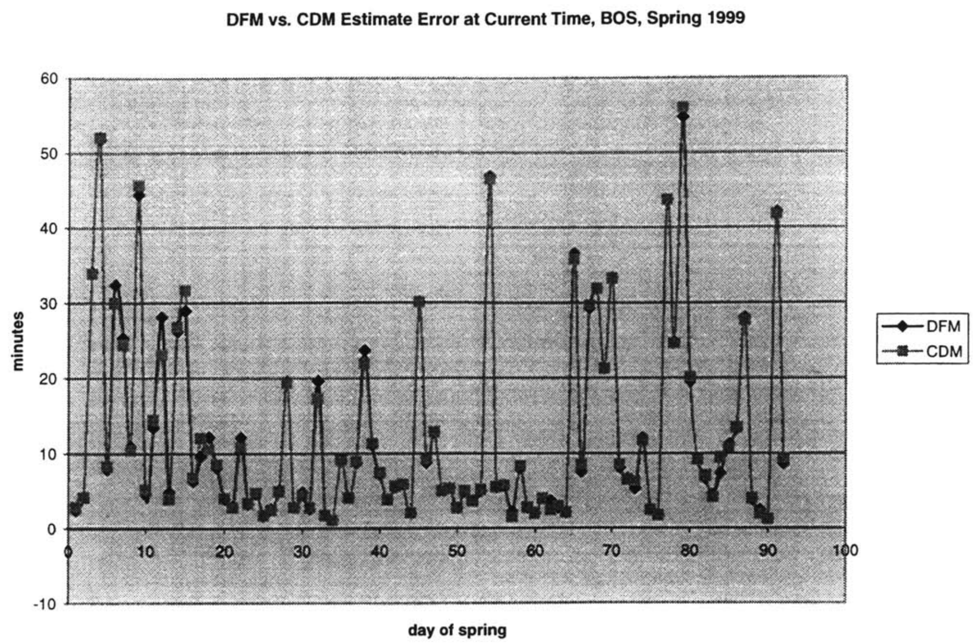

Figure E-6: BOS Error for Current Time, Spring 1999

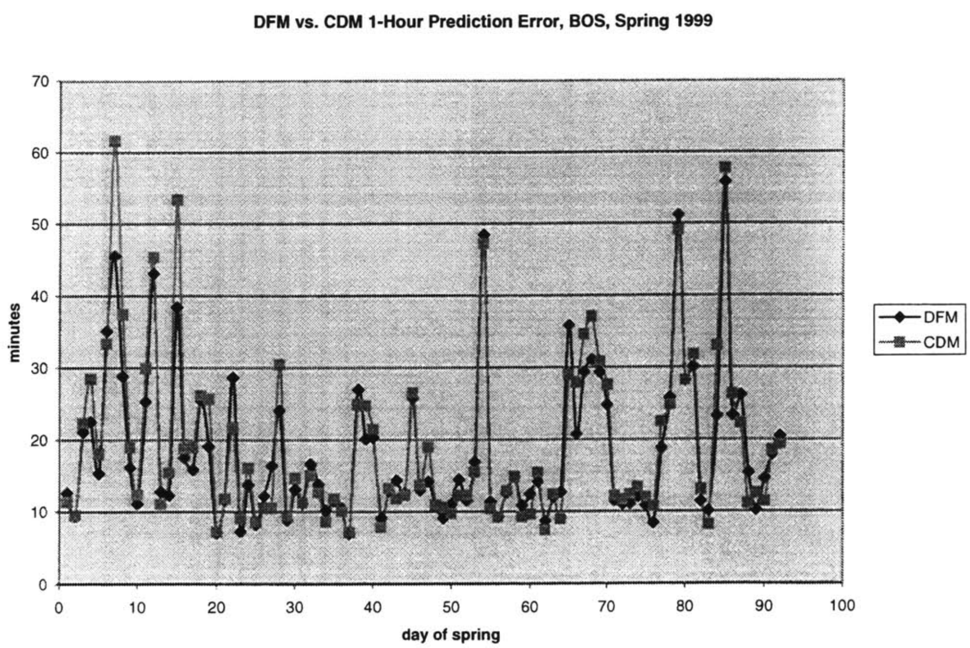

Figure E-7: BOS Error for 1-Hour Prediction, Spring 1999

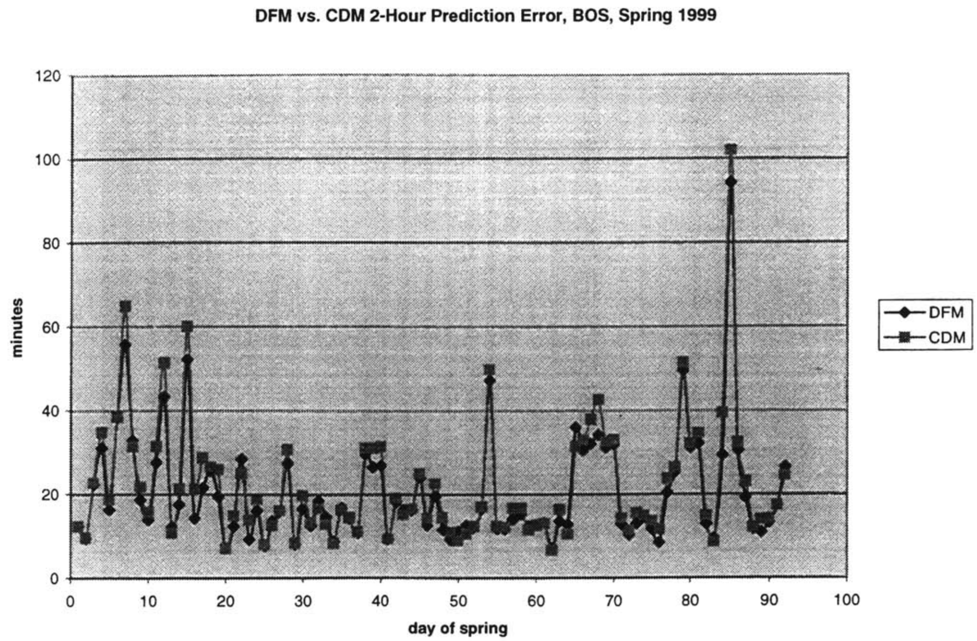

Figure E-8: BOS Error for 2-Hour Prediction, Spring 1999

