
Collaborative Decision Making in Air Traffic Management: A Preliminary Assessment

*Michael Ball, Robert Hoffman
William Hall, Alp Muharremoglu*

Prepared by

NEXTOR

The National Center of Excellence for Aviation Operations Research

Report #RR-99-3

August 1998

Project Coordinator

Robert Hoffman

Project Team

University of Maryland Massachusetts Institute of Technology

Michael Ball
Robert Hoffman
Tasha Inniss
Thomas Vossen
Chien-Yu Chen
Daniel Darr
Joseph Previte

Amedeo Odoni
William Hall
Alp Muharremoglu
Ioannis Anagnostakis
Ryan Rifkin
John-Paul Clarke
John Jensen

Preface

This report gives the findings of the Spring, 1998 NEXTOR analysis of the first five months of operation of the joint FAA-industry initiative, Collaborative Decision Making (CDM). The analysis sought to identify

- realized and potential benefits derived from CDM,
- operational changes within the air traffic community as a result of CDM,
- those components of CDM that have been particularly successful,
- and those components of CDM that require special attention.

The report is structured as follows: Chapter 1 gives an introduction and background to CDM. Chapter 2 delivers the results of the analyses. Chapter 3 summarizes the results. The Appendices contain the details of the analyses and other relevant information such as a description of the analytic models and metrics that were employed. The Glossary lists terms and acronyms used throughout the report.

An executive summary of this report is also available.

Acknowledgements

This report was prepared by the National Center of Excellence for Aviation Operations Research, under Federal Aviation Administration Research Grant Number 96-C-001 and contract number DTFA03-97-D-00004. This document has not been reviewed by the Federal Aviation Administration (FAA). Any opinions expressed herein do not necessarily reflect those of the FAA or the U.S. Department of Transportation.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 1.1 | Background of Study | 1 |
| 1.2 | Introduction | 2 |
| 1.2.1 | Status of CDM | 4 |
| 1.2.2 | Databases | 5 |
| 1.2.3 | The AOCnet | 6 |
| 2 | Results of Analysis | 8 |
| 2.1 | Initial Implementation | 8 |
| 2.1.1 | Quality of Information | 8 |
| 2.1.2 | Information Distribution and Situational Awareness | 18 |
| 2.1.3 | Airline Resource Allocation Decisions | 20 |
| 2.1.4 | Delay Reduction | 21 |
| 2.2 | Assessment of Future Benefits | 29 |
| 2.3 | Effectiveness of CDM Components | 30 |
| 2.4 | Special Issues | 30 |
| 3 | Summary and Conclusions | 32 |
| A | Data and Source | 36 |

| | |
|---|-----------|
| B Airborne Holding Analysis | 38 |
| B.1 Background | 38 |
| B.2 The Model | 38 |
| B.3 Accuracy of the Model | 40 |
| C Flight Counts | 42 |
| D Cancellation Notices | 43 |
| D.1 Background | 43 |
| D.2 Details of the Analysis | 44 |
| E IPE Metric | 48 |
| E.1 What is IPE? | 48 |
| E.2 Variations on IPE | 50 |
| E.3 Units of IPE | 51 |
| F Command Center Interviews | 53 |
| F.1 Structure of the Interviews | 53 |
| F.2 Screening for Objectivity | 53 |
| F.3 FSM at the ATCSCC | 54 |
| F.4 Results of the Interviews | 55 |
| F.5 Summary of Interviews | 59 |
| G Glossary | 62 |
| References | 65 |
| Index | 66 |

List of Figures

| | | |
|-----|---|----|
| 1.1 | CDM Databases and Processing Strings | 6 |
| 1.2 | The CDM Flow of Information | 7 |
| 2.1 | Key Elements in Predicting Arrival Demand | 9 |
| 2.2 | More Accurate Predictions of Departure Time | 11 |
| 2.3 | The Number of Cancellation Notices Per Day at SFO | 13 |
| 2.4 | Cancellations per Hour at EWR and SFO on a Typical Day | 14 |
| 2.5 | Distribution of Time Before OETD that First CNX Was Known | 15 |
| 2.6 | Comparison of Error in Prediction of Arrival Demand | 17 |
| 2.7 | Delay Savings at EWR Due to Compression | 25 |
| 2.8 | Delay Multiplier as Function of Base Delay | 26 |
| B.1 | Model of Airborne Holding | 39 |
| E.1 | Computing the Integral for IPE | 49 |
| E.2 | A Dampening Function for IPE | 50 |

List of Tables

| | | |
|-----|--|----|
| 2.1 | Comparison of C and CDM IPE-6 Values | 10 |
| 2.2 | Airborne Delay Averages for Three Periods | 27 |
| A.1 | Delta Files | 36 |
| D.1 | Cancellation Notice Results for a Typical Day at SFO | 46 |

1. Introduction

1.1 Background of Study

Collaborative Decision Making (CDM) embodies a new philosophy for managing air traffic. It was initially conceived in the mid-1990s within the FADE (FAA-Airlines Data Exchange) project. That project, as well as the initial operational implementation of CDM, has been aimed at the development of new operational procedures and decision support tools for implementing and managing ground delay programs (GDPs). However, a much broader class of problems in air traffic management is envisioned for the application of CDM.

CDM prototype operations commenced in January of 1998. At the request of the FAA, NEXTOR (National Center of Excellence for Aviation Operations Research) undertook the task of assessing the current status of CDM. This assessment should be viewed as preliminary, given that, at the start of this analysis, CDM had been operational for less than five months and some of our analysis employs data only from the first three months. The assessment has sought to answer the following questions:

- Has the initial implementation of CDM led to substantial, beneficial operational changes?
- Have the CDM-based GDPs produced tangible benefits?
- What are the prospects for future CDM benefits?
- Have certain components of the initial implementation been more effective than other components?
- Are there problem areas that deserve special attention or new dedication of resources?

It is important to distinguish between CDM, the philosophy, and the initial implementation of CDM in the context of ground delay programs. CDM, the philosophy, can be potentially applied in a wide range of contexts and its success

in any particular area depends on a range of specific conditions. Our analysis is based on results of the initial implementation and, thus, primarily applies to the context of ground delay programs. Additionally, we try to point out where the initial implementation demonstrates the underlying CDM philosophy.

In order to carry out the analyses, operational data sets covering conditions both before and after the implementation of CDM were gathered. These data sets were analyzed in a variety of ways to estimate the effects of the initial CDM implementation on the quality of data and operations. In addition, input was solicited from airline operations personnel. The input provided by airlines typically consisted of either the results of internal airline benefits analysis or a qualitative description of changes in operational procedures that have resulted from CDM. The final source of inputs for the report consisted of a set of interviews conducted with specialists from the Air Traffic Control System Command Center (ATCSCC).

1.2 Introduction to CDM

CDM is a joint FAA-Industry initiative aimed at improving air traffic flow management. CDM strives to achieve such improvements by

- generating better information, usually by combining information generated by the FAA with information generated by National Airspace System (NAS) users;
- distributing the same information both to FAA traffic flow managers and to NAS users;
- creating tools and procedures that allow NAS users to directly respond to capacity/demand imbalances and to collaborate with FAA traffic flow managers in the formulation of flow management actions.

The initial concepts for CDM were developed in the early 1990s with the FAA-Airline Data Exchange (FADE) project. These initial efforts, as well as the initial implementation of CDM in January of 1998, have been aimed at creating better procedures for generating Ground Delay Programs (GDPs).

A GDP is a control action taken by the Air Traffic Control System Command Center (ATCSCC) in response to degraded arrival capacity at one or more airports. In the typical scenario, bad weather forces the reconfiguration of runways

at an airport and the use of instrument flight rules (IFR). Under such conditions, the number of flights that can land at an airport per hour can be reduced by a factor of 2 or more. During busy periods, such reductions do not allow for the timely accommodation of the airlines published schedules. In order to balance the flow into an airport with its arrival capacity, the ATCSCC can issue a GDP, which consists of a set of revised departure times for flights destined for the impacted airport. For example, a flight that was scheduled to take off at 2 PM might be issued a ground delay of 25 minutes, resulting in a revised departure time of 2:25 PM. When the flight takes off at 2:25, it would be able to fly directly to the impacted airport without experiencing any airborne delay. On the other hand, if a ground delay were not applied to the flight, then, when the flight reached the airspace adjacent to the destination airport, it would be placed in an airborne holding pattern, resulting in a larger consumption of fuel and greater demand on the air traffic control system.

