IMPLEMENTING AND EVALUATING ALTERNATIVE AIRSPACE RATIONING METHODS

by

Jason Matthew Burke

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Master of Science
2002

Advisory Committee:

Professor Michael O. Ball, Chair
Dr. Robert L. Hoffman
Assistant Professor David J. Lovell

ABSTRACT

Title of Thesis: IMPLEMENTING AND EVALUATING ALTERNATIVE AIRSPACE RATIONING METHODS

Degree candidate: Jason Matthew Burke

Degree and year: Master of Science, 2002

Thesis directed by: Professor Michael O. Ball
Institute for Systems Research

While airport congestion has long been viewed as a major air traffic management problem in the United States, congestion in the en route airspace is drawing an increasing amount of attention. Sources of en route congestion, such as severe weather, often cause the Federal Aviation Administration (FAA) to delay and reroute aircraft in order to ensure safety. Most current research into methods for managing en route congestion seeks to reduce delay or aid in aircraft rerouting. However, there has been less attention paid to delay and reroute allocation methods – an area in which there appears to be a pressing, practical need.

Collaborative Decision Making (CDM) is a movement within the air traffic management community that has combined the interests of the FAA and industry to develop a universally-accepted resource rationing process for congested airports. There is high expectation that CDM can achieve similar success in developing a parallel rationing process for the en route airspace.

The CDM-inspired research underlying this thesis led to the development of a software tool, the En Route Resource Allocation Prototype (ERAP), that supports the analysis of alternative en route airspace rationing methods. In this thesis, we define a basic en route traffic flow management scenario, conduct experiments, and derive ERAP results which provide insight into rationing resources in the en route airspace domain. It is hoped that ERAP can serve as a baseline for future comparison and help lead to final industry acceptance of an ideal rationing solution.

DEDICATION

To my father, a model of altruism and perseverance

ACKNOWLEDGMENTS

Above all, I would like to thank my advisor, Dr. Michael Ball, for his guidance and extreme patience throughout the course of developing this thesis. I extend my gratitude to my unofficial advisor, Dr. Robert Hoffman, for his insight and for the countless hours he spent helping me to plow through the details. I would also like to thank Dr. Jasenka Rakas for helping to develop the criteria for initial traffic class implementation and Dr. David Lovell for participating in my thesis committee. In addition, I thank everyone at the NEXTOR lab for their assistance, support, and friendship throughout the past two years.

Special acknowledgement is due to Mike Feldman, Mark Klopfenstein, Travis White, and others at Metron Aviation, Inc. for providing and supporting software items that are essential to this thesis. Metron contributed the C-Flow software that is used in ERAP for displaying flight tracks. Metron also granted access to their POET database from which all flight data in this thesis is derived.

This work was supported in part by the Federal Aviation Administration through NEXTOR, the National Center of Excellence for Aviation Operations Research.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

AOC: airline operational control center

ATM: air traffic management

ARTCC: air route traffic control center

ATCSCC: Air Traffic Control System Command Center

CCFP: Collaborative Convective Forecast Product

CDM: Collaborative Decision Making

CDR: Coded Departure Route

CR: Collaborative Routing

CRCT: Collaborative Routing and Coordination Tools

CTA: controlled time of arrival

CTD: controlled time of departure

ERAP: En Route Resource Allocation Prototype

ETMS: Enhanced Traffic Management System

ENAD: Equity via Net Arrival Delay

FAA: Federal Aviation Administration

FACET: Future ATM Concepts Evaluation Tool

FCA: flow constrained area

FSM: Flight Schedule Monitor

GA: general aviation

GDP: ground delay program

GUI: graphical user interface

LAADR: Low Altitude Arrival and Departure Routes

NAS: National Airspace System

POET: Post Operations Evaluation Tool

RBS: ration-by-schedule

TOAD: time-ordered accrued delay

TFM: traffic flow management

Chapter 1.

Introduction

Airlines generally do not account for air traffic congestion when planning their flight schedules [1]. As business enterprises, they must schedule "aggressively" in order to compete financially and meet customer demand. However, unpredictable events such as convective weather can severely limit the capacity of entire regions of the National Airspace System (NAS). When such events occur, the demand for the use of en route airspace often exceeds capacity, and, in the interests of safety, the Federal Aviation Administration (FAA) must counter the imbalance by imposing delays and/or rerouting aircraft.

Delays cost the airline industry and its passengers an estimated 5.4 billion dollars in the year 1999 [18], and roughly 70 to 75 percent of all airline delays are caused by weather [1]. These two facts alone make it obvious that improvements in handling weather-induced airspace congestion would generate significant benefits for the air transportation industry and its passengers.

There is a substantial amount of research currently dedicated to this problem. Industry, government, and academic agencies are developing various software tools, technologies, and operational procedures to help improve the problems of en route

airspace congestion. A recent movement in the air transportation industry, called Collaborative Decision Making (CDM), is a mode of problem solving that has shown proven potential for uniting the interests of all the various stakeholders. The problem solving methods of this thesis are influenced by a sub-activity within CDM called Collaborative Routing (CR) that is responsible for improving the en route airspace congestion problem.

Many CR improvements are already in place and many others are in development. In general, these enhancements are designed to reduce delays or enhance the function of routing aircraft. However, there is not yet a universally agreed-upon method for the actual allocation of delays and reroutes to specific flights. The purpose of this thesis is to incorporate concepts proposed by the Long-Term CR Group into a highly transparent software prototype that can be used to further the goals of CR and help lead to future agreement upon a resource rationing algorithm. This research represents an important stepping stone in the development of CR, as it is the first implementation of newly proposed alternative concepts for rationing en route resources.

1.1 Collaborative Decision Making

CDM is a joint FAA-industry initiative that began in the mid-1990s. In general, CDM represents a symbiotic relationship of sharing near real-time operations information between the FAA and airline operational control centers (AOCs) in order to improve the NAS. CDM is one of the key tenets in the FAA's Free Flight program, which is in the process of redefining the FAA's role in air traffic management (ATM). The long-term goal of Free Flight is to give airlines near-total control over their operations. For further

information regarding Free Flight, see the program's official web site (http://ffp1.faa.gov/home.asp).

CDM really amounts to a philosophy. It shifts the role of the FAA in ATM from an absolute control authority to a service provider. It asks what information can be shared and what mechanisms and operations can be enforced in order to promote safer, more efficient, and more equitable usage of the NAS. The CDM philosophy recognizes that the sharing of accurate information is necessary for competent decision making, and it provides incentives for user participation. It distributes appropriate airline operations decisions to the airlines, and it attempts to make the best overall use of independent airline decisions to increase the net benefit for all NAS users.

The need for CDM became apparent as a means for reducing inefficiencies in ground delay programs (GDPs). A GDP is a standard ATM practice used during periods of congestion to reduce the incoming air traffic demand for a specific airport. A GDP is enforced by delaying flights destined for the chosen airport at their airports of origin. The premise is that delaying flights on the ground reduces the workload upon air traffic controllers and saves fuel that would otherwise be wasted in an airborne holding pattern.

The original GDP operations paradigm was highly dysfunctional [21]. The resource rationing algorithm, called Grover-Jack, was proven to be inequitable, and it actually discouraged the airlines from providing accurate information. For example, due to a misaligned incentive structure, airlines would neglect to notify the FAA of flight delays or cancellations, and valuable airport arrival slots would often go unused.

The advent of CDM brought about major changes that dramatically improved the efficiency of GDPs. The FAA and the airlines agreed upon resource rationing algorithms

called compression and ration-by-schedule (RBS) as equitable means for improving GDPs. Compression is an algorithm that credits airlines for reporting delays and cancellations and allows for the fair redistribution of arrival slots. RBS is the equitable priority scheme for resolving the competition for limited resources. An extranet called CDMnet was deployed to enable information sharing, and a common decision support tool, the Flight Schedule Monitor (FSM), was deployed at the FAA facilities and the AOCs. For further information regarding CDM history see [20] and [21].

CDM-inspired solutions have proven highly successful. According to the FAA, the implementation of CDM solutions in GDPs saved more than four million minutes of scheduled ground delay between September 1998 and December 1999 [15]. For detailed analyses of how CDM has led to more effective GDPs, see [3] and [4].

In practice, CDM functions as a cyclic process of information sharing between the NAS users and the FAA as shown in Figure 1.1. Using information interfaces that are

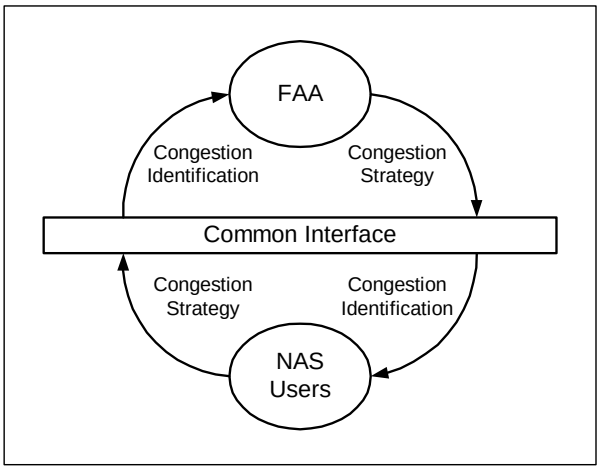

Figure 1.1: Information Sharing in CDM

common to both the FAA and the airlines, the FAA identifies a source of congestion in the NAS. The FAA then forms a strategy for dealing with the congestion (such as a

GDP) and provides the NAS users with information describing the strategy (for example, affected flights). The NAS users, in turn, incorporate the FAA's strategy with the congestion information to make their own operational decisions (such as cancellations). Then, the FAA updates the congestion status, revises the traffic management strategy, shares the updated information with the NAS users, and the cycle continues.

This thesis deals with the area of CDM called Collaborative Routing. It is anticipated that much of the goodwill and achievements used to improve GDPs can be utilized to improve congestion in the en route airspace.

1.2 Collaborative Routing

Just as CDM used near real-time information sharing to improve GDPs, CR hopes to use the same basis as a means for improving en route airspace management. The goals of CR are to improve NAS safety and efficiency and to minimize delays in ways that promote equity and distribute appropriate decision-making to the NAS users. This thesis explores concepts that have arisen in the area of CR for rationing resources during periods of congestion.