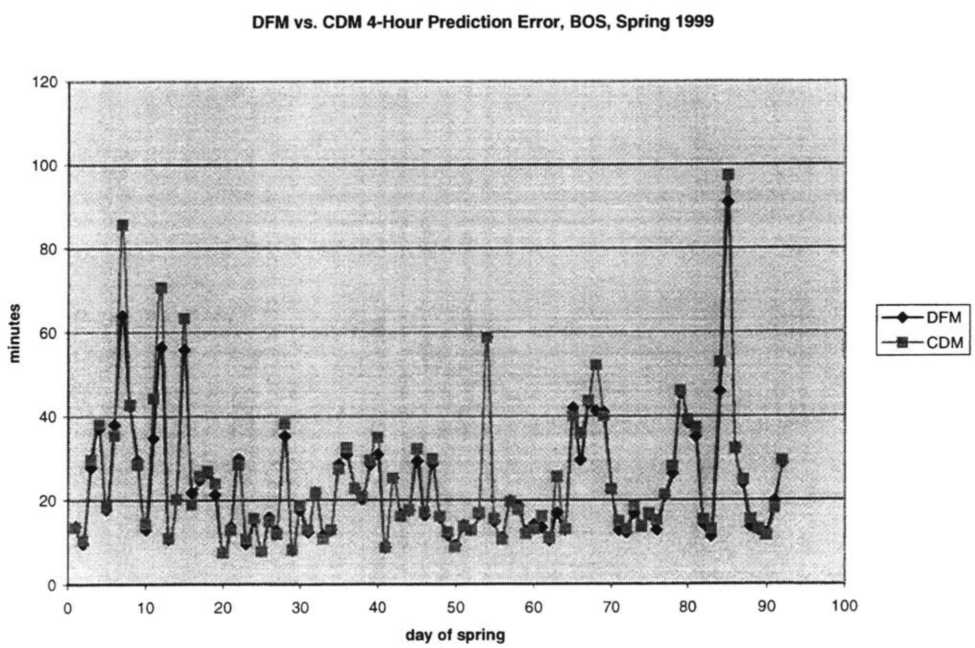

Figure E-9: BOS Error for 4-Hour Prediction, Spring 1999

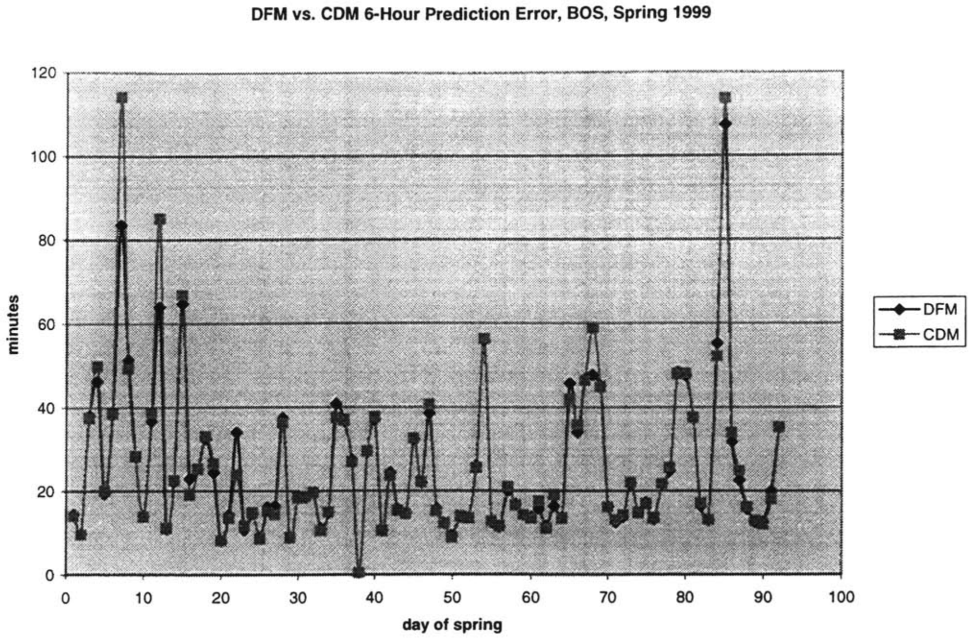

Figure E-10: BOS Error for 6-Hour Prediction, Spring 1999

E.3 New York La Guardia Airport (LGA)

Figures E-11, E-12, E-13, E-14, and E-15.