CDM grew out of a desire on the part of both the airlines and the FAA for improvements in the manner in which GDPs were generated. The FAA, or more specifically, the ATCSCC, had a desire to receive more up-to-date information on the status of flights and the intentions of the airlines. Such information might include the status of flights that were delayed or canceled due to mechanical difficulties, or due to delays encountered at prior stops on the aircraft itinerary. Previously, the ATCSCC did not have access to this information or did not receive it in a timely manner. In fact, the slot allocation process actually discouraged the airlines from providing the most up-to-date information. With this improved information, the specialists can gain a more accurate estimate of the future demand on the impacted airport and are able to run more effective GDPs. The airlines did not feel the allocation procedures used by the ATCSCC were always fair. In addition, they had a desire to gain more control over how delays were allocated to their flights. For example, because of their position in the original schedule, the FAA's GDP procedures might allocate AAL 255 a delay of 15 minutes and AAL 405 a delay of 30 minutes. It could easily happen that, because of connecting flights for either the crew or equipment used for AAL 405, the cost per minute of delay to AAL 405 was much higher than the cost per minute of delay to AAL 255. Thus, American Airlines would like the capability to switch the slot assignments so that ALL 405 received a 15-minute delay and AAL 255 a 30-minute delay. Even before the advent of CDM, such slot assignment switching could be carried out. The intention of CDM is to provide a higher level of flexibility and comprehensiveness when compared to the previous capabilities.

1.2.1 Status of CDM

The current status of CDM can be summarized as follows:

- Agreement on new GDP paradigm: The airlines and the FAA have agreed on a new resource allocation concept embodied within two procedures: ration-by-schedule and compression. These provide mechanisms for a “fair” rationing of arrival time slots under degraded conditions. A key driver to all of CDM is that all parties have agreed on the fairness of the allocation. Moreover, the ability of a specific airline to achieve its fair allocation is not hindered by that airline providing the most up-to-date information on the status of its flights.
- Regular meetings of all CDM players: All CDM players, including airline representatives, FAA traffic flow managers, developers and planners, software developers, consultants, and researchers, meet regularly (usually monthly). These meetings have allowed for the CDM agenda to be pushed forward very rapidly and have led to greater mutual understanding among the parties involved.
- Flight Schedule Monitor (FSM), the CDM decision support tool: FSM has undergone an evolutionary development over the past few years and it now embodies the requirements expressed by both FAA and the airlines. The tool is in place at both the ATCSCC and the operations centers of all CDM participating airlines.
- AOCnet: The participating airline operations centers (AOCs), the ATCSCC, the Volpe National Transportation Systems Center (the hub site of the Enhanced Traffic Management System – ETMS), Metron (the developer of FSM) as well as certain other parties are all interconnected via a private intra-net, the airline operations center AOCnet. This network is used to exchange CDM operational information.
- Prototype implementation: The initial implementation of CDM commenced in January of 1998. All participating airlines, the Volpe Center and the ATCSCC are exchanging information via the AOCnet. Furthermore, as of January 20, 1998, CDM procedures were used to implement all GDPs at San Francisco (SFO) and Newark (EWR) airports. On April 20, 1998, this list was expanded to include La Guardia (LGA) and St. Louis (STL) airports.

- Work on future applications of CDM: Although CDM began as a way of improving ground delay programs, it embodies basic concepts that have wider applicability. Two CDM working groups, NAS Status and Collaborative Routing, have been working to extend CDM into the following areas:
 - NAS Status: The goal of this group is to use the AOCnet to distribute a variety of informations on the status of the National Airspace System (NAS) to various NAS users. Its efforts have been geared toward identifying information beneficial to NAS users and to defining sources for the information and procedures for delivering that information to the AOCnet.
 - Collaborative Routing: The goal of this group is to apply CDM procedures and concepts to problems involving the routing of aircraft, particularly issues related to routing around congested and weather-impacted airspace.

1.2.2 Databases and Processing Strings

The implementation of CDM required a significant enhancement to the data management and distribution capabilities that support traffic flow management. Figure 1.1 illustrates that enhanced system. The key elements are described below.

- ETMS (Enhanced Traffic Management System): Prior to CDM, this was the only flight information database available to the ATCSCC. Ground delay actions taken by the ATCSCC have been traditionally based upon data from ETMS.
- ATMS (Advanced Traffic Management System): ATMS is the CDM counterpart of ETMS. ATMS has all the ETMS data and functions plus additional data and functions that are being prototyped. In particular, ATMS merges into the ETMS flight database real-time schedule updates flowing from the airline operations centers (AOCs). At the time of this writing, both ATMS and ETMS feed flight information to the ATCSCC
- A, B, and C strings: A “string” is a group of about ten computers that perform all the ETMS hub site processing; a string receives input data from many different sources, e.g., NAS, OAG (official airline guide), weather

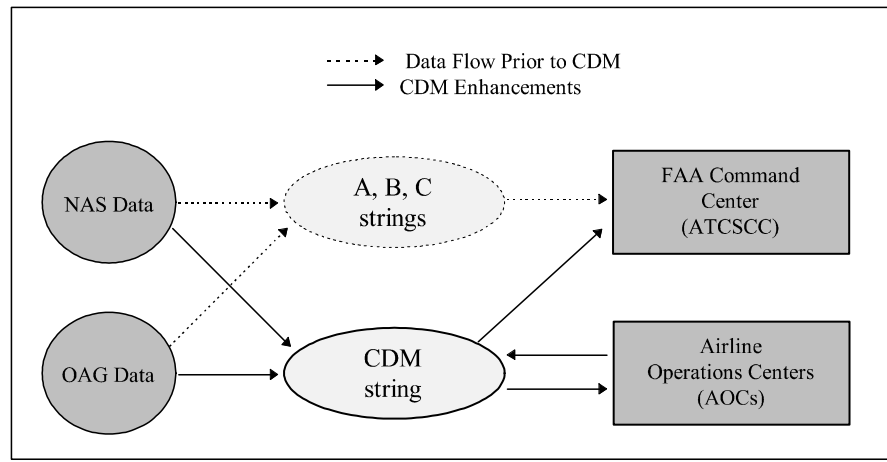


Figure 1.1: CDM Databases and Processing Strings

providers, specialists, and airlines, and then integrates this data in a database. This data is then made available to the various ETMS sites. There are two operational strings called A string and B string; there is also a test string called C string; these three strings should have the same data except for alterations to C string made as part of system tests. All three of these strings transmit ETMS data to the ATCSCC.

- CDM string: This is the string that receives all the input that is received by A, B, and C strings but that, in addition, receives the data feed of real-time airline schedule developed by the CDM working group.

1.2.3 The AOCnet

Figure 1.2 describes the flow of information through the AOCnet. The Volpe hub site (shown at the top) forms ADLs (aggregate demand lists) based on three sources of flight information:

1. the Airline Operations Centers (AOCs);
2. the National Airspace System (NAS);
3. the Official Airline Guide (OAG).

The ADLs are then pushed out over the AOCnet to the Air Traffic Control System Command Center (ATCSCC) and the AOCs. The ADLs are used by

both the ATCSCC and the AOCs to generate the information displayed through the CDM decision support tool, FSM.