CR is a focus that exists within the traffic flow management (TFM) domain of air traffic management. TFM is responsible for balancing demand and capacity in the NAS. Within the scope of TFM, there are three main entities: the Air Traffic Control System Command Center (ATCSCC), the air route traffic control centers (ARTCCs), and the NAS users. The ATCSCC and the ARTCCs are FAA organizations. The ATCSCC is responsible for forming strategies to deal with major NAS congestion, and the ARTCCs are in charge of regional routing problems. Major congestion events, such as large scale

convective weather, require close coordination between the ATCSCC and the appropriate ARTCCs. The airlines are the most prominent of the NAS users, with the resources to best participate in CDM innovations. However, there are other NAS users such as general aviation (GA) that cannot be ignored.

The system shown in Figure 1.2 demonstrates a vision for the future of Collaborative Routing. The image depicts common congestion predictions for weather

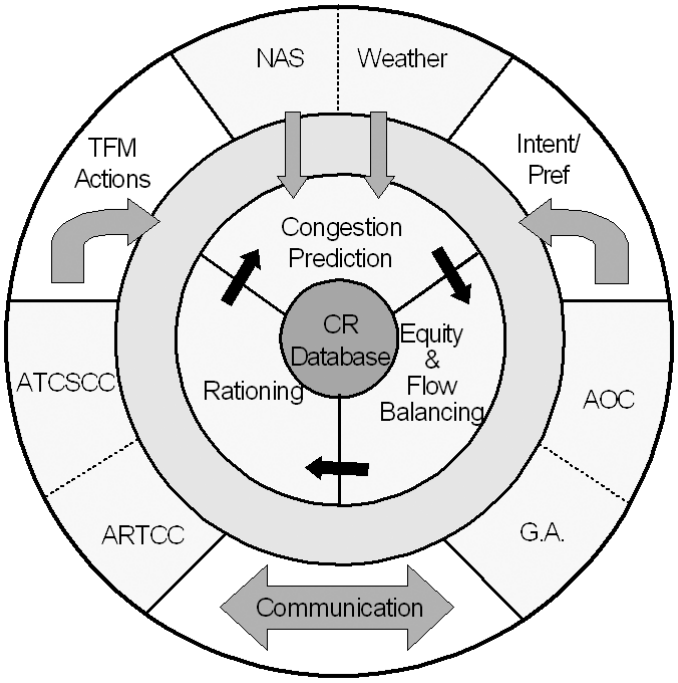

Figure 1.2: A Vision for Collaborative Routing

and NAS status as input to the system. The NAS users (AOCs and GAs) share intent information with the ATCSCC, the ARTCCs, and each other. The FAA entities enact TFM strategies for relaxing NAS congestion. The CR database provides common situational awareness for all participants. The Figure also shows that the system provides a resource rationing function as part of the TFM cycle. It is exactly this CR function that this thesis is concerned with.

1.2.1 Current Collaborative Routing Efforts

There are a number of tools in the development and deployment stages to help further the goals of CR. Already, the FAA has deployed a National Playbook, Coded Departure Routes (CDRs), Low Altitude Arrival and Departure Routes (LAADR), and a Collaborative Convective Forecast Product (CCFP) in support of CR. Summaries of these operational mechanisms appear below [12]:

- The CCFP is a weather forecast product that is generated by a number of collaborative sources. It exists as a common source of weather data for all participants in CR.
- The National Playbook is a document published by the FAA that contains standard routes used to handle common weather scenarios. These routes help to facilitate communication and expedite coordination of rerouting strategies.
- CDRs are a database of alternate standard routes, used during rerouting to aid communication, that the airlines and the FAA can manage using a software package called the Route Management Tool.
- LAADR are low altitude alternate flight procedures that are available during periods of congestion and are used to improve airspace efficiency.

There are also several CR software tools employed by the CR research and development community including the Collaborative Routing and Coordination Tools (CRCT), the Future ATM Concepts Evaluation Tool (FACET), and the Post Operations Evaluation Tool (POET) that are described below. See [12] for more information regarding these achievements in CR.

- POET is useful for viewing and evaluating congestion strategies post facto.
- FACET is a tool that can be used to simulate futuristic congestion strategies.
- CRCT is a prototype designed specifically for testing and implementing CR strategies. CRCT capabilities include identifying aircraft affected by a region of reduced capacity and tactically rerouting or delaying those aircraft [16]. The airspace rationing methods discussed in this thesis purposely form a component of a CRCT-like system.

1.2.2 Long-Term CR Group

Building upon the success of CDM in improving GDPs, a Long-Term CR Group was formed to propose methods for improving congestion management in the en route airspace. This group was comprised of representatives from the FAA, airlines, industry, and academia. Several alternative rationing concepts were proposed in these meetings. The rationing schemes in this thesis are based largely upon the output of these CDM meetings (from [8], [9], and [14]).

1.3 Project Motivation and Objectives

The underlying problem posed in this thesis is to reduce the demand upon an airspace region where demand is predicted to exceed capacity in an efficient manner that promotes equity among the NAS users. This thesis addresses demand reduction by assigning delays to lower the rate of airspace usage, rerouting aircraft around problem airspace, or a delay/rerouting combination. There are possibly other mechanisms such as altering aircraft speed or altitude that are not used in our solution.

As already described, there are currently tools under development, such as CRCT, to aid in tactical rerouting and delaying of aircraft. There is also some prior research in the literature on optimization-based methods for solving en route congestion problems (see [6] and [7]). However, this research employs basic priority rules rather than complex optimization models in order to maintain system transparency. This approach is taken to facilitate alternate priority scheme experimentation by the user community so as to draw out fundamental fairness issues.

The objective of our en route rationing research is to develop methodologies for allocating reroutes and delays based on CDM principles, in such a way that an acceptable balance of fairness and system efficiency can be achieved. We hope that our work can be viewed as a natural extension of RBS and compression in GDPs. As of yet, there is no agreed-upon method for this rationing process. This thesis presents the En Route Resource Allocation Prototype (ERAP) software to help accomplish this goal. The core goals of this thesis are listed below:

- Investigate concepts that surfaced in the deliberations of the Long-Term Collaborative Routing group.
- Refine operational concepts and resource allocation principles through prototype implementation.
- Develop a platform that allows users to view and manipulate possible decision models leading to decisions upon tool requirements.
- Provide a mechanism for strategic comparison of alternate resource allocation methods.

1.4 Organization of Thesis

This thesis continues in Chapter 2 by laying the foundation for a solution to the CR problem already posed. The scenario examined in this thesis is presented and operational concepts are explained. Chapter 3 introduces the global algorithm used in ERAP. It then explains the CR concepts chosen for inclusion in the prototype, and it provides details of the various components that can be used to create user-defined rationing algorithms. Chapter 4 gives an overview of ERAP's capabilities. Chapter 5 presents a set of experiments to exhibit important properties and to demonstrate resource rationing options. This thesis ends in Chapter 6 with conclusions and recommendations for future work.

Chapter 2.

Problem Approach

A software prototype was developed and an operational scenario was selected as a method of investigating the en route airspace rationing problem. This chapter gives the details of both the prototype and the scenario.

2.1 General Collaborative Routing Requirements

A major goal of this thesis work is to evaluate the feasibility of proposed en route resource rationing schemes. ERAP is not designed to be a comprehensive traffic flow management tool. However, in order to generate useful insight, the rationing schemes must exist within the context of a general CR platform. ERAP uses a simplified model for representing the functions that are fundamental to an en route traffic flow management tool. The fundamental functions for a CR platform are derived from CRCT documentation [16] and are shown below:

1. Flow constrained area (FCA) definition.
2. Identification of aircraft affected by an FCA.

3. Definition of alternate routes.
4. Rerouting and/or delaying aircraft.

Currently, ERAP simplifies the first three of the above steps and concentrates upon the fourth. ERAP assumes the first three functions as input to the system and does not provide specific features to support them. The rationing schemes used in ERAP could form a crucial component of CRCT or similar en route TFM tools.

2.2 En Route Resource Allocation Prototype Overview

ERAP is a software prototype designed using the Java programming language. The system uses flight data stored in Microsoft Access databases. The data requirements for these databases appear in Appendix A. ERAP is designed to possess a significant level of flexibility. Resource rationing scenarios are defined by the ERAP input data, and ERAP can handle any scenario that the data can portray.

The data used in this evaluation is drawn from the POET database maintained by Metron Aviation, Inc. POET data is derived from Enhanced Traffic Management System (ETMS) data. The ETMS carries much of the operational information used by the FAA for ATM. A subset of the POET data for July 11, 2001 is used for this analysis.

ERAP is designed to provide insight into the selection of resource rationing methods for en route airspace. In support of this goal, ERAP provides the following core capabilities:

- User-defined resource characteristics.
- User-defined resource allocation schemes.

- Flight track displays before and after resource rationing.
- Statistical and graphical analysis of results.