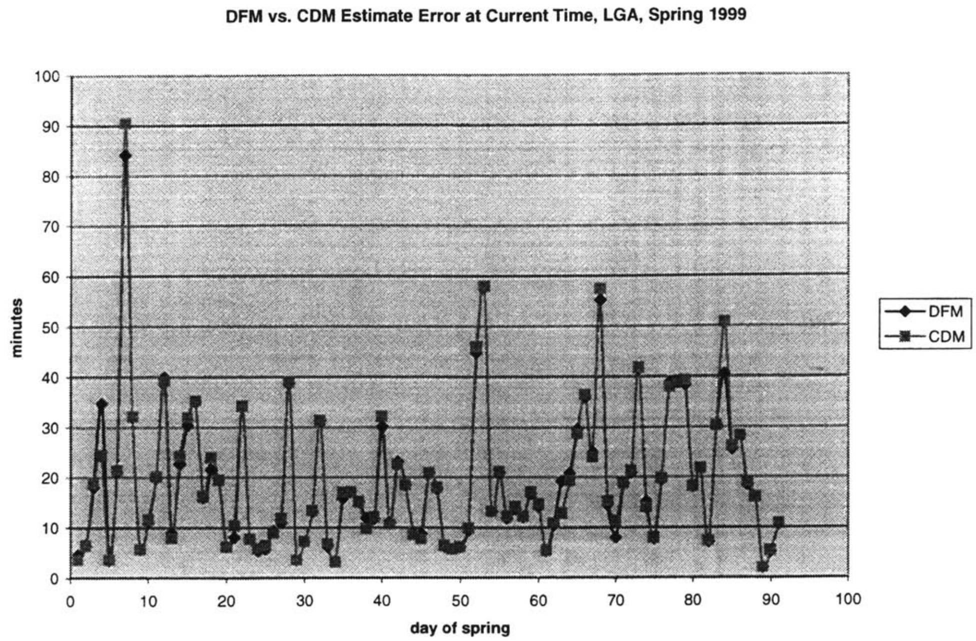

Figure E-11: LGA Error for Current Time, Spring 1999

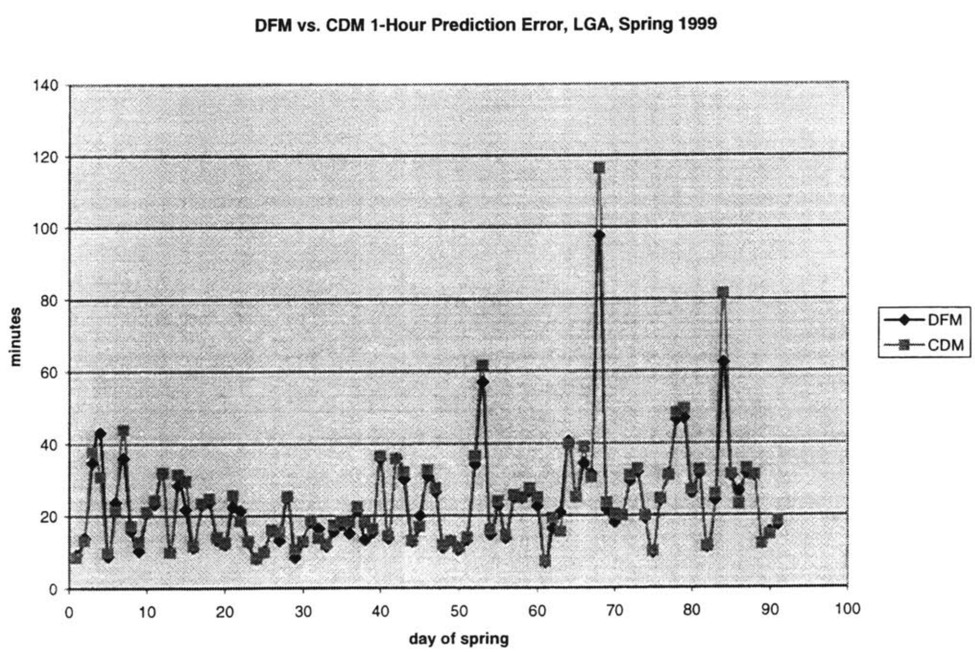

Figure E-12: LGA Error for 1-Hour Prediction, Spring 1999

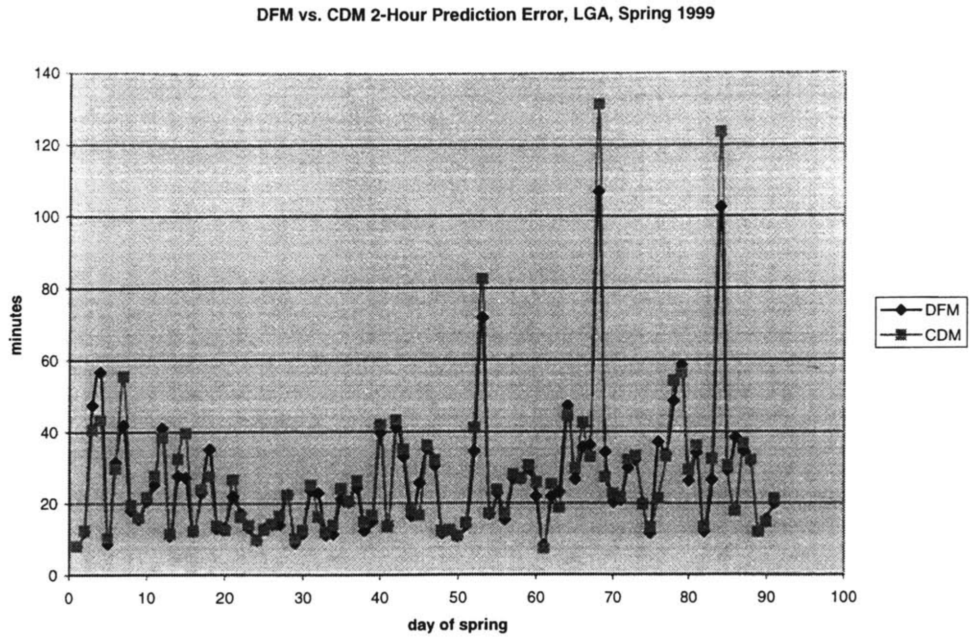

Figure E-13: LGA Error for 2-Hour Prediction, Spring 1999