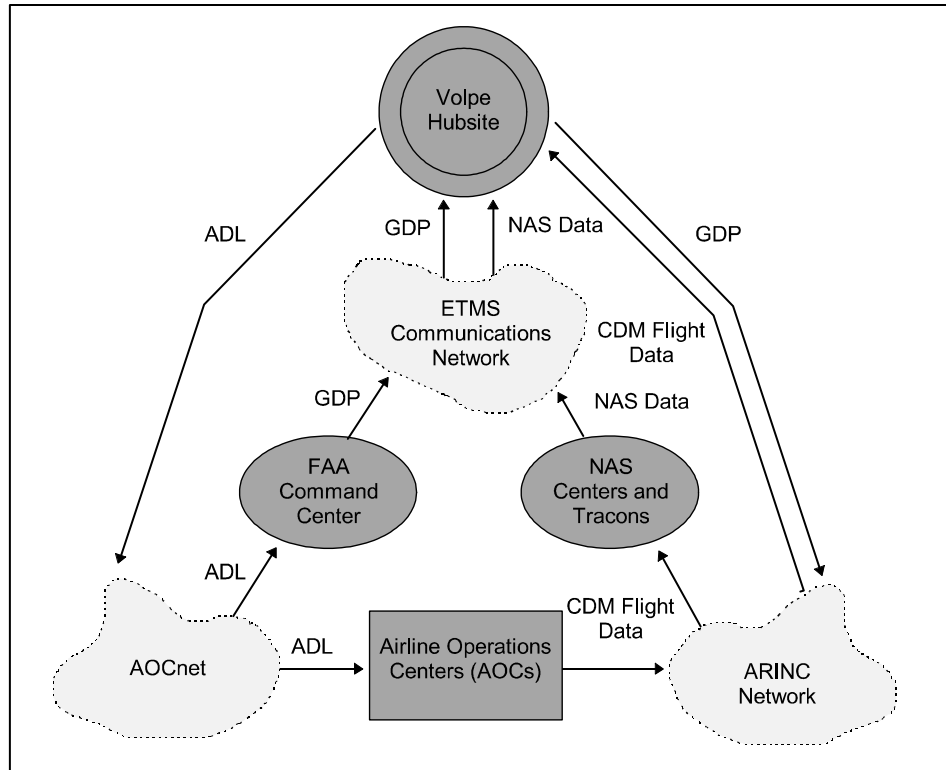


Figure 1.2: The CDM Flow of Information

2. Results of Analysis

2.1 Evaluation of Initial CDM Implementation

CDM has certain basic philosophical components that its proponents believe to be fundamental principles whose desirability is inherently obvious. These might be compared to the inherent desirability of free markets within a modern economy. In fact, to the extent that a general philosophical component of CDM is to increase the airlines control over their own economic decisions, the comparison to free market economies is a good one. If one believed in this point of view to the extreme, it would not be necessary to demonstrate that the important CDM effects have produced positive benefits but merely to demonstrate that the effects have taken place. With this in mind, we consider whether the initial implementation of CDM has led to operational changes consistent with its philosophical underpinnings. However, in addition, where practical, we also try to quantify specific benefits. Sections 2.1.1, 2.1.2 and 2.1.3 assess whether the initial implementation of CDM indicates that CDM can achieve certain fundamental operational changes. Section 2.1.4 assesses whether the initial implementation of CDM has produced improved GDPs. Specifically, estimates of delay reductions are provided. The source of the data used in our analyses is described in Appendix A. A detailed report of the ATCSCC interviews, which are referred to throughout this section, is given in Appendix F.

2.1.1 Changes in Quality of Information

As was stated in section 1.2 of this report, one of the original motivations for CDM was the desire on the part of the ATCSCC to receive up-to-date information on airline schedule changes and intentions. The merging of real-time airline operational information with FAA information has been a fundamental goal of CDM. In fact, this goal has clearly been achieved.

All CDM airline participants, including American Airlines, Continental, Delta, Northwest Airlines, Southwest Airlines, Transworld Airlines (TWA), United Airlines and US Airways have implemented data feeds from their operations systems

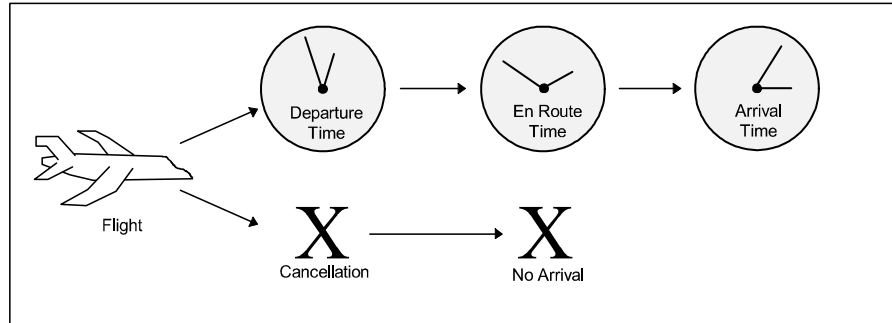


Figure 2.1: Key Elements in Predicting Arrival Demand

into the AOCnet. This information is merged with FAA-generated information by systems at the Volpe Center into a real-time data feed known as the “CDM String”. Thus, the answer to the question of whether CDM has resulted in the availability of a new real-time information source generated by the merger of FAA and airline information streams is, clearly, yes.

The second question to ask is whether this new information source is in fact an improvement over existing information sources. That is, does the merging of FAA data and airline-supplied data produce an information source of higher quality than existing FAA sources. This is a complex question, given that information can be used in so many ways. In addressing this question, we focused on the accuracy of elements essential to the prediction of arrival demand, since this is a key parameter in designing a GDP.

Figure 2.1 illustrates the key components in the prediction of the arrival demand profile at an airport. Since many flights are canceled on a day-to-day basis, the starting point is determining the number of flights that are canceled and the departure time of the remaining flights. Once this is done, an enroute time together with the departure time produces an estimated arrival time for each flight. The arrival times aggregated over all flights produce an arrival demand profile. With this in mind, we assessed the impact of CDM on data quality relative to the predictive accuracy of flight departure times, the timeliness of cancellation notices, the predictive accuracy of flight arrival times and the accuracy of estimates of arrival demand.

In assessing the predictive accuracy of both flight arrival time and flight departure time estimates, we used the IPE (Integrative Predictive Error) metric, which was

| GDP Days | | | | | | | |
|--------------|--------------------|-------|-----------------------|-------|-------|----------------------|-------|
| | Avg IPE-6 (min/ft) | | % of Flights Improved | | | Avg Improve (min/ft) | |
| | C | C | C | Equal | CDM | C | CDM |
| SFO | 30.86 | 27.21 | 15.89 | 39.23 | 44.88 | 14.44 | 13.27 |
| EWR | 35.40 | 28.19 | 21.69 | 31.63 | 46.69 | 06.21 | 18.71 |
| Non-GDP Days | | | | | | | |
| | Avg IPE-6 (min/ft) | | % of Flights Improved | | | Avg Improve (min/ft) | |
| | C | CDM | C | Equal | CDM | C | CDM |
| SFO | 12.51 | 09.74 | 08.92 | 68.00 | 23.08 | 04.21 | 09.58 |
| EWR | 14.72 | 13.34 | 08.96 | 67.49 | 23.56 | 12.13 | 10.99 |

Table 2.1: Comparison of C and CDM IPE-6 Values for Departure Time Estimates 2/2/98 - 3/16/98

developed by NEXTOR specifically for this purpose. The metric is described in detail in Appendix E. We specifically used the IPE-6 metric, which measures a stream of predictions coming in over a 6-hour period. The units of the IPE-6 metric are minutes. As an illustration of the meaning of the metric for an estimate of flight departure time, one could interpret an IPE-6 value of 15.5 minutes as, “over a 6-hour period, the flight departure time estimate had an average error of 15.5 minutes”.

Flight Departure Times

We wish to evaluate whether CDM enhanced information improves the accuracy of flight departure time predictions. To do so, we compare the predictive accuracy of the streams of departure time estimates provided by the CDM string and by the C string. (The CDM string is the real-time CDM data feed generated from the ATMS database; the C string is the ETMS-based test string that closely mirrors the ETMS-based operational strings that feed the ATCSCC). Table 2.1 gives a comparison of the IPE-6 metrics for the estimates produced by the CDM string and the C string. To provide a fair comparison, the numbers in the tables are based on those flights common to the two databases for which a valid 6-hour stream of data was present in both strings.

To illustrate the meaning of this data, one could interpret the first row of numbers

in Table 2.1 as follows.