2.3 Operational Scenario

The scenario chosen for analysis in this thesis was recommended by an ATCSCC specialist [2] and is depicted in Figure 2.1. This scenario shows the New York

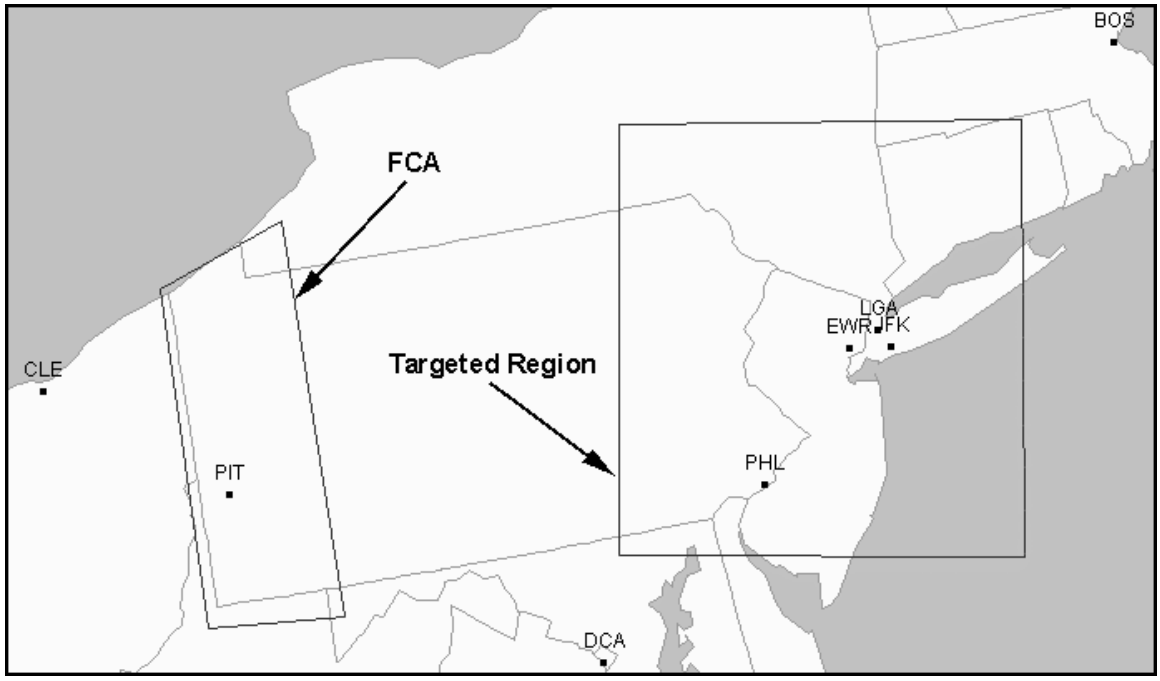

Figure 2.1: Operational Scenario Used in This Analysis

airspace as a targeted region requiring TFM control strategies due to a large FCA to the west. The FCA constrains the New York area's westerly inbound and outbound traffic. An FCA is an ATM concept used to describe a region of airspace with reduced capacity (caused by events such as severe weather). The New York area airspace, as defined by the specialist, is roughly the size of the box shown in Figure 2.1, and it includes major airports such as Newark, LaGuardia, John F Kennedy, and Philadelphia. The ATCSCC

specialist indicated it would be useful if a tool could determine controlled times of departure (CTDs), controlled times of arrival (CTAs), and reroutes for the aircraft affected by the FCA.

2.4 Problem Formulation

ERAP uses priority-based, user-defined rationing schemes to allocate flights to resources. At this point, it is useful to describe operational concepts that are fundamental to the resource rationing procedure.

2.4.1 Target Flights

A CR platform, as described in Section 2.1, must identify the aircraft that will be targeted for TFM initiatives. Given the targeted region and the FCA, the dataset used for ERAP analysis includes all flights with a preferred route that intersects the targeted region and the FCA. Figure 2.2 shows the preferred flight plans for 71 targeted flights expected to fly through the FCA in a one-hour time period. The other flights that pass though this geographic area (called peripheral flights) are not included.

While the omission of peripheral flights does not detract from the core goals of this thesis (to elucidate the resource rationing process, using alternative CR concepts, in the en route domain) future ERAP implementations should incorporate these flights in order to generate more practical resource rationing outcomes. The decision will have to be made, however, as to how the peripheral flights should be handled. For example, these flights could be given priority over the target flights, or they could compete evenly with the target flights for airspace resources.

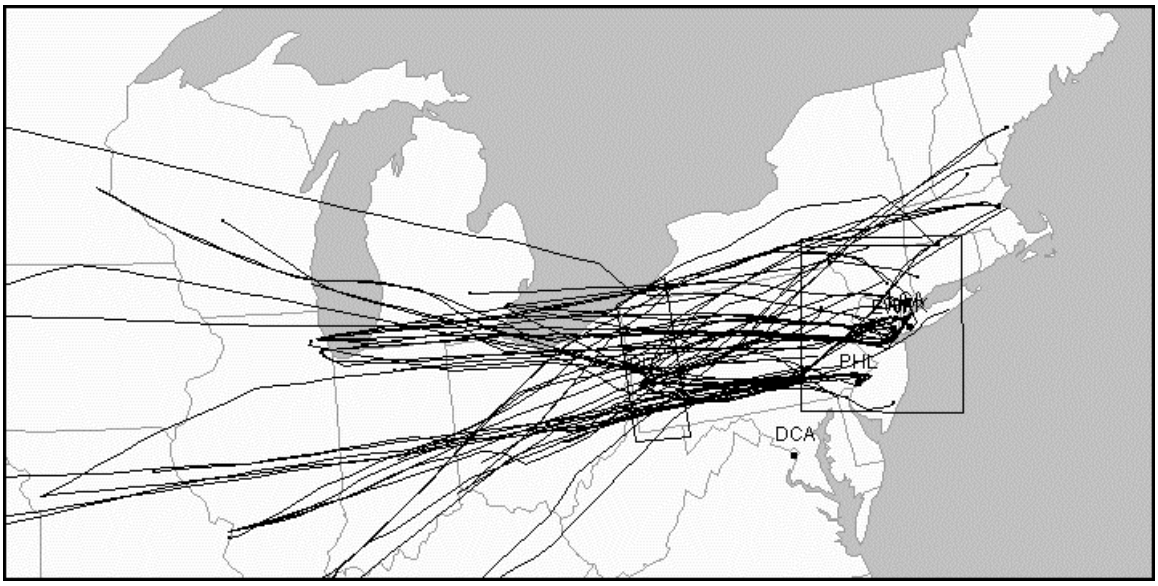

Figure 2.2: Preferred Routes of Targeted Flights

2.4.2 Resources

ERAP is designed to allocate resources. The term "resource" is used very broadly in this context. A resource could be anything that is "consumed" by a flight's en route trajectory (such as an arrival fix or sector). Resource usage by a flight is characterized by an occupancy time. ERAP uses resources as metering points, where a metering point is a location in the NAS through which air traffic flow is regulated. In ERAP, resources are modeled as a series of time slots. Each time slot can be allocated to a single flight, and a resource's capacity is defined by the number and size characteristics of its time slots.

shows the resources used in this analysis. Here, there are five resources represented as geometric planes in the sky. There are two resources to the north and two resources to the south of the FCA. Of these resources, one to the north and one to the south are for east-to-west traffic, and one north and one south are for west-to-east traffic. The FCA itself is also a resource that can handle bi-directional traffic. The

FCA in this analysis was chosen arbitrarily to demonstrate ERAP functionality, and the resources follow as a natural paths around the FCA that an ATCSCC specialist might define.

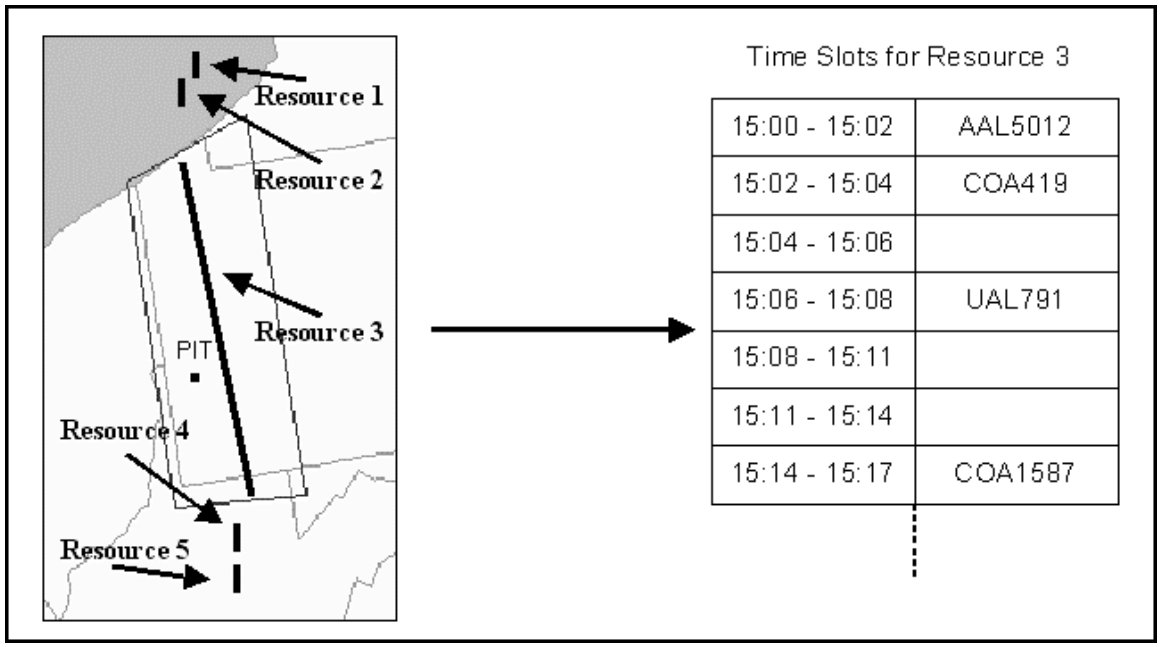

Figure 2.3: Resources and Time Slots

In this analysis, each flight may use one of the five parallel-structured resources. Future work might wish to address interdependent resource networks, as there certainly are other resource allocation scenarios that cannot be modeled using the approach taken in this thesis.

illustrates how ERAP represents a resource's capacity using time slots. In this implementation, flights use a resource for a single instant of time. A common air traffic flow management scenario involves managing sector load (the number of flights allowed in a sector during a period of time). In ERAP, a sector resource could be modeled by making the following alterations to the resource model:

1. Permit multiple flights to occupy a time slot (up to the maximum number of flights allowed to simultaneously occupy a sector).
2. Allow a single flight to be assigned to multiple time slots in order to represent a period of sector occupancy. For example, a flight that occupies a sector for ten minutes could be assigned to five two-minute time slots.

2.4.3 Choosing the Best Route

Another important function in CR is the ability to reroute aircraft. ERAP possesses the capability to move flights from their preferred route to an alternate route. Figure 2.4 this shows an example of two possible routes a flight could use for the scenario discussed in thesis. The process of choosing among multiple routes is further discussed in the next chapter.