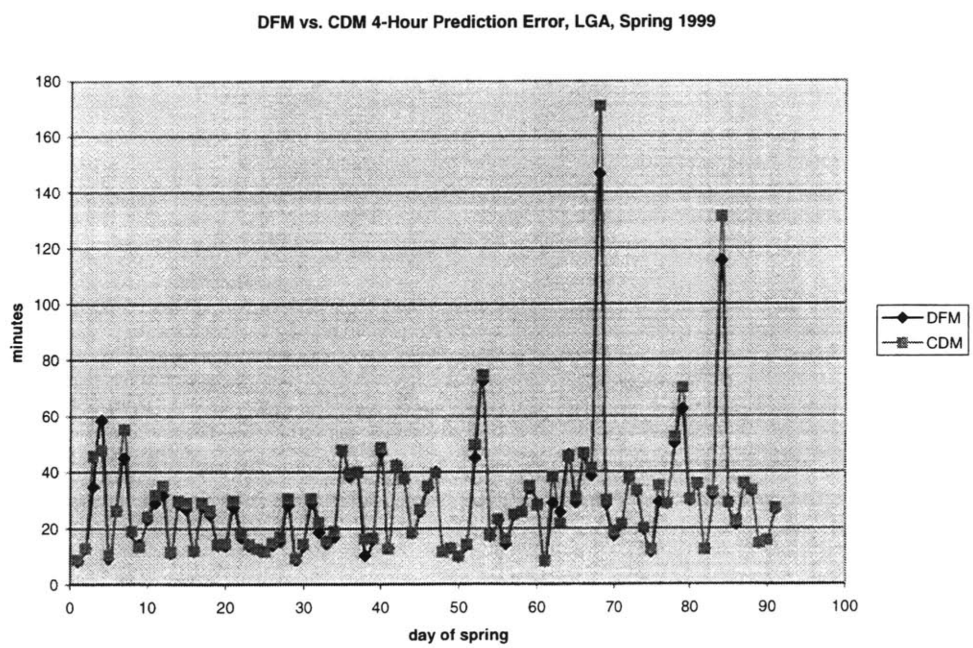

Figure E-14: LGA Error for 4-Hour Prediction, Spring 1999

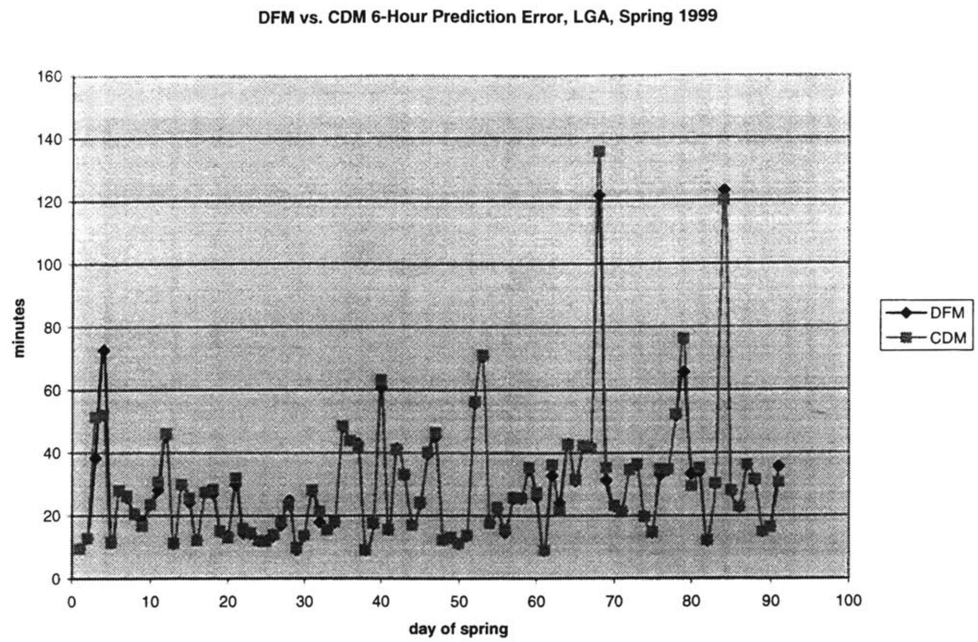

Figure E-15: LGA Error for 6-Hour Prediction, Spring 1999

E.4 Piladelphia Intenational Airport (PHL)

Figures E-16, E-17, E-18, E-19, and E-20.