On GDP days at SFO, the average value of the IPE-6 metric for C was 30.86 and the average value of the IPE-6 metric for CDM was 27.21 (lower is better). For about 39% of the flights, the IPE-6 metric values were equal for the two strings; for 45 % of the flights the CDM string was better and for 16 % of the flights the C string was better. When the CDM string was better, it was better by an average of 13.27 minutes; when the C string was better, it was better by an average of 14.44 minutes.

It is significant that, during GDP days, when accurate data is most important to guide flow management actions, the CDM string produces better information for a greater percentage of the time. This phenomenon is illustrated in Figure 2.2.

In summary, the results of this section indicate that, when averaged over all flights and over both non-GDP and GDP days, CDM-based information provides a moderate improvement in the accuracy of departure time predictions. It is significant to note that CDMs impact is greatest on GDP days, which is the time when such predictions are of greatest importance.

It is arguable that the accuracy of arrival time predictions is equally (if not more) important to flow management decisions than departure time predictions since arrival times have a more direct impact on accurate arrival demand profiles. We conducted IPE-6 studies on arrival time predictions and found no significant difference between the CDM and non-CDM cases. Certainly, further study is




Figure 2.2: Percentages of Flights at SFO for which CDM and C Strings Give More Accurate Predictions of Departure Time

warranted to understand this phenomenon. However, there is an obvious reason that at least partially explains this effect. Once a flight is enroute, arrival time predicting is done by ETMS-based algorithms that have (in theory) been unaffected by CDM. Thus, for a significant percentage of the cases covered by the IPE-6 metric, there is no difference between the CDM and the non-CDM mechanisms.

Cancellation Notices

Airlines routinely cancel flights in response to crew and equipment problems and, most significantly, in response to capacity reductions and traffic flow management actions such as GDPs. To illustrate the significance of cancellations, consider Figures 2.3 and 2.4. These illustrate, respectively, the variations in cancellation count by day and by hour. These statistics are notable relative to the number of cancellations as well as to the high degree of variability. Thus, it seems clear that accurate estimation of daily demand requires accurate and timely cancellation information.

To further illustrate the importance of cancellation information, we note that during the January through May 1998 period, there was at least one period of heavy cancellations on 79% of the days at SFO and on 62% of the days at EWR. A period of heavy cancellations is defined as a period of 4 consecutive hours with 28 or more cancellations.

We now compare the timeliness of cancellation notices before and after CDM. The method used for this analysis is rather complex and is described in Appendix D. Appendix C gives the process used to uniquely identify flights in the data files. Figure 2.5 gives the distribution of the time before OETD (original estimated time of departure) at which the cancellation was known with and without CDM for SFO. Estimation was required to determine when timeout notices would have occurred under ETMS without the presence of CDM. Note that for both the conservative default cancellation time of 0 minutes before OETD (top portion of Figure 2.5) and the more likely default cancellation time of 120 minutes after OETD (bottom portion of Figure 2.5), the effect of CDM is to push the distribution of notice time (before OETD) to the right, meaning that cancellations are generally known earlier under CDM. The results for EWR are very similar. The results at SFO and EWR can be summarized as follows:

During the January through May 1998 period, flight cancellation notices were received under CDM, on average, at least 76 minutes earlier at EWR and at least 63 minutes earlier at SFO, than they would have been without CDM. We estimate that a more likely Figure is that the average cancellation notice was received 113 minutes earlier at EWR and 93 minutes earlier at SFO (see Figure 2.5).

We view the improvement in cancellation notice times as very significant. Furthermore, timely knowledge of cancellations is vital in GDP planning. Improved cancellation information was cited by one of the ATCSCC specialists as the most important benefit of CDM.




Figure 2.3: The Number of Cancellation Notices Per Day at SFO, for the period January - May 1998

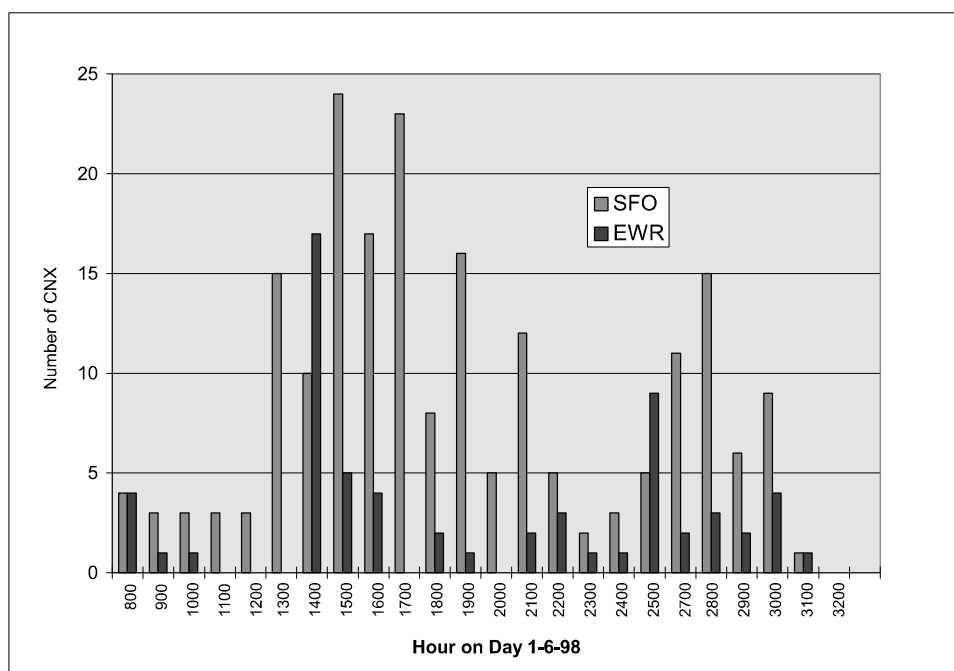


Figure 2.4: The Number of Cancellations per Hour at EWR and SFO on a Typical Day, 1/6/98.


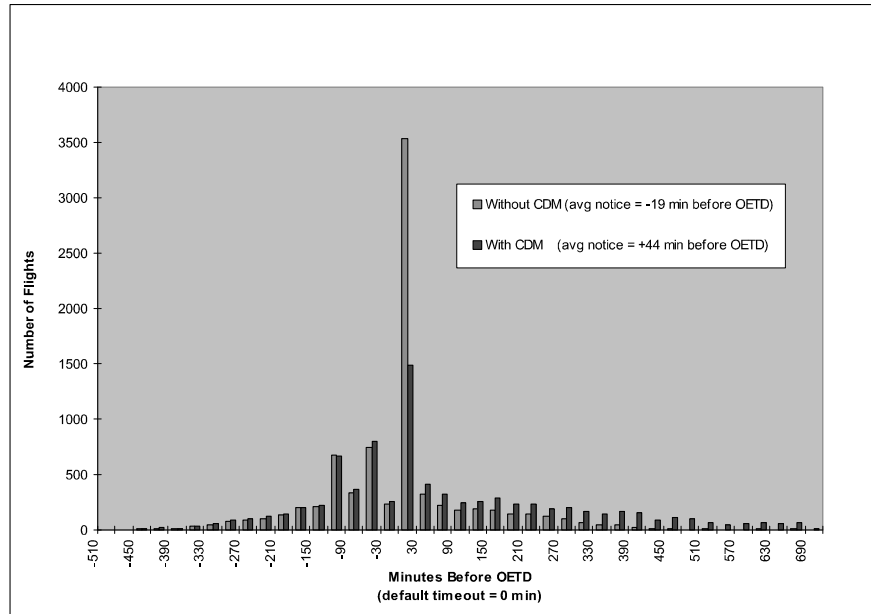


Figure 2.5: Distribution of Time Before OETD that First CNX Was Known at SFO, Jan - May 1998. (Top: conservative figures show a 63 minute improvement. Bottom: most likely figures show a 93 minute improvement.)

Arrival Demand Profile at an Airport

Accurate arrival demand forecasts are vital to the development and successful employment of ground delay actions taken by the ATCSCC. There is growing consensus among the Traffic Management Coordinators at the ATCSCC that the CDM-based demand profiles provided by FSM are more accurate than those provided by the ETMS-based system. We designed an analysis to assess the predictive performances of the CDM and ETMS-based systems with respect to arrival demand. Specifically, we measured the number of flights that would be predicted by each system to arrive at SFO during select hours against the number of flights that would actually arrive. The “actual” number of arrivals was determined by the traffic count worksheets maintained by the FAA at SFO. Since these traffic count worksheets are recorded directly by the air traffic controllers at the airport tower at SFO, we believe them to be accurate and independent of the systems under study.