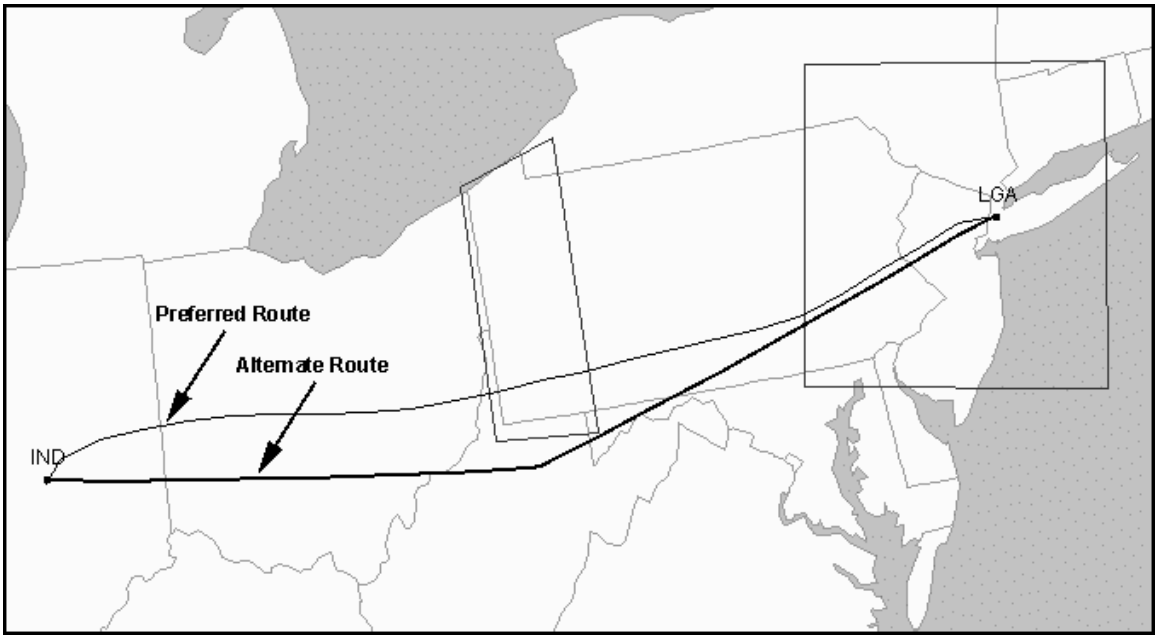

Figure 2.4: Rerouting a Flight to an Alternate Route

2.5 Metrics

In order to assess the equity and efficiency of different rationing schemes, ERAP provides a suite of evaluation tools. These tools can be used to quickly weigh the advantages and disadvantages of different schemes.

Equity is measured by calculating the distribution of delay across various groups of interest. There are six different types of delay:

- Preprogram Air Delay: Air delay accumulated prior to an ERAP allocation.
- Preprogram Ground Delay: Ground delay accumulated prior to an ERAP allocation.
- Assigned Air Delay: Airborne holding assigned by an ERAP allocation.
- Assigned Ground Delay: Ground delay assigned by an ERAP allocation.
- Longer Route Air Delay: Air delay incurred by flying (in accordance with ERAP) a route that is longer than the originally planned route.
- Total Delay: The sum of all types of delay.

Statistics (such as minimum, maximum, average, and standard deviation) can be associated with each type of delay. ERAP further permits the analysis of equity by parsing delay statistics into categories such as airlines, airports, and classes of aircraft traffic. Another useful metric calculated by ERAP is the average delay of a fraction of the most delayed flights. This calculation is designed to characterize the plight of a resource allocation's most penalized flights.

The efficiency of a resource allocation describes how optimally resources are used. Given resources with defined time slot characteristics, the efficiency of different

rationing schemes can be evaluated based upon resource utilization, average total delay, and by the percentage of aircraft that are assigned to alternate routes.

2.6 Major Assumptions

As already mentioned, ERAP makes several simplifications in order to focus on its rationing scheme evaluation goals. Some assumptions were made due to insufficient airline data, and others are made to bring this project within the scope of a master's thesis. Overall, it is felt that these assumptions do not damage the goals of ERAP. The major assumptions appear below:

- In the scenario examined in this thesis, all background traffic (such as flights through the FCA but not through the targeted region) is disregarded. It is assumed that the background traffic could be handled separately or will be addressed in future ERAP revisions.
- ERAP is viewed as a strategic tool, and it lacks micro-level tactical abilities. As such, ERAP solutions do not necessarily account for the timing considerations that would be used in a tactical domain to ensure NAS safety (for example, ERAP solutions do not attempt to prevent flight collisions).
- Historical flight radar data is used in this analysis to represent preferred flight paths.
- In this analysis, alternate flight paths are created using a three point trajectory (airport-resource-airport), and the alternate route flight times are estimated using the average flight time of the preferred route.

- Delays assigned by ERAP are based solely on en route resource allocations. ERAP disregards any need to coordinate feasible airport departure and arrival times. For example, it is theoretically possible for ERAP to assign the same CTA to two flights destined for the same airport.
- ERAP resource allocations are completely deterministic. It does not account for the stochastic processes that are inherent to the en route resource rationing domain (such as weather and aircraft arrival and departure times).

Chapter 3.

Resource Allocation

ERAP is a platform for demonstrating and analyzing new en route resource rationing concepts. The ERAP design and implementation represent a significant contribution of this thesis, as ERAP offers a baseline for future agreement upon standardized en route rationing procedures. This chapter provides details regarding the implementation and integration of the resource rationing concepts that exist in ERAP.

3.1 Overall Resource Allocation Process

ERAP is capable of assigning delays and rerouting aircraft in order to ease congestion upon resources where demand exceeds capacity. ERAP assigns flights to resources via user-defined priority functions. The global resource allocation algorithm, shown in Figure 3.1, operates as a priority-based assignment loop. Given a list of flights to assign and descriptions of available air space resources, ERAP first determines the resource that each flight would use if selected next for assignment. ERAP then determines the flight that has the highest priority for its selected resource and assigns the flight to that resource. ERAP allocates one flight to one resource's time slot per iteration through the

loop until all flights have been assigned. Each flight allocation alters the makeup of the resources for the remaining flights.

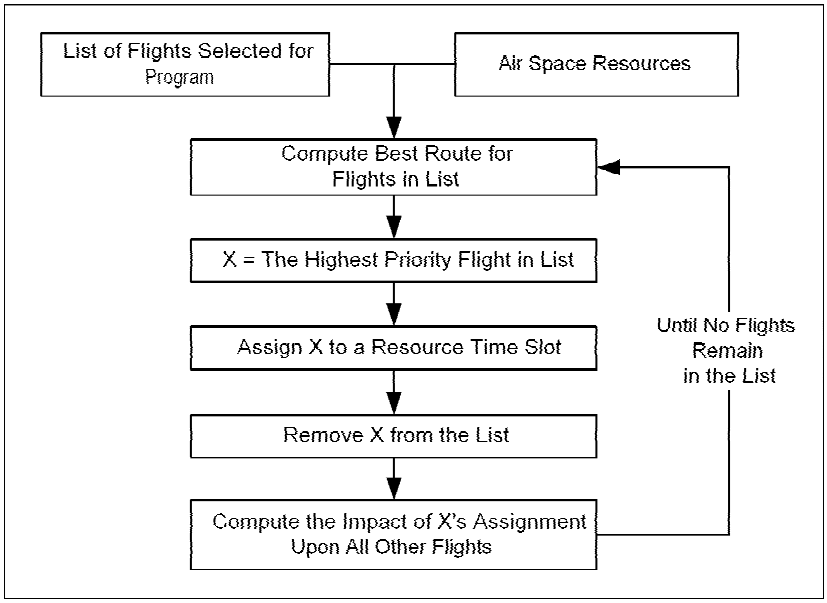

Figure 3.1: The Global Resource Allocation Algorithm

Since one might naturally seek optimization-based solutions, the priority-driven assignment loop technique used in this research requires justification. One reason for the greedy algorithm approach to resource allocation, as already mentioned, is to allow for a high degree of transparency, so that broad en route resource allocation principles can be quickly and easily derived. Another argument for our approach stems from the RBS algorithm that is currently viewed as the highly lauded CDM solution to improving the effectiveness of GDPs. RBS is a priority-based solution, and it is only natural to begin the search for the RBS equivalent within the en route domain by emulating RBS itself.

The remainder of this chapter is arranged to explain the various components and operational concepts that are necessary to fully define the ERAP resource allocation process. Section 3.2 describes how ERAP handles the option for flights to file multiple

routes. Then, Section 3.3 explains another important concept in ERAP known as traffic classes. In Section 3.4, the actual priority functions are described in detail.

3.2 Alternate Routes

Continental U.S. flights currently file a single flight plan before takeoff. However, severe weather events can intersect the original flight plans and lead to a need for FAA-imposed reroutes. ERAP allows flights to file multiple routes. Flights can specify, a priori, alternate routes accompanied by rules to indicate their willingness to switch. This represents a new operations paradigm, described in [8], that could have major benefits for the airline industry. Some of these benefits are listed below:

- Keeping with the goals of Free Flight, airlines can exert more control over their own business operations. They are encouraged to incorporate their own economic models for balancing delay against fuel consumption into their flight plans.
- From the point of view of the FAA, alternate route filing adds "free" intelligence to easing congestion. Ideally, the airlines could anticipate congestion and use alternate routes to alleviate the problem without the need for FAA intervention. Also, alternate route filing adds structure to the process of selecting which flights to reroute when TFM initiatives are enacted.
- The capacity to file multiple routes lessens the impact of imperfect weather or congestion data. Flights can cope with congestion uncertainty by filing routes that react to a number of potential congestion outcomes.

A foundation for filing multiple routes is currently exemplified in the Pacific Track Advisory Program [19]. Currently, airlines indicate preferences for flying predefined routes over the Pacific Ocean on an individual flight basis. The Oakland ARTCC then attempts to honor airline requests when awarding route assignments.

ERAP approaches multiple route filing as shown in Figure 3.2. A flight files a preferred route and a number of alternate routes. Each route requires access to a resource that is dependent upon the route's trajectory. In ERAP, each alternate route is accompanied by a delay threshold. The delay threshold is a rule for the amount of total delay the flight would need to save in order to switch to an alternate route. For example, in Figure 3.2, the flight might be willing to fly the longer, more fuel consuming, Alternate Route 1 instead of the Preferred Route if it would save 30 minutes in reaching its destination. The delay threshold procedure used in ERAP is a single implementation, and it is not the only possible solution to handling alternate route preferences. Another potential solution might be to use flight-specified air delay to ground delay tradeoffs.