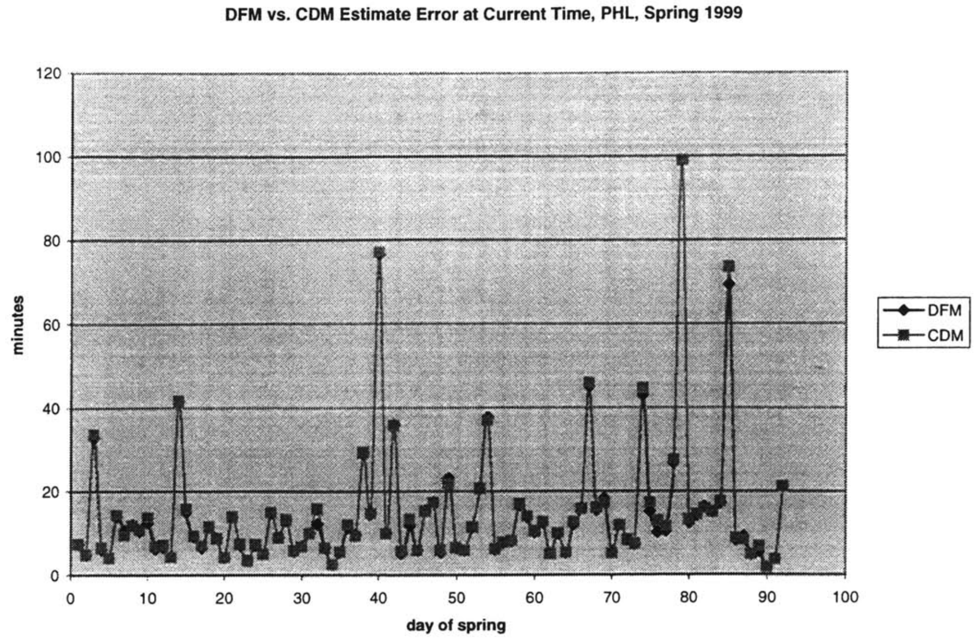

Figure E-16: PHL Error for Current Time, Spring 1999

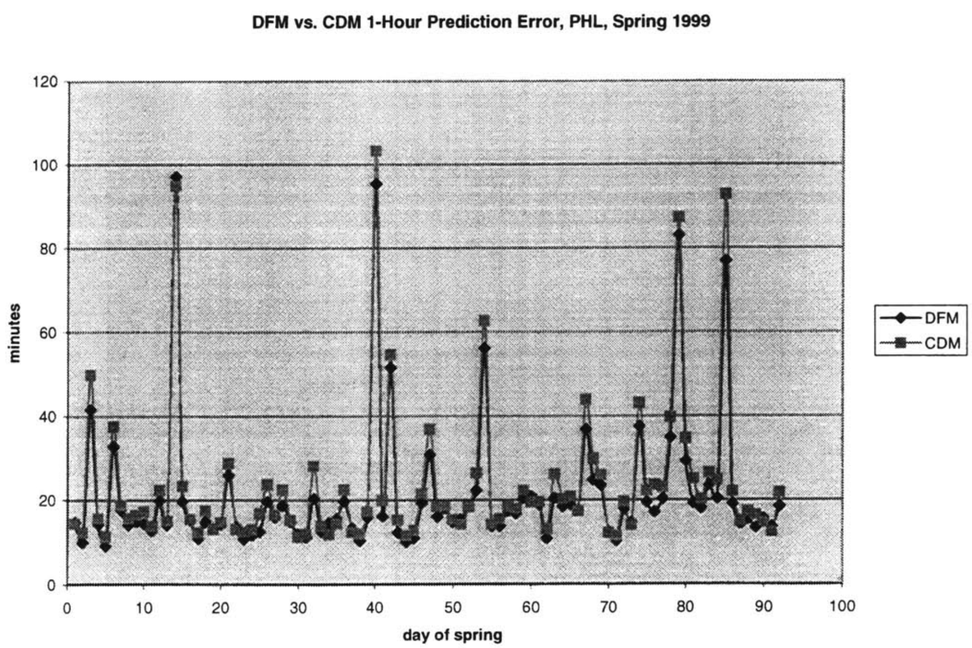

Figure E-17: PHL Error for 1-Hour Prediction, Spring 1999

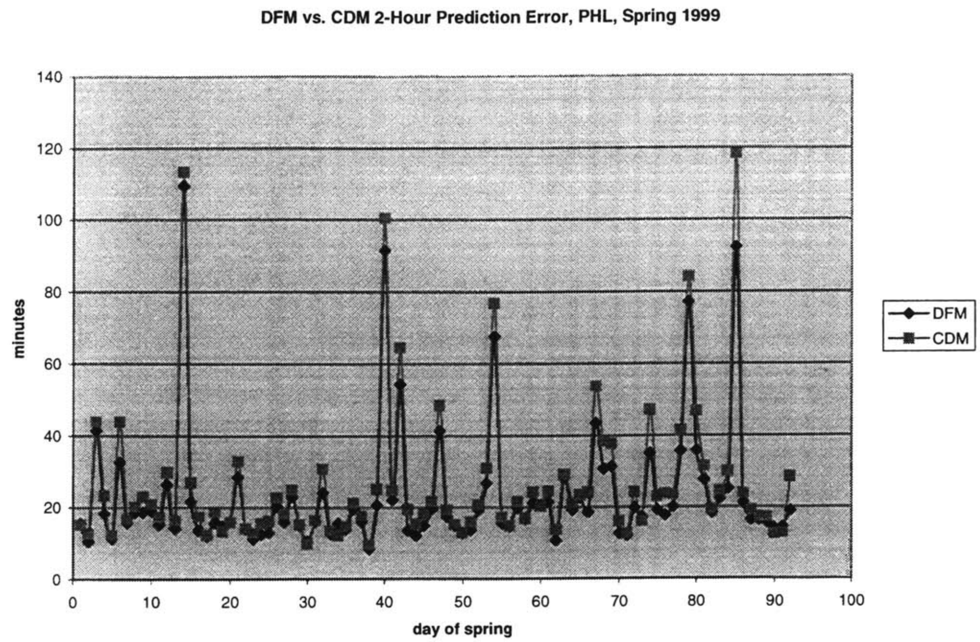

Figure E-18: PHL Error for 2-Hour Prediction, Spring 1999