Demand predictions for a given hour naturally vary over time as cancellations, delays and other such information are received. Thus, a demand prediction made several hours in advance of a future hour is bound to differ from the actual number of arrivals (traffic count). For instance, suppose that at 1200z, a system predicts that 50 flights will arrive during the 1500 hour (1500-1559z) and that, in fact, only 40 flights arrive. Then it seems that the system has over-predicted by 10 flights. If those 10 flights were canceled at 1000z, then (ideally) the system should have known about them by 1200z and has made a predictive error of 10 flights. However, if those 10 flights were not canceled until 1400z, then the system has made a “correct” prediction at 1200z, after adjusting for future demand changes.

Based on the data at our disposal, we found no practical way to separate predictive error in demand predictions from natural variation. Since most flights are more than one hour in duration, the status of flights that will arrive in the next hour should be stable and any difference between the actual number of arrivals and predicted number of arrivals should be primarily attributable to predictive error. Thus, in our analysis, we chose to capture demand predictions at a time “0 hours” in advance, meaning, at a time just before the start of the hour in question.

Over the period January 13 - February 25, 1998 we studied demand predictions made at SFO by the CDM and ETMS systems, as communicated over the “CDM” and “C” strings, respectively (there were no traffic count worksheets available for

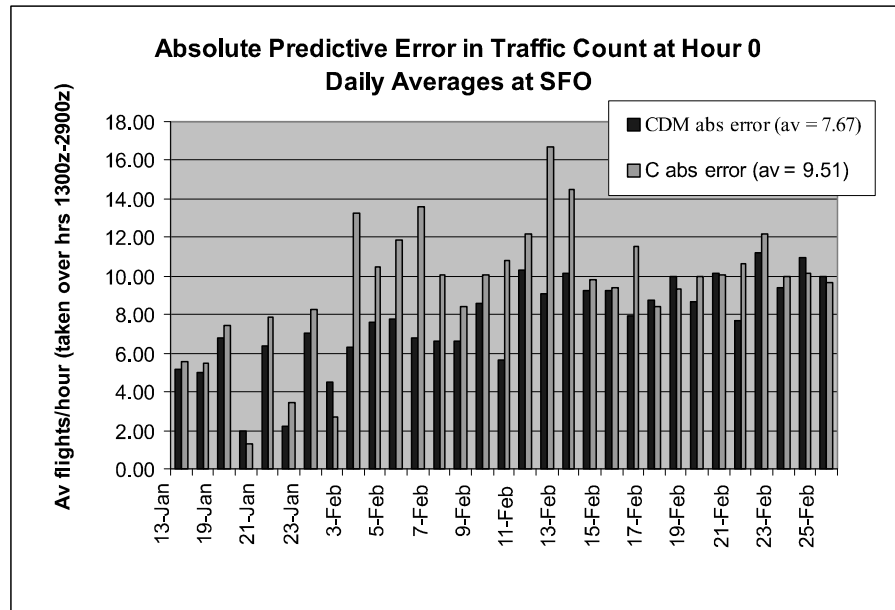


Figure 2.6: Comparison of Error in Prediction of Arrival Demand

EWR). For each day at SFO, the absolute predictive errors made just before each of the hours 1300 - 2900z were averaged to arrive at a single value for each day and under each system (CDM and C). The two daily values appear as juxtaposed bars in Figure 2.6.

We found that on 24 out of the 31 days, the average (over all hours) absolute predictive error made by the CDM string was strictly less than that of the C string. On these 24 days, the CDM string outperformed the C string by a margin of 2.57 flights, on average, while on days that the C string outperformed the CDM string, it did so only by 0.67 flights, on average. The (non-absolute) predictive errors averaged over all days were 7.67 and 9.51 for the CDM string and the C string, respectively. One could interpret this to mean that,

under CDM, the accuracy of the prediction for the number of flights that will arrive in the next 60 minutes has improved by approximately two flights per hour.

We should point out that both strings are consistently over-predicting the arrival demand on average and that, although the CDM string generally has less error associated with its predictions, its average error (7.67 flights per hour) is still a

substantial percentage of the average number of arrivals during any hour. The number of arrivals at SFO is typically in the range of 30-60 flights per hour, so this average error is (approximately) in the range of 12-25% of total arrivals.

Lastly, we have observed on a daily basis that the CDM string tends to be more sensitive than the C string, meaning that it often has greater variation over the day in the number of flights predicted to arrive during a fixed future hour. We consider this to be a positive feature since the ATMS (CDM) database is probably reacting to an influx of information that is not available from the ETMS database. We did not quantify this variation, however, since there is no base line of variation against which to measure the two strings. This requires a further study which should include an examination of the flight databases and a categorization of the information that each string is reacting to.

2.1.2 Impact of Information Distribution and Common Situational Awareness on Decision Making

Given that a new, improved information source has been generated, the natural follow-up question is whether this information is being made available to appropriate NAS users and air traffic managers and whether the availability of this information has led to improved decision making. The answer to this question is clearly yes.

In particular, the ATCSCC and all participating airlines have real-time CDM data feeds via the AOCnet. This information is accessible, via FSM, the CDM decision support tool, to all specialists in the ATCSCC and in each of the airlines operations centers (AOCs).

A second question to be asked is whether this new information source has impacted in a positive way the decision-making of the ATCSCC and of the airlines. Specialists at the ATCSCC have indicated that the information provided by CDM helps them to coordinate flow management actions such as ground stops and GDPs. In particular, they cite the usefulness of the cancellation information, which was previously unavailable.

We note that there was near-universal agreement among the ATCSCC specialists interviewed that GDP planning is more time consuming under CDM. They felt that this is due to the fact that the specialists are experiencing a learning curve under the new system and that working with more information and a wider range of options induces more procedures and requires a more extensive assessment of

options.

With regard to impact on the airlines, we have not attempted to independently quantify the extent of the improvements. Rather, we provide relevant descriptions provided by various airlines.

GDP Planning for Non-CDM Airports

- United Airlines has used information provided by CDM/FSM to help determine the number of flights to cancel at ORD (Chicago's O'Hare airport) under conditions of degraded capacity. According to United Airlines, on at least two separate occasions, the number of canceled flights was reduced by 25% over the number that would have normally been canceled; the estimated total cost savings for these two occasions alone was \$1.5 M [5][6].
- US Airways routinely uses FSM for GDP planning at airports where GDPs are not CDM-based. It can access any cancellations from FSM, including cancellations made by its affiliated commuter carriers and cancellations not reported to the US Airways ATC (air traffic control) desk. Using this information, which previously was not accessible in a timely manner, US Airways is able to run a substitution program, which involves regular US Airways flights and flights from its affiliated carriers [3].

Control and Management of Flow into Hub Airports

- Dispatchers at US Airways use FSM to plan operations based on the size and characteristics of arrival banks at their hubs [3].
- In times of very high demand, US Airways has used FSM to implement its own "internal GDP" to prevent gridlock at hub airports [3].
- Delta Airlines uses FSM to determine the necessity of diversions by estimating anticipated airborne delays for flights arriving at their hubs during peak periods. Delta reports that the more accurate information provided by FSM has allowed them to preserve the destinations of flights that normally would have been diverted to other airports [7].

Management of Flow over Arrival Fixes

- Both US Airways and Delta Airlines use the demand forecasts provided by FSM to adjust the balance and arrival priority of flights over arrival

fixes for various airports. In some cases, these changes require actions on the part of FAA controllers, which are jointly determined via airline/FAA discussions. Thus, the common viewpoint provided by FSM has facilitated joint FAA-airline decisions outside the domain of GDPs [3], [4].

Fuel Planning

United Airlines, US Airways and TWA report using FSM to anticipate airborne delay and the subsequent fuel requirements for their aircraft [3], [4].