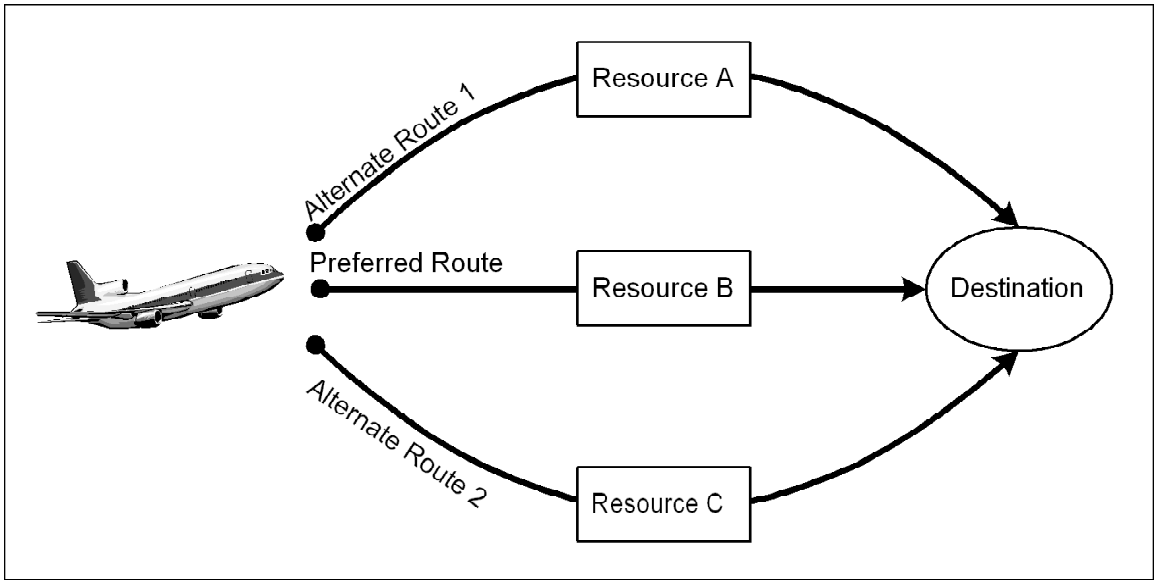

Figure 3.2: Alternate Routes for a Hypothetical Flight in ERAP

Through every iteration of the resource allocation loop, ERAP determines each flight's best route based upon the status of the resources. ERAP does this by calculating the amount of destination delay that would be incurred for each of a flight's filed routes separately and then using the delay thresholds to select the flight's specified best route. A flight's best route corresponds to the route with the minimum summation of delay plus the threshold (the preferred route's threshold is assumed to be zero). Through every iteration of the allocation loop, flights compete for the resource used by their best route. The best route algorithm is described in pseudocode below:

```
Best Route ← Preferred Route
Best Route Delay ← Preferred Route Delay
For Each Alternate Route i {
    If (Alternate Route_i Delay + Threshold_i < Best Route Delay){
        Best Route ← Alternate Route_i
        Best Route Delay ← Alternate Route_i Delay + Threshold_i
    }
}
```

Refer to Table 3.1 as a brief example of selecting the best route for a flight. The delay thresholds imply that the flight prefers Route B to Route A because it saves over 30 minutes of total delay.

Table 3.1: Example of Using Delay Thresholds

Route	Air Delay	Ground Delay	Threshold	Total
A	0 min.	60 min.	0 min.	60 min.
B	15 min.	5 min.	30 min.	50 min.

3.3 Traffic Classes

The traffic class concept was created to provide a mechanism for measuring resource allocation equity and for guiding algorithms into generating equitable solutions. The Long-Term CR subgroup describes traffic classes as an approach to handling aggregate equity by separating air traffic into meaningful, user-defined sets of flights. Then, these sets are treated differently when competing for the use of en route resources [9].

A good argument for traffic classes occurs regularly at Chicago O'Hare airport. Convective weather can often cause the FAA to reroute air traffic through a dense channel over O'Hare airport [11][17]. In order to maximize throughput in this scenario, traffic controllers regularly limit the amount of flights departing from O'Hare. This practice occurs because the departing flights require extra resource usage time to merge into the overhead traffic stream.

While the method of limiting O'Hare departures maximizes resource throughput, the practice is viewed by some airlines as unfair to the airports that are repeatedly selected for this TFM initiative [17]. Traffic classes could be used to trade off a certain amount of system efficiency in exchange for greater equity in the treatment of the flights departing O'Hare relative to the overhead traffic stream.

In ERAP, the traffic class procedure is based upon setting resource utilization goals for each traffic class. For example, a goal of 10% of flights per hour could be set for flights departing O'Hare, and a goal of 90% of flights per hour for the overhead stream. Then, the resource allocation mechanism can account for these goals when distributing access to the congested resources.

The traffic class concept is not entirely revolutionary, as a partial basis for traffic classes already exits within the traffic flow management community. Air traffic controllers often mentally group air traffic into separate flows such as Chicago arrivals from the east [13]. Examples of how traffic controllers group traffic flows are also evident in the FAA's National Severe Weather Playbook routes [10]. Several playbook scenarios are designed to manage the rerouting of specific categories of flights in the event of severe weather. For example, a playbook route for eastern arrivals to Chicago that are affected by an area of unusable airspace (the shaded region) is shown in Figure 3.3. We anticipate that traffic classes represent an air traffic management operations paradigm that will be natural for air traffic specialists to understand and interact with.

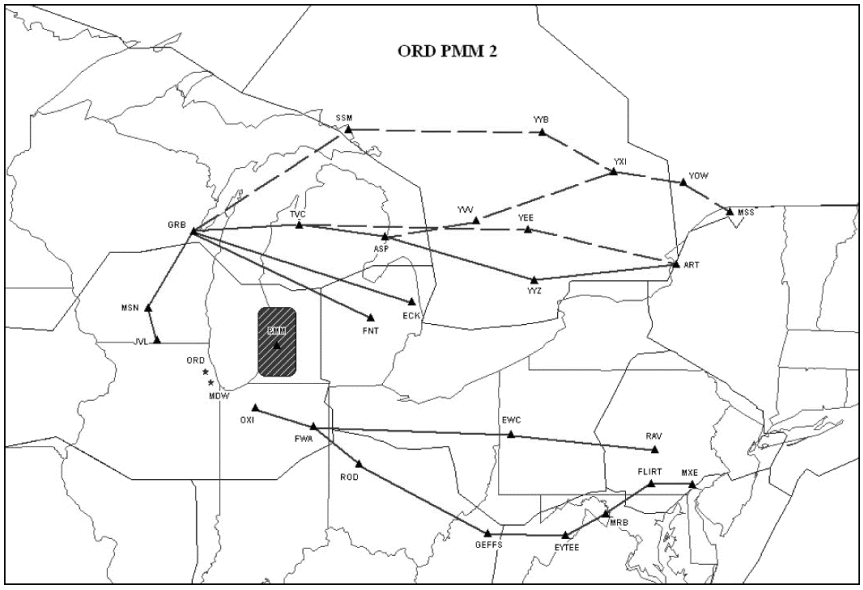

Figure 3.3: Playbook Route For Airspace Near Chicago

Traffic classes are incorporated into the overall rationing workflow of ERAP using the following three-step process: traffic class definition, resource goal definition,

and goal deviation as a rationing function. The traffic class definition and resource goal definition steps occur before resource allocation. The ensuing sections explain these steps in further detail.

3.3.1 Traffic Class Definition

In the ERAP traffic class definition process, a user combines flight criteria to define traffic class sets. Table 3.2 lists the flight criteria chosen for implementation in ERAP.

Table 3.2: Traffic Class Criteria in ERAP

Flight Criteria	Values
Operation	Region Departure, Region Arrival, Region Departure/Arrival, Region Overflights, Any Airport Combination
Departure Airport	Airport code
Arrival Airport	Airport code
Physical Class	Jet, Piston, Turbo
User Class	Air Taxi, Cargo, Commercial, GA, Military, Other
Weight Class	Small, Medium, Heavy

These criteria were selected to demonstrate the practicality of traffic classes and are by no means assumed to be exhaustive. They represent an initial implementation of meaningful flight characteristics that can be used to distinguish among flights, and they were chosen so that no criteria could be dually interpreted as an en route airspace resource.

The Operation flight criterion is relative to a particular region. Recall that the targeted region for this analysis is the New York area. Thus, in this analysis, a Region Overflight is a flight that flies through New York airspace but does not originate from or terminate at an airport in the New York area. The other five flight criteria are taken directly from the POET database and are defined by their possible values.

The user can employ the above flight criteria from Table 3.2 to define sets of traffic classes. Flights can exist in multiple traffic classes, and ERAP uses colors to represent these traffic class sets in order to enhance visibility. ERAP automatically assigns flights that lack a user-defined traffic class to an "Other" traffic class. As an example, the user could define the following five traffic classes:

- Red: Newark, JFK, and LaGuardia Airport Arrivals
- Blue: Newark, JFK, and LaGuardia Airport Departures
- Orange: Philadelphia Airport Arrivals
- Green: Philadelphia Airport Departures
- Cyan: All Region Overflights

The output of this traffic class assignment is described visually in Figure 3.4 and in Figure 3.5, which show a two-hour period of air traffic before and after traffic class aggregation.