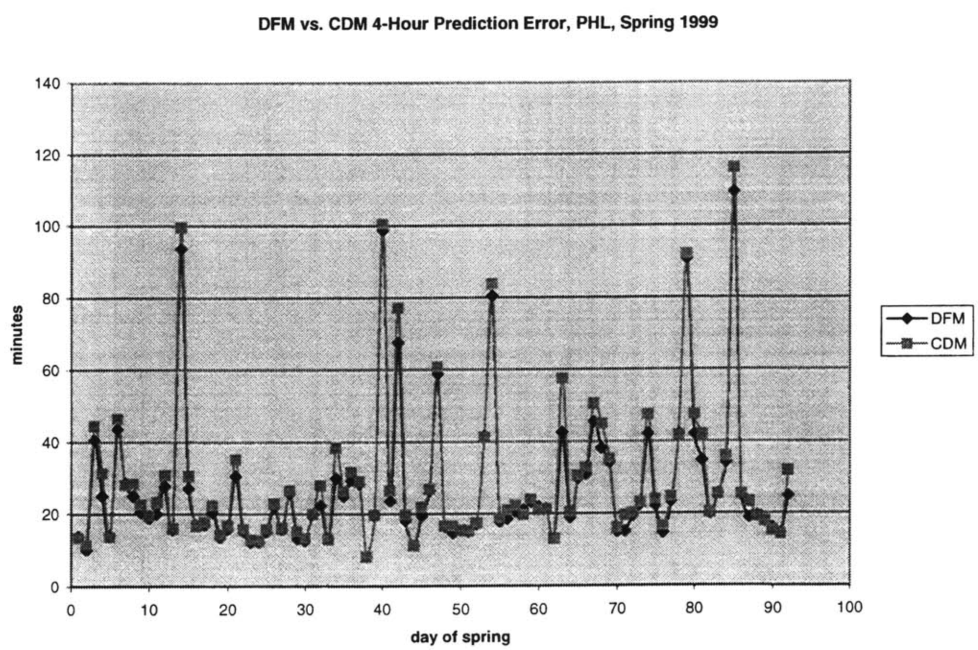

Figure E-19: PHL Error for 4-Hour Prediction, Spring 1999

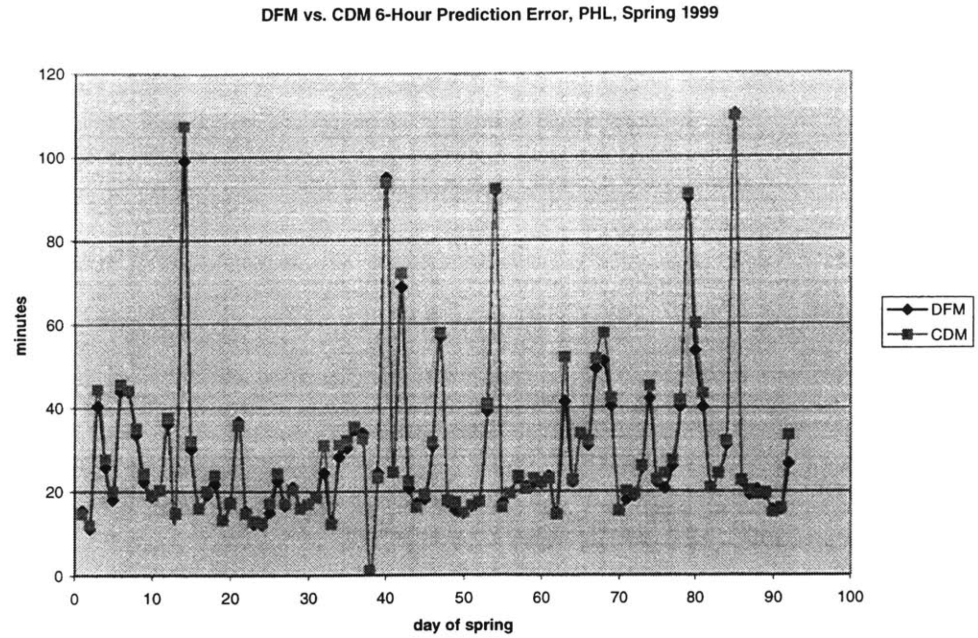

Figure E-20: PHL Error for 6-Hour Prediction, Spring 1999

E.5 Pittsburgh Intenational Airport (PIT)

Figures E-21, E-22, E-23, E-24, and E-25.