While these reports do not represent a comprehensive study, we can conclude that

CDM has had a definite positive impact on airline decision-making. What is particularly surprising is that CDM information has been used to support a variety of decision-making outside the initial intended application area of ground delay programs.

2.1.3 Assessment of Increase in Ability of Airlines to Make Economic Resource Allocation Decisions and to Solve Capacity-Demand Imbalances

One of the fundamental principles of CDM is that, ideally, the FAA should identify any potential demand-capacity imbalances and give the airlines the opportunity to determine how to reduce or reallocate demand so as to satisfy all capacity constraints. Thus, the airlines can take into account their own economic considerations in determining the appropriate reallocation of resources. Of course, this ideal cannot always be achieved, since in many cases there are several airlines competing for the same resource; there must be an intermediary (the FAA) who rations the resource among the competing parties. However, one goal of CDM is that the process of FAA constraint identification followed by voluntary airline demand reduction should solve problems, thereby reducing the number of times when the FAA has to take explicit action (e.g., running a ground delay program). While it is sometimes difficult to quantify the benefits associated with voluntary demand reduction when compared with demand reduction mandated by the FAA, a strong philosophical tenet of CDM is that it is inherently preferable that the airlines are able to solve a problem without FAA intervention. Below we describe some evidence provided by airlines, that

indicates that the initial implementation of CDM has resulted in a reduction in the need for the FAA to take explicit control actions.

- Continental Airlines reports that on several occasions, they have used the demand predictions provided by FSM to reduce demand at EWR to a level that eliminated the need for a GDP [4].
- TWA reports using FSM to plan its cancellations for GDPs in St. Louis; they feel this has led to more effective matching of demand and capacity and has avoided the need for ground stops on certain occasions [2].

2.1.4 Delay Reduction and Other Benefits Associated with CDM-based Ground Delay Programs

A CDM-oriented GDP differs in several ways from a traditional GDP. The primary difference is that both the airlines and the ATCSCC use the information exchange and decision support capabilities in FSM to arrive at their GDP decisions. Some of the important impacts of this approach are:

- 1) The ATCSCC and the airlines have more accurate demand information. This information can allow the airlines to respond to demand-capacity imbalances independently of the ATCSCC, which, in some cases, can eliminate the need for the ATCSCC to implement a GDP. This information enables the ATCSCC to better assess the need for a GDP and to set GDP parameters.
- 2) The ATCSCC gives the airlines timely information on proposed actions thus allowing the airlines to plan their responses to such actions.
- 3) The “power run” feature of FSM allows the ATCSCC to explore alternative GDP parameters, which can lead to a more effective balance between the arrival and acceptance rates at an airport.
- 4) FSM allows the ATCSCC to postpone the decision to implement a GDP until such a time that more accurate weather information is available. This feature can enable the ATCSCC to cancel the initiation of a GDP based on new weather information.
- 5) FSM allows the ATCSCC to revise a GDP after its inception. Such revisions either reduce ground delays when weather conditions have improved or assign additional ground delays if weather conditions have degenerated or demand has been underestimated. In the former case, a revision has the potential to reduce

ground holding while, in the latter case, it has the potential to reduce airborne holding and the need for diversions.

6) The compression algorithm fills the vacant arrival slots that arise when an airline cancels a flight or substitutes a flight forward but is unable to fill the vacated slot with another flight due to schedule constraints. The use of compression results in a reduction in the total assigned ground delay.

7) Together with the exchange of information between the airlines and the FAA, the ration-by-schedule and compression algorithms provide the airlines greater flexibility in the redistribution of assigned ground delay among their flights. In turn, this enhances the ability of the airlines to reduce their respective delay costs.

Items 1), 2) and 4) can all lead to a reduction in the number of GDPs implemented by the ATCSCC. It is difficult to estimate the amount of this reduction under CDM and even more difficult to quantify the savings associated with such a reduction. For example, suppose that the dominant airline at a weather-impacted airport has eliminated the need for a GDP through cancellations and diversions. It could be that the airline has canceled or diverted more flights than it would have if the GDP were in effect. However, the airline may have actually minimized the detrimental impact of the capacity restriction on their operations by allowing them to take action early on. It is clear that the benefits of such an independent airline action could be computed (with considerable difficulty) only by the airline itself. Although we were not able to quantify the benefits in this area, we were able to determine from information provided by the airlines and the ATCSCC that CDM did result in instances where GDPs, which would otherwise have been implemented, were not implemented. Section 2.1.3 cites such instances identified by airlines. In addition, a specialist at the command center cited two instances in which a GDP had been canceled prior to its inception because FSM showed that arrival demand had dropped to a safe level through airline flight cancellations.

Thus, we conclude the following.

CDM has led to instances in which a GDP that ordinarily would have been run was totally avoided. These GDPs were determined to be unnecessary either because the airlines had sufficiently reduced demand or because CDM allowed postponement of the GDP decision until a better weather forecast was available.

Items 1), 3), 5) and 6) all can help the specialists from the ATCSCC design better GDPs. The quality of a GDP cannot be measured easily. The design of a GDP can impact assigned ground delays, the amount of possible airborne holding in the airspace adjacent to the destination airport, and the throughput at the destination airport (number of flights that land per unit time). To illustrate some of the tradeoffs, if a specialist assigns too much ground delay, then the rate of aircraft arriving at the airport is too low and the airport is “starved”. On the other hand, if too little ground delay is assigned, then the rate of aircraft arriving at the airports airspace is too high and airborne holding is necessary to reduce the flow of aircraft landing at the airport. This leads to unnecessary airborne delay and, eventually, to possible flight diversions.

The ATCSCC specialists used FSM to design essentially all GDPs for SFO and EWR during the period of study. In addition, FSM was regularly consulted for information in designing GDPs for other airports. The consensus opinion seems to be that FSM allows the specialist to generate more effective GDPs.

A unique feature of CDM-based GDPs is the use of the compression algorithm. This allows the specialist to fill holes in the sequence of assigned arrival time slots created by cancellations. This has the net effect of reducing the total assigned ground delay. Furthermore, it has the benefit that the airlines can implement tradeoffs in delay not previously available to them. For example, compression allows an airline with a relatively small presence at an airport to cancel an early flight and, in return, receive a reduction in the delay of its flights much later in the schedule. As one way of quantifying the distinct impact of CDM, we estimated the total savings in assigned ground delay achieved by compression during the initial CDM implementation.

In assessing savings obtained by compression, we calculated the net delay reduction achieved by the compression algorithm each time it was used in the first quarter of 1998.

In order to measure the net reduction in delay, we first calculated the total savings in assigned ground delay achieved by the compression algorithm. However, a portion of these savings could also have been obtained had the compression algorithm not been run.

More specifically, each airline could have attempted to fill the holes in their own sequence of assigned arrival slots using additional substitutions. Therefore, we consider the net delay reduction to be the total savings in assigned ground

delay minus the savings that each airline could have made on its own. In other words, the net delay reduction might be viewed as savings that are unique to CDMs collaborative resource allocation mechanisms. This was accomplished by running FSM (V1.05) in historical mode on archived data.

Figure 2.7 shows the results for all executions of the compression algorithm at EWR during the period of January through mid-April. Similar results were obtained for SFO. We can summarize as follows:

During the period of January through mid-April, the use of compression resulted in an average reduction in assigned ground delay of approximately 13% at SFO and 12% at EWR. Slightly half of this reduction could have been obtained by the airlines through substitutions. The remainder could only be obtained using compression. Using an industry-accepted value of \$25 per minute of delay, the compression savings was \$26,000 per GDP at San Francisco and was \$29,000 per GDP at Newark. The respective average monthly savings were \$269,000 and \$100,000.

It should be noted that we have not included the results for the period of mid-April through June. The reason for this is that, starting April 19, 1998, airlines were allowed to exempt themselves from compression. The open arrival time slots of exempted airlines were not available to the compression algorithm to be filled with other flights. United Airlines at SFO and Continental Airlines at EWR were exempt from compression. A more involved analysis is required to assess compression under these conditions. Preliminary results show that savings are much lower in this case. This is to be expected as the exemptions reduce the options available to the compression algorithm. The reasons for allowing exemptions are due to limitations of the current substitution procedures. We should note that the apparent negative impact that exemptions have on the overall savings indicates that allowing airlines to vary their degree of participation in CDM procedures can have significant adverse impacts on the overall benefits. Thus, it is important to develop rules and standards for participation. Plans call for eliminating exemptions when newer substitution procedures are available (known as "Simplified Substitution Rules").