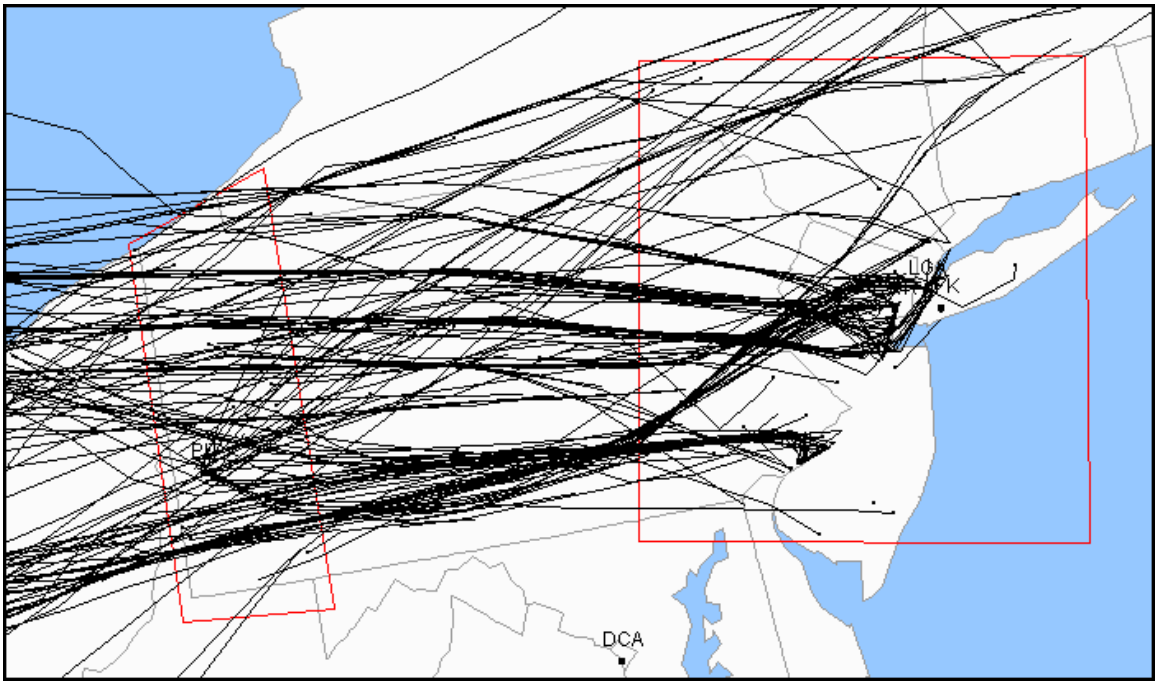

Figure 3.4: Flight Tracks for Two-Hour Period Before Traffic Class Aggregation

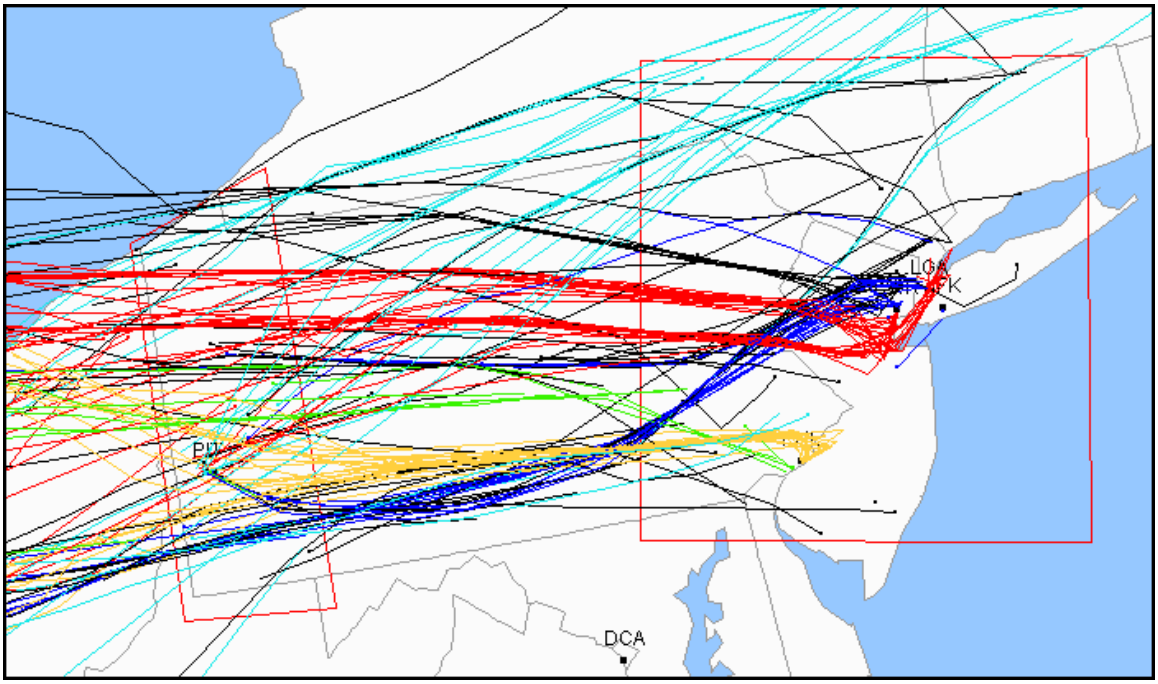

Figure 3.5: Flight Tracks for Two-Hour Period After Traffic Class Aggregation

3.3.2 Resource Goal Definition

The next step in the ERAP traffic class process is for the user to set resource-specific goals for traffic class flow rates. Goals represent the percentage of flights of particular traffic classes that should get assigned to a resource per hour time period. Thus, a traffic class goal is specific to a single resource. ERAP uses goals for resource usage instead of specifying strict ownership in order to maintain resource usage efficiency. Ideally, these resource usage goals should be based upon formal equity principles. The derivation of formal goal definition models is another research problem.

The "Other" traffic class used in ERAP is a special case. Every resource that is made available for rationing has a default "Other" traffic class goal of 0% so that it does not conflict with user-defined traffic class goals. The user can, however, change the "Other" goal value from its initial value of 0%. A flight is assigned as an "Other" traffic class if it is not a member of any traffic class that matches a goal definition for resource it is assigned to.

As an example of traffic class goal definition, refer to the traffic classes created in the previous section and shown in Figure 3.5, and assume that the FCA is a resource with some limited capacity. A user could decide (for some arbitrary reason) that the New York Overflight traffic class (colored light blue) has minimal priority for using the FCA resource. The user could enforce this principle by setting goals at the FCA resource for 0% Overflights and 100% Other flights. This goal structure gives priority for the FCA resource to all flights in a non-Overflight traffic class. However, because ERAP treats traffic classes as a goal-based solution, a member of the Overflight traffic class could still

be assigned to the FCA resource if no non-Overflight traffic class member is capable of using a particular time slot.

3.3.3 Goal Deviation

The final step in ERAP's traffic class implementation is to incorporate the user's traffic class goals into the resource allocation. As ERAP assigns flights to resources, it tallies the number of flights from each traffic class that have been assigned to each resource per hour time period. This bookkeeping methodology allows ERAP to calculate the deviation of traffic classes from their resource usage goals. If a flight belongs to multiple traffic classes, ERAP assigns the flight as its traffic class with the highest amount of goal deviation. The deviation function represents the number of flights of a particular traffic class that are absent from resource's goal for a time period. The function is shown below for resource *r*, traffic class *tc*, and time period *t*:

Goal Deviation (*r*, *tc*, *t*) =

[Total Assigned (*r*, *t*) × Goal (*r*, *tc*, *t*)] - Total Assigned (*r*, *tc*, *t*)

where Total Assigned (*r*, *t*) represents the sum of all aircraft assigned to a resource within a time period, Goal (*r*, *tc*, *t*) is a user-defined traffic class goal for a resource within a time period, and Total Assigned (*r*, *tc*, *t*) is the number of flights of a specific traffic class that have been assigned to a resource within a time period. The goal deviation function returns a negative value if a traffic class goal is exceeded, and it is positive when more flights from a traffic class are required to satisfy a resource goal.

In order to respect the fact that resources can have vastly different characteristics, (such as the number of time slots per hour) ERAP multiplies the goal deviation by a

normalization factor that is resource and hourly time period specific. These normalization factors are constant throughout the resource allocation process. For a given resource time period, the normalization factor is equal to the ratio of the capacity of the most abundant resource in that time period to the resource's time slot capacity in the same period. The deviation function is shown below for resource *r* and time period *t*:

Norm (*r*, *t*) = Maximum Total Time Slots (*t*) / Total Time Slots (*r*, *t*)

where Maximum Total Time Slots (*t*) is the highest number of time slots in a time period among all resources and Total Time Slots (*r*, *t*) is the number of time slots in a particular resource within a time period. The normalization factor used in ERAP implies that a single time slot in a resource with 10 time slots per hour is equivalent in value (with respect to goal deviations) to 10 time slots in a resource with 100 time slots per hour.

Refer to Table 3.3 for an example of a traffic class deviation calculation. Before normalization, the deviations at both resources for the Other, Green, and Blue traffic

Table 3.3: Example of Calculating Traffic Class Deviation for a Time Period

		Traffic Class Goals			**Traffic Class Assignments**		
Resource	**# Slots**	**Other**	**Green**	**Blue**	**Other**	**Green**	**Blue**
A	100	0%	50%	50%	0	3	1
B	10	0%	50%	50%	0	2	0

classes are 0, -1, and 1 respectively. For example, in Resource A, four slots have been assigned. A goal for Blue of 50% means that of the four assignments, two should be Blue. The resulting deviation is 1 because Blue is one short of its traffic class goal. Before normalization, the Blue traffic class deviations (valued at 1 for both resources) tie for the highest value. The normalization process breaks the tie by increasing the deviation of Blue at Resource B by a factor of 10 in order to compensate for the larger number of slots in Resource A. The final values returned by the normalized traffic class

deviation function for the Other, Green, and Blue traffic classes are, respectively, 0, -1, and 1 for Resource A and 0, -10, and 10 for Resource B.

This section has explained the ERAP implementation of traffic classes. ERAP's approach to traffic classes culminates in the calculation of deviation from user-defined goals for handling aggregate flows of air traffic. The next section explains how this goal deviation function is used during resource rationing.

3.4 Resource Allocation Priority Functions

ERAP assigns flights to resources via user-defined priority functions. The output of a single priority function is the flight or set of flights that have the highest priority. Figure 3.6 shows that flight priority in ERAP is based on a multi-leveled, hierarchical approach. Priority functions are ordered, with higher level functions feeding into lower level functions. A combination of individual priority functions works to filter a flight

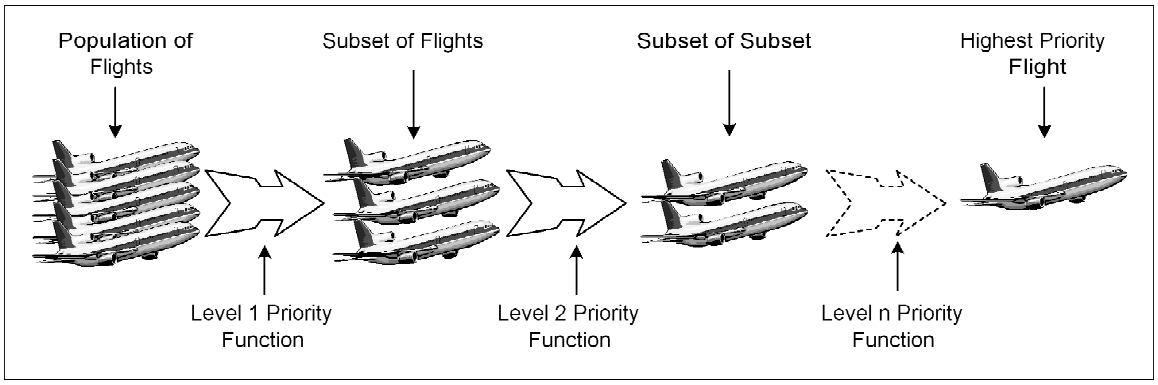

Figure 3.6: Priority Function Usage In ERAP

population down to the highest priority flight. This leveled approach allows priority functions (with vastly different objectives) to work together to solve a resource allocation within the greedy algorithm context of ERAP. If a tie exists after the last priority

function, ERAP assumes equality and picks a flight for resource assignment based upon no formal selection process.