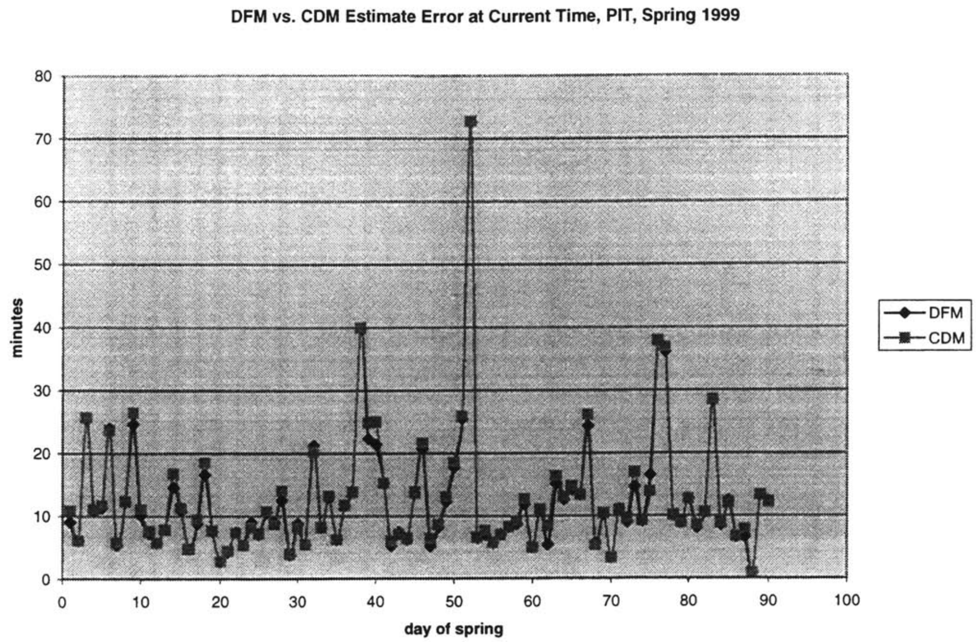

Figure E-21: PIT Error for Current Time, Spring 1999

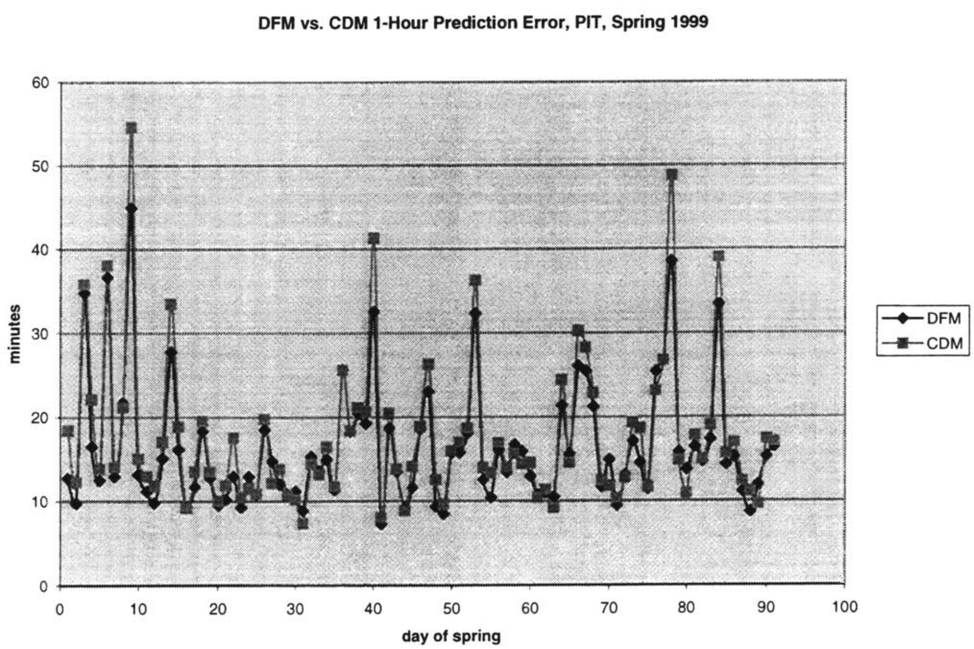

Figure E-22: PIT Error for 1-Hour Prediction, Spring 1999

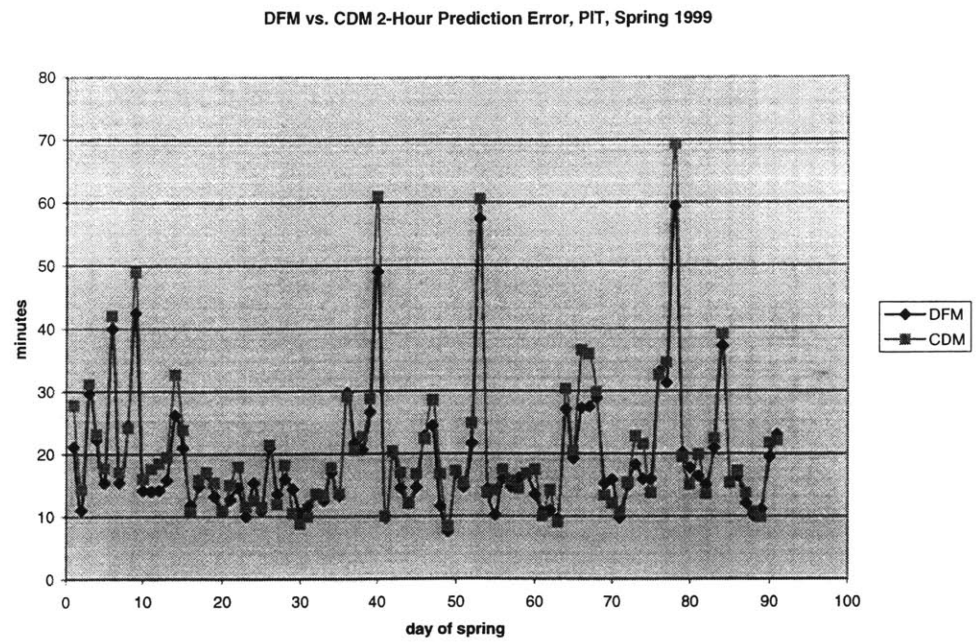

Figure E-23: PIT Error for 2-Hour Prediction, Spring 1999

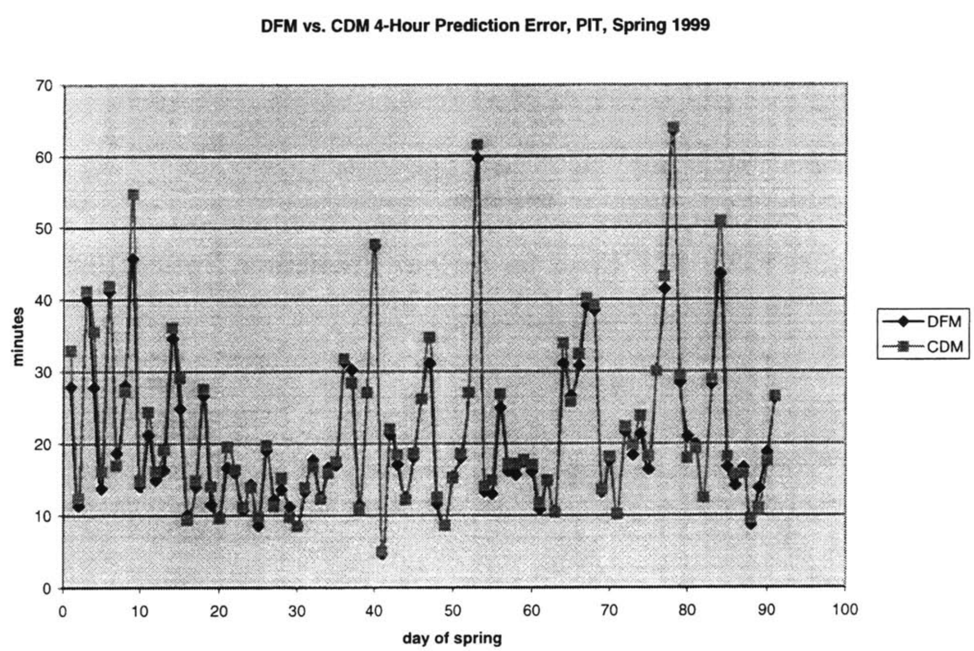

Figure E-24: PIT Error for 4-Hour Prediction, Spring 1999

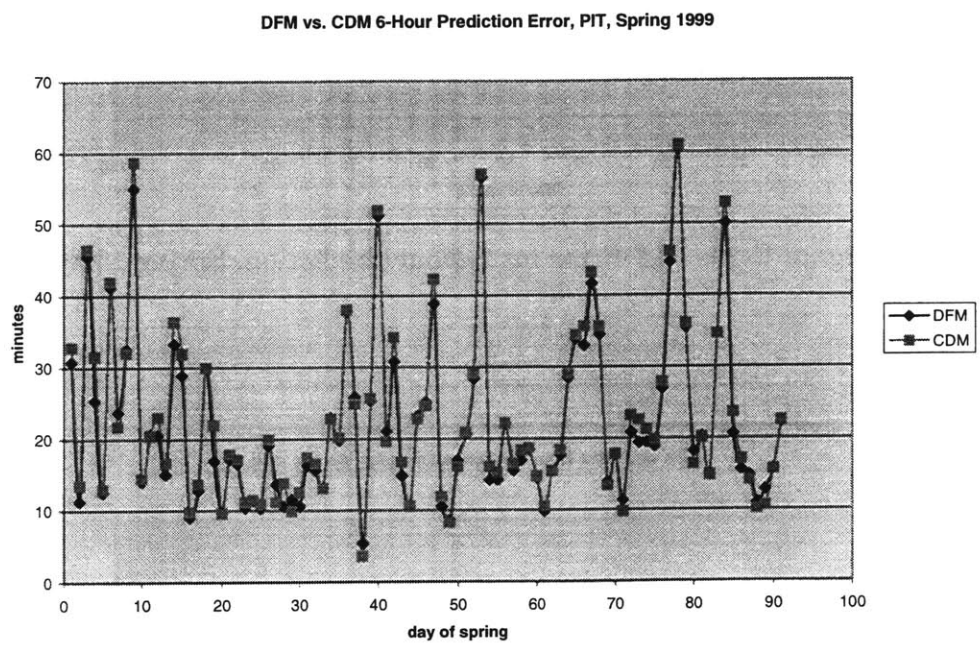

Figure E-25: PIT Error for 6-Hour Prediction, Spring 1999

Bibliography

[1] Kari Andersson. Potential benefits of information sharing during the arrival process at hub airports. Master's thesis, Massachusetts Institute of Technology, 2000.

[2] Air Traffic Control System Command Center. Collaborative decision making (cd-m) website. http://www.atcscc.faa.gov/Information/CDM/cdm.html, 1999.

[3] Rayson Comiskey, Nelson. Original airline desktop guide. Richard A. Nelson, Publisher, 1999.

[4] Jody Hoffer Gittell. *Crossfunctional Coordination and Control: Influencing Employee Behavior and Process Outcomes through Organization Design in the Airline Industry.* PhD thesis, Massachusetts Institute of Technology, 1995.

[5] James Ott. Air travel in crisis. *Aviation Week*, 151(17):42–44, October 1999.

[6] Robert Arthur Shumsky. *Dynamic Statistical Models for the Prediction of Aircraft Take-off Times.* PhD thesis, Massachusetts Institute of Technology, 1995.

[7] CDM Workgroup. Cdm adl data format update, march. Memo, 1999.

[8] CDM Workgroup. Ground delay program summary for july 1999. Memo, 1999.