It is well known that delays to a single flight can propagate throughout a carriers operations causing sometimes-substantial downstream delays. This occurs because of the multiple connectivities related to equipment and crews within the




Figure 2.7: Delay Savings at EWR Due to Compression (as a percentage of the total assigned ground delay)

operations of the major US carriers. For example, suppose flight ABC200 received a delay of 60 minutes and that the aircraft for ABC200 was designated to be used on ABC320, which was scheduled to depart 45 minutes later than the scheduled arrival time of ABC200. If the minimum time needed to turn the aircraft were 30 minutes, then 15 minutes of the 60-minute delay would be absorbed by the 15-minute buffer that was built in between the arrival of ABC200 and the departure of ABC320. However, the remaining 45 minutes of delay would be transferred to ABC320. Thus, the 60-minute delay of flight ABC200 would cause a 45-minute downstream delay of flight ABC320, for a total of 105 minutes of delay incurred by airline ABC. If this were the only downstream delay effect, then we could associate a multiplier of 1.75 with the initial delay. That is, the initial “base delay” of 60 minutes would lead to a $1.75 \times 60 = 105$ minute total system delay.

In conjunction with the CDM project, American Airlines and Oak Ridge National Laboratories undertook an effort to quantify the delay multiplier [1]. Figure 2.8 represents a data excerpt from that study provided by Oak Ridge. This figure shows the delay multiplier for a set of American Airlines flights into Dallas-Fort




Figure 2.8: Delay Multiplier as Function of Base Delay (based on American Airlines operations into DFW for one day)

Worth airport. The study calculated multiplier values as a function of both the initial base delay and the flight departure time. In the figure, the vertical axis gives the multiplier value and the horizontal axis gives the initial base delay. The multiple curves correspond to separate, initial flight departure times. Note that, in all cases, the multiplier value grows with the base delay, since it is harder for larger delays to be absorbed by the system. Also, note that the rate at which the multiplier grows decreases throughout the day because later in the day there are fewer downstream flights to be impacted. The multiplier values can be quite substantial, often exceeding 2.0 and, for some times of the day, getting close to 4.0.

When considering the impact of delay savings, one should keep in mind this multiplier effect, which implies that a certain amount of base delay savings is converted into a greater system-wide delay savings. Delay multipliers can be used to estimate downstream savings due to compression. The results of the American Airlines-Oak Ridge National Labs study are dependent upon both the

| | GDP Days | Non-GDP Days | All Days |
|--------------|----------|--------------|----------|
| Jan-Mar '97 | 2.07 | -0.45 | -0.83 |
| Jan-Mar '98 | 7.88 | 2.11 | 0.26 |
| November '97 | 6.23 | 1.65 | 1.17 |

Table 2.2: Airborne Delay Averages for Three Periods

connectivity characteristics of the American Airlines hubbing operation and the flights selected for the hypothetical assignment of initial delay. Until estimates of delay multipliers based on a broader range of flights and airlines can be found, we did not think it appropriate to apply these numbers directly within this analysis. Nonetheless, we feel this study strongly suggest that the actual delay savings due to are several times higher than those that we have reported.

As mentioned earlier, there is a complex set of interactions among assigned ground delays, airborne holding and the rate of arrivals into the airport. Of particular concern is the possibility the compression savings in ground delay is merely a transfer of ground delay into airborne delay. In the event that the arrival demand at an airport were equal to the capacity (airport acceptance rate), then decreasing ground delay would increase the rate of flights into the airport, and the reductions in ground delay would be converted directly into increases in airborne delay.

In order to investigate this issue, we carried out an analysis of airborne delays before and after CDM. We have been able to develop new procedures for estimating airborne delays that are described in Appendix B. Unfortunately, due to the sensitivity of delays to weather and other factors, we have been unable to determine within the time frame of this study the impact of CDM GDP procedures on airborne delays. Table 2.2 shows our estimates of airborne delay over three periods.

The first period (January-March '97) is prior to CDM prototype operations while the second period (January-March '98) is during CDM prototype operations. Given the rise in average airborne delay on GDP days from the first period to the second period (from 2.07 Min. to 7.88 Min.), one might be inclined to conclude that CDM has had a negative impact on airborne delays. However, note that there was also substantial airborne delay (6.23 minutes) on GDP days during

the control period, November of 1997. Moreover, the average airborne delay on non-GDP days and on all days was substantially higher in November of 1997 than during the first period. This strongly suggests that the increase in airborne delays from period 1 to period 2 may be due to factors other than CDM, such as an increase in the number of scheduled arrivals and adverse weather conditions, e.g., El Nino. Complex, data-intensive analyses would be required to determine the exact causes of the increases in airborne holding.

The observations of airline operational personnel and ATCSCC specialists are that CDM-based GDPs have not led to any significant changes in airborne delays. We concur with this, based on analysis of a small number of individual days. Thus, we feel that, for most flights, the savings in assigned ground delay have been realized, meaning that earlier departures have actually led to earlier arrivals.

In the future, it is important to carefully monitor assigned ground delay, airborne delay and the arrival throughput at the airport in assessing the quality of ground delay program procedures.

2.2 Assessment of Future Benefits

There is clear evidence of strong potential for future CDM benefits. Below we list elements from which future benefits could be derived.

- Movement out of prototype operations: It should be recognized that to date CDM is still in a prototype phase. Thus, once the system goes into operation, based on a refined set of operational procedures and software, one can expect more consistent performance and greater user benefits.
- Extension of CDM-based GDPs to all major US airports: For most of the period covered, CDM-based GDPs have been generated only at EWR and SFO. Recently, LaGuardia (LGA) and St. Louis (STL) airports have been added. However, these airports still represent a relatively small subset of the US total.
- Improved ability of airlines to take advantage of CDM capabilities: CDM features are geared to provide the airlines with much greater control over resource allocation and decision-making. Comments from the airlines indicate that, in many cases, the procedures are not in place for the airlines to take full advantage of the control offered by CDM.
- Improvements in CDM-based GDPs through better data quality and new FSM features: CDM-based GDPs depend critically on accurate data from both FAA sources and the airlines. We have determined that, while the initial implementation of CDM clearly has increased the predictive accuracy of the data, there is also much room for improvement. As the data quality improves we expect improvements in the decision-making based upon it. In fact, the airlines are currently working to improve the accuracy of schedule update information sent to the CDM string. The quality of CDM-based GDP planning also depends on the decision support capabilities provided by FSM. The procedures embedded within FSM as well as other CDM procedures are being enhanced and revised based on the initial experience.
- NAS Status, Collaborative Routing: Future plans call for applying CDM concepts in other contexts. Efforts are already underway in the areas of the distribution of NAS Status information and in collaborative routing.

2.3 Effectiveness of CDM Components

In considering whether certain aspects of the initial CDM implementation have been more effective than other aspects, we have broken down the initial implementation into the following four elements.

1. GDP-E: The new set of operational concepts and automatic procedures for computing ground delay programs.
2. FSM: This decision support tool allows for the display and manipulation of information in new and more flexible ways.
3. Improved Information: Real-time feed and database derived by combining ETMS data with airline-supplied data.
4. Environment: Improved relationship and mutual understanding between FAA and NAS Users.

It is quite clear that all of these elements are intimately related. For example, in order to make use of the new procedures for constructing GDPs, specialists at the ATCSCC need a decision support tool and improved information. Furthermore, the improved relationship between the ATCSCC and the airlines has been fundamental in getting the GDP procedures to work on a practical level. On the other hand, when viewing CDM in this fashion, it might be possible to note that one of the above four elements dominated the positive aspects of CDM or that one the four elements was particularly disappointing and should be eliminated or replaced. After seriously considering these questions we have concluded that:

The success of the initial implementation of CDM is intimately tied to the combination of the new GDP-E procedures, the effectiveness of the decision support tool, FSM, the information derived by combining airline and FAA data sources and the improved relationship between the FAA and the airlines. While attention should be paid to improving each of these elements, none is noticeably stronger or weaker than the others in contributing to the overall success.