The priority functions available in ERAP can be divided into two categories: high level functions and low level functions. Most functions may be accompanied by a user-specified tolerance value in order to enlarge the set of returned flights. These functions are explained in the ensuing sections.

3.4.1 High Level Priority Functions

The high level priority functions are designed to return a relatively large set of flights with shared characteristics. These functions should, generally, be used before low level functions in a resource allocation priority scheme.

3.4.1.1 Airborne Take-Off Status

In the interest of safety, many of today's TFM control strategies give airborne flights priority for airspace usage over flights on the ground. The airborne take-off status priority function gives the user the option of replicating this condition in ERAP.

3.4.1.2 Highest Traffic Class Deviation

This priority function uses the traffic class deviation calculation (from Section 3.3.3) to return flights belonging to traffic classes with the highest priority for their desired resource. The user can specify a percentage range from the highest deviation to return a greater number of flights.

3.4.1.3 Earliest Demanded Time Slot

The earliest demanded time slot function searches through all unassigned flights' demanded time slots, and it returns the flights that share demand for the earliest time slot. This function can be used to enforce a policy of assigning time slots in increasing time order. A user can indicate a tolerance (in minutes) from the earliest time slot in order to expand the output from this function. This tolerance value can be especially useful if the user wishes to allow competition among multiple resources that have staggered time slot starting times.

3.4.2 Low Level Priority Functions

Low level priority functions operate in the same manner as high level functions. The only difference lies in how these functions are used. The design intention for low level priority functions is to return a small set of flights or to converge upon an individual flight.

3.4.2.1 Random Flight Selection

This function returns a single, randomly selected flight. It can be used to resolve priority ties among flights or as a baseline for experiments.

3.4.2.2 Earliest Expected Meter Time

The ERAP database maintains an expected resource meter time for every filed route, where a meter time is the discrete time of resource usage. An expected meter time accounts for all delays that a flight may have incurred. As the name implies, this function

returns the flight (or flights) with the earliest expected meter time. The user may specify a tolerance from the earliest expected meter time (in minutes) when using this function. This priority function emulates the Grover-Jack resource allocation algorithm that was previously used for pre-CDM resource allocations in GDPs. In these implementations of GDPs, Grover-Jack determined resource priority based upon expected airport arrival time.

3.4.2.3 Earliest Scheduled Meter Time

This function operates in similar fashion to the earliest expected meter time, except that the meter times are based upon flights' scheduled resource usage times. A scheduled meter time is the time at which a flight's filed route intersects with a resource (ignoring any delays). A user may specify a range of minutes from the earliest scheduled meter time as an input parameter to the function. The earliest scheduled meter time priority function is modeled after the ration-by-schedule algorithm that is currently used in GDPs. In GDPs, RBS rations limited airport arrival slots in the order they are scheduled for use (as published in the official airline guide).

3.4.2.4 Most Accrued Delay

The most accrued delay function gives priority to the flight or flights with the largest amount of total accrued delay. Accrued delay includes all forms of delay (such as ground delay, mechanical delay, and rerouting delay) that detract from a flight's original, scheduled time of arrival at its destination. Once again, a user may indicate a range (in minutes) from the maximum in order to return more flights. This priority function can be

used to emulate the Equity via Net Arrival Delay (ENAD) resource allocation concept described in [9]. The intent is to recognize that arrival delay is the measure of success that matters most to all parties (especially the passengers), regardless of the causes for delay.

3.4.2.5 Multi-Objective Priority Function

The final priority function available in ERAP is a multi-objective function that combines the highest traffic class deviation with the most accrued delay function. The user can weight each of the individual functions in order to achieve the desired balance between the two. It is an attempt to strike a balance between equity concerns involving categories of flights and individual flights.

3.5 Resource Allocation Example

This section provides an example that illustrates the approach ERAP takes to determining the highest priority flight for resource allocation. The example described in Figures 3.7, 3.8, 3.9, and 3.10 begins in the middle of a resource allocation for five flights (A, B, C, D, and E) and two resources (Y and Z). The Figures show the time slots that each flight would use if selected next for assignment as well as the amount of accrued delay that would follow from such assignments. Assume that the user has defined the following priority function hierarchy:

- Level 1: Earliest Demanded Time Slot (Range 0 Minutes)
- Level 2: Most Accrued Delay (Range 0 Minutes)

Figure 3.7 shows all five flights selecting the time slot that each would use if selected next for assignment (the crosses indicate occupied time slots). Due to the defined priority structure, flights A, B, C, and D are filtered by the earliest demanded time slot function (because they share demand for the 1:05 time slot). Then, flight A is awarded its chosen time slot because it has the highest amount of accrued delay (20 minutes).

Figure 3.8 shows the same resources after the assignment of flight A. Note that the assignment of flight A has affected the best time slots for the remaining flights, and flight B has even switched resources. This can happen when a flight's delay threshold indicates a switch in the preferred route. In choosing the highest priority flight in this step, ERAP first filters this flight set down to flights B and C because of the highest demanded slot time function. Then, flight B is awarded the 1:05 time slot in Resource Z because it has the highest amount of accrued delay in the filtered flight set. After the assignment of flight B, the example continues in Figures 3.9 and 3.10, which depict the competition among flights C, D, and E for the remaining two time slots. Figure 3.10 shows the result of assigning flight C, and the Figure indicates that flight D would receive the 1:20 time slot in Resource Y if the example were to continue.