2.4 Issues Deserving Special Attention

Below we describe issues that we feel deserve special attention in future CDM planning and development.

Information Quality:

CDM has clearly improved information quality. However, it is also clear that there is room for substantial additional improvement. Effective use of the power of FSM requires accurate information. Significant resources should be invested into improving information quality. Effort is needed on the part of all parties. Good data quality requires proper reporting and formatting of the data by the airlines; proper error checking, data collection and data modeling on the part of the FAA (and Volpe Center) and proper error checking and analysis by FSM.

Formalization of Procedures:

CDM-based GDPs are much more complex than their predecessors. As such it is important that they be run in a consistent, well-understood manner. This will allow the airlines to develop standard methods for reacting to FAA actions and will allow the FAA to develop smooth procedures for carrying out their activities. For example, there should be well-defined schedules for running compression and clearly defined time intervals for airline responses. In the context of CDM, it is important that the overall procedures continue to be fair to all participating parties. For example, there should be precise rules defining participating airline exemptions from CDM procedures. Opportunities for airlines to unfairly “game” the system should be identified and eliminated. All of the preceding points call for more formalization and standardization of CDM GDP procedures.

3. Summary and Conclusions

To summarize, we now answer the questions posed in at the beginning of this report.

Q: Has the initial implementation of CDM led to substantive, beneficial operational changes?

The use of CDM has produced new information by combining FAA and airline data sources. Furthermore, this information is definitely of higher quality. The most dramatic area of information quality improvement is in the timeliness of cancellation notices. For example, during the January through May 1998 period, flight cancellation notices were received, on average, at least 76 minutes earlier at EWR and at least 63 minutes earlier at SFO, than they would have been without CDM. We have measured significant improvements in information quality in other areas as well. We have observed, however, that there is still room for significant improvement in this area.

Through the AOCnet, the CDM-enhanced information has been distributed in an unprecedented fashion, giving airline operational personnel and FAA traffic flow managers a common view of air traffic in the NAS. On at least two separate occasions United Airlines used CDM-based information to reduce the number of flights canceled in anticipation of a GDP by 25% over the number that would have normally been canceled; the estimated total cost savings was \$1.5 M. There have been instances where airlines have solved capacity-demand imbalances by reducing demand in response to CDM information, thereby eliminating the need for an FAA action.

We have also found that the airlines have used this information for a variety of purposes outside the intended application domain of GDP planning. For example, airline operations managers have used CDM-supplied information to support fuel planning, diversion decisions and management of flow into hubs.

Q: Have the CDM-based GDPs produced tangible benefits?

There was consensus among ATCSCC specialists interviewed that CDM procedures produce more effective GDPs through richer, more timely information,

through the ability to revise a GDP once it is in effect, through the ability to weigh strategies, and through the compression algorithm. The compression algorithm is a procedure unique to CDM. It eliminates vacant slots and reduces overall delays by altering the assignment of slots to airlines in a way that treats all airlines fairly. During the period of January through mid-April, the use of compression resulted in an average reduction in assigned ground delay of approximately 13% at SFO and 12% at EWR. Slightly half of this reduction could have been obtained by the airlines through substitution processes available without CDM. The remainder could only be obtained using compression and, thus, represents savings only obtainable under CDM.

It is important to emphasize, however, that the major component of CDM benefits for GDPs should lie in the reduction of the cost of delays to airlines. One of the principal objectives of CDM is to provide each airline with greater flexibility to control the allocation of delay among their flights. In this manner, an airline can reduce the cost of delays by shifting delay away from flights where the delay has the most detrimental or costly impact. As an indication of the benefits in this area, we note that United Airlines has reported that it has achieved significant delay reduction based on the use of GDP enhancements at SFO and EWR and also the use of FSM to plan its responses to GDPs at ORD. They estimate the total savings between the January 20th start of CDM and March 2nd to be 11,000 minutes with a value of \$3 to \$4 M [6].

Q: What are the prospects for future CDM benefits?

The quantification of future benefits is outside the scope of our analysis. We do feel that there are good prospects for future benefits far greater in magnitude than those observed so far. Future additional benefits should accrue from the following:

- movement out of prototype operations
- extension of CDM-based GDPs to all major U.S. airports
- improved ability of airlines to take advantage of CDM capabilities
- improvements in CDM/GDP-E through better data quality, new FSM features, etc.
- application of CDM to other areas, including distribution of NAS status information and collaborative routing.

One should remember that CDM is still in a prototype phase. Thus, one should expect the bulk of its benefits to be evident in the future. At the same time, it is important to evaluate the results of the prototype to test whether it exhibits the stated promise of CDM. This is essentially the purpose of our analysis and report.

Q: Have certain components of CDM been more effective than others?

We considered the four major components of the initial CDM implementation to be:

- CDM-based GDPs;
- CDM decision support tool (FSM);
- Improved information;
- ATM decision-making environment.

Our conclusion is that none of these components has had a substantially greater impact than any other and that all have played a significant role.

Q: Are there issues that deserve special attention or new dedication of resources?

Data quality: CDM-based GDP planning requires accurate data and predictions. CDM-enhanced information is definitely better than the information provided by prior systems. However, on a more objective level, we feel that the quality of the data can and should be better.

Formalization of CDM procedures: there is a need to standardize and formalize CDM concepts of operations and procedures. This will provide smoother operations and interactions between the FAA and airlines and reduce the likelihood that airlines will be able to unfairly “game” the system.

Final Thoughts

Our analyses of the prototype phase to date indicate that the initial implementation has clearly demonstrated the promise of CDM. Even in a prototype mode, CDM procedures have had a significant positive impact on the decision-making of the airlines and the FAA. Furthermore, this improved decision-making has led to substantial quantifiable savings. What may be most surprising is that CDM data and tools have been used to support decisions outside the realm of the initial

intended application domain of ground delay programs. While the results of the prototype phase are optimistic, we must extend a cautionary note. CDM is in its infancy. It is vitally important that both the FAA and airlines maintain a strong commitment to improving the operational concepts, computer based tools and data sources. These require substantial refinement before widespread, consistent benefits will result.

Appendix A. Data and Source

The analyses for this report were based on “delta files” obtained from Metron Scientific Consulting Co., a prime contractor on the CDM project, and the developer of FSM (see figure A.1 for a listing of the delta files).

Each delta file spans one day and one airport and is based on ADLs compiled from one of two databases: ETMS (over the C string) or ATMS (over the CDM string). Below, we explain the structure of the ADL and delta files.

ADL: (ASCII text file) An aggregate demand list (ADL) is an airport-specific list of data for flights bound for that airport over a rolling time horizon, typically 12 hours. The ADL is comprised of a sequence of flight records (one for each flight) that contain information such as a flight call sign, ETD (estimated time of departure), ETA (estimated time of arrival), and so on. A fresh ADL is distributed by Volpe approximately every five minutes over the AOCnet to the ATCSCC, the airline operations centers and to Metron.

Delta Files: (ASCII text file) Using FSM, Metron has archived ADLs for the prototype operations airports since the inception of the AOCnet. A daily file is archived for each airport. In order to reduce the size of the archived files, FSM substitutes for each ADL an *update*, which is an abbreviated form of an ADL. Like the ADL, the update is a list of flight records. However, a flight is included in the update only if some of the information on the flight (e.g., a new ETD) has changed from the previous update (ADL). Thus, a delta file is an FSM-processed text file that contains a collection of changes (deltas) in ADL information for a given day and a given airport.

| Airport | Underlying Database | Dates (one file per day) | Type |
|---------|---------------------|--------------------------|-------|
| SFO | ETMS/CDM String | Jan 13 - May 30 | Delta |
| SFO | ATMS/C String | Jan 13 - May 30 | Delta |
| EWR | ETMS/CDM String | Jan 13 - May 30 | Delta |
| EWR | ATMS/C String | Jan 13 - May 30 | Delta |

Table A.1: Delta Files