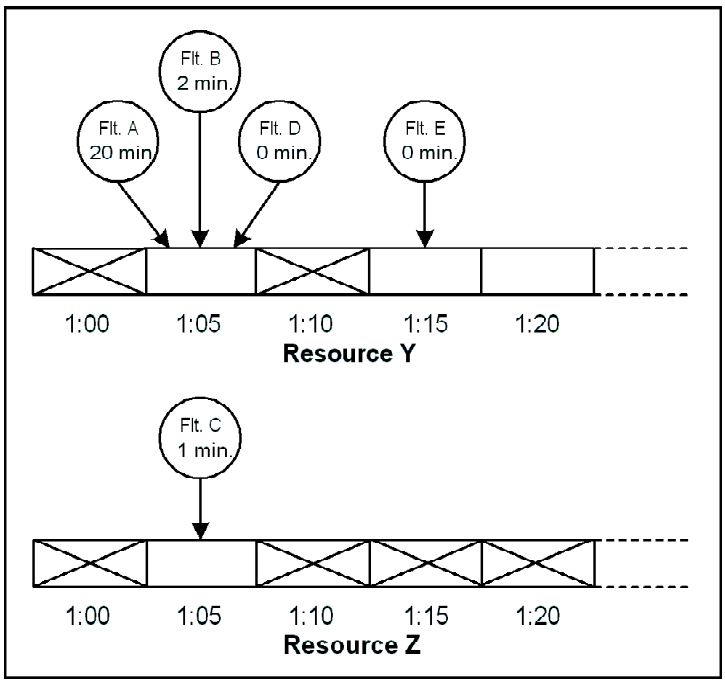

Figure 3.7: Resource Allocation Example Before Assignments

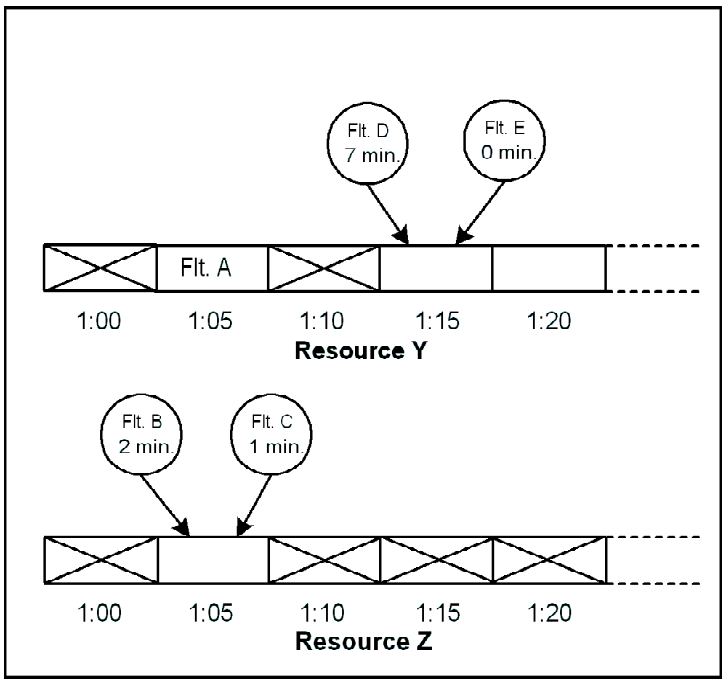

Figure 3.8: Resource Allocation Example After Assignment of Flight A

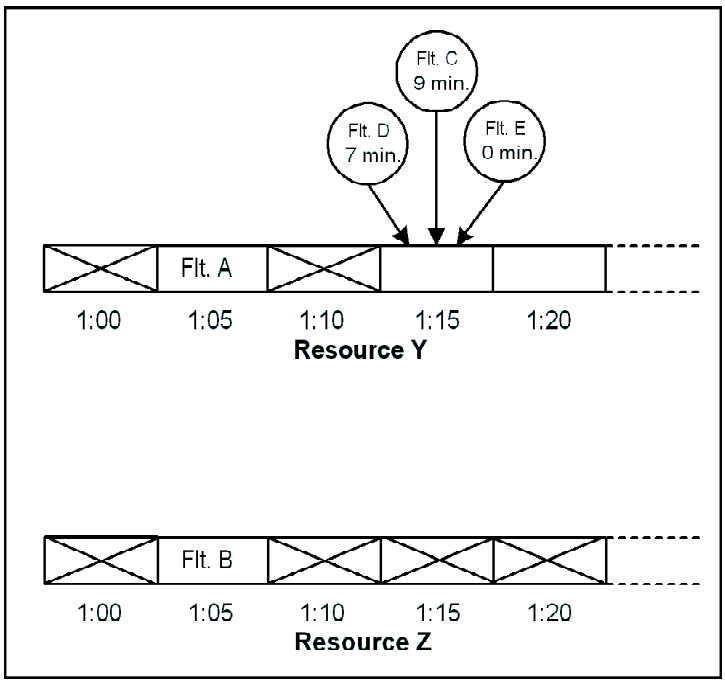

Figure 3.9: Resource Allocation Example After Assignment of Flight B

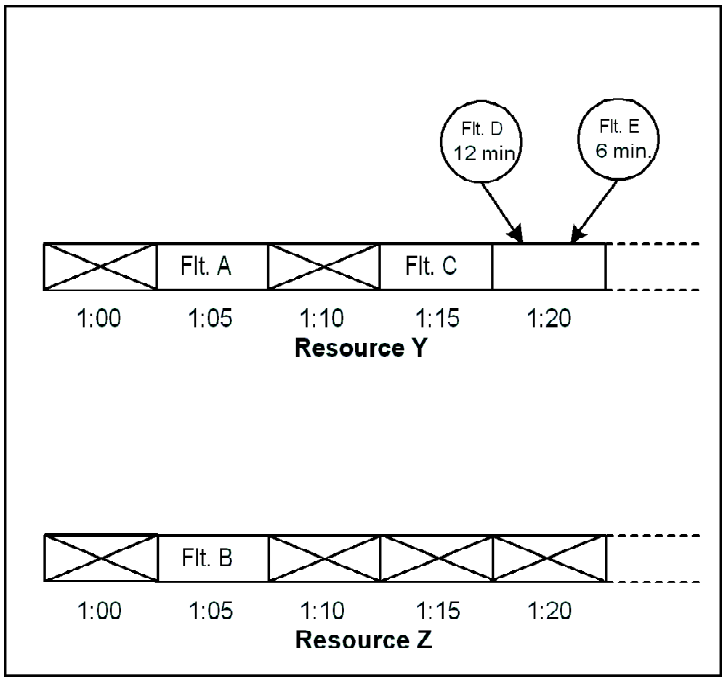

Figure 3.10: Resource Allocation Example After Assignment of Flight C

3.6 Comparing RBS to Accrued Delay

ERAP implements the en route parallel to RBS by using the earliest scheduled meter time priority function. When assigning flights to a single resource, such as in a GDP, it has been proven that RBS lexicographically minimizes the maximum resource delay from schedule [5]. This means that no flight can be given a position of lesser delay without increasing the delay of a second flight to an amount that is equal to or higher than the original delay of the first flight.

Due to the use of time values as the means of comparison, RBS has inherent time ordering information. The most accrued delay priority function lacks this information, as it uses numerical values for comparison. However, time intelligence can be imposed by adding the earliest demanded time slot function to form the time-ordered accrued delay (TOAD) function. TOAD can be viewed as nearly equivalent to RBS when rationing a single resource (a slight difference is shown in Section 5.1.2). Both functions credit flights for delays from schedule, and for a single resource and proper tie-breaking rules, both algorithms lexicographically minimize the maximum delay [5]. There are, however, some substantial differences between the two functions that point to TOAD as being a better solution within the en route domain.

The ERAP implementation of RBS bases priority upon delay to a resource. TOAD gives priority that is based upon destination airport delay. One outcome of these differences, illustrated in Figure 3.11, is that TOAD increases the priority for flights that fly a longer alternate route (Alternate Route 2) because the extra flight time adds to the accrued delay. The RBS priority scheme does not recognize priority for flying a longer

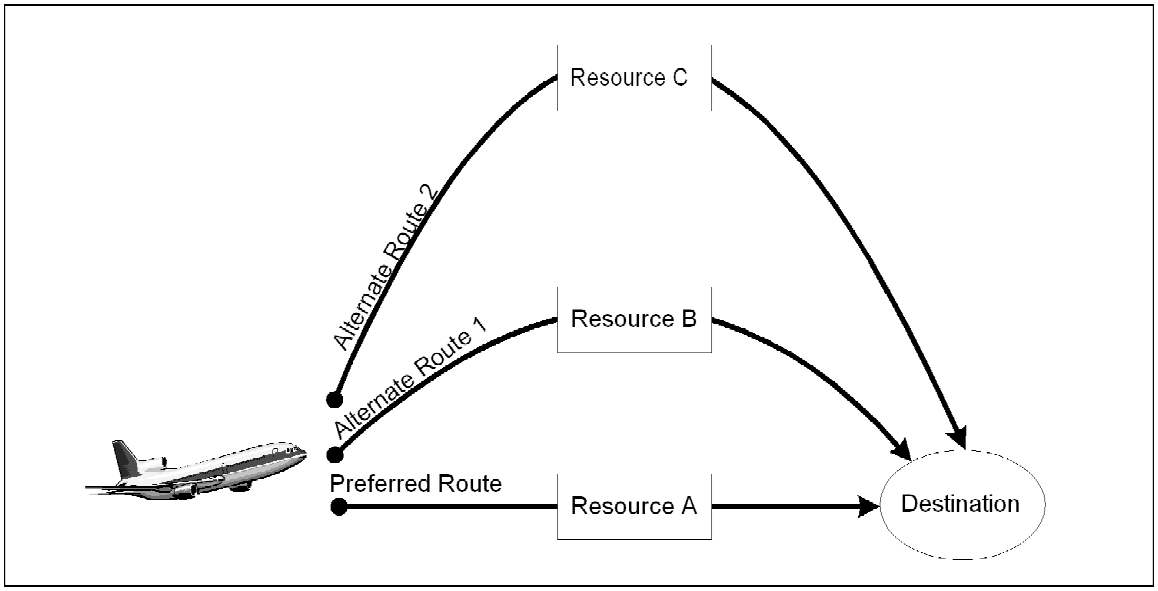

Figure 3.11: Resource Priority Differences Between TOAD and RBS

route and, in fact, would determine that the flight in the Figure would have the highest priority (relative to itself) for its Preferred Route. In general, when rationing multiple en route resources, the ERAP implementation of RBS bases priority for a resource upon geographic proximity to that resource. This priority principle contrasts with the goal of CR to encourage flights to voluntarily file routes around congested airspace. The TOAD algorithm, in effect, rewards flights that file a longer route by increasing their priority for the resource used by the longer route.

A further distinction between the two priority functions lies in the greater robustness and implementability of TOAD. The scenario previously discussed in this thesis addresses resources as geometric planes in the sky, however, consider the scenario depicted in Figure 3.12. Here, three flights are competing for a sector resource. It is not intuitively clear how an RBS implementation would define the all-important scheduled time of resource usage. For example, it could be the sector edge, some point in the middle, or an arbitrary plane through the sector. None of these solutions seem more

plausible than using accrued delay. TOAD would base priority upon the amount of destination delay a flight would incur as a result of flying through the sector. Accrued delay can be implemented in a number of equivalent ways, such as calculating the difference between expected and scheduled destination arrival time or associating a "counter" with every filed route that increases any time delay is incurred.

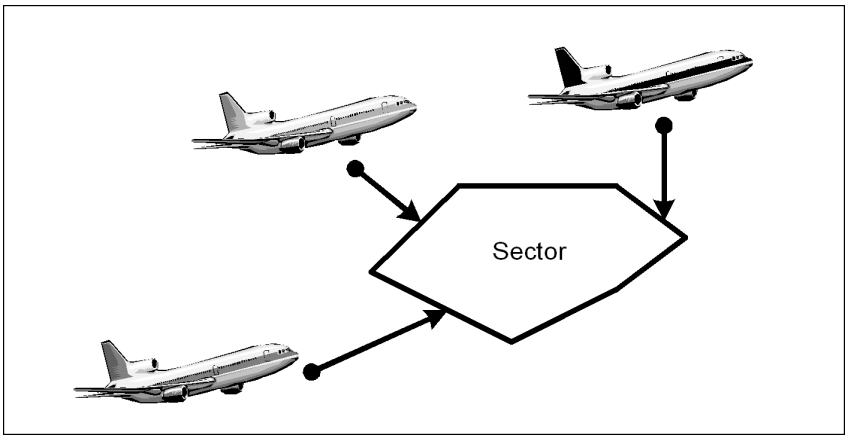

Figure 3.12: Three Flights Competing for a Sector Resource

Chapter 4.

Using the En Route Resource Allocation Prototype

Chapter 2 provided a brief description of the ERAP system architecture. The purpose of this chapter is to introduce the tools available in ERAP and describe how they are put into practice. ERAP is a software platform that can be used to design and evaluate en route rationing schemes. It exists as a suite of individual functions. These functions are separated into the following three categories: database initialization, resource allocation, and results analysis. The remainder of this chapter discusses the functions of each category and how they relate to ERAP capabilities. All screenshots of ERAP graphical user interfaces (GUIs) appear in Appendix B.

4.1 Database Initialization

ERAP maintains a flight operations database (described in Appendix A) with information that is necessary for en route resource allocation. The functions explained in this section allow an ERAP user to manipulate the ERAP database.

4.1.1 Simulate Ground Delay Program

ERAP provides a simple mechanism for simulating a GDP in order to establish preprogram ground delay for a group of flights. The GUI for this function is shown in Figure B.1 of Appendix B. The user inputs a GDP airport, time range for a GDP, original airport arrival acceptance rate, and controlled airport arrival acceptance rate. ERAP uses the input parameters to assign preprogram ground delay to flights. This models the delay flights would get in an actual GDP by FSM software, and it can be useful in evaluating how different rationing schemes treat flights with pre-allocation delay.

4.1.2 Set Alternate Route Delay Thresholds

This feature allows a user to quickly instantiate the airline delay thresholds (in the ERAP database) based upon an airline-generalized economic parameter. Recall that these thresholds represent a flight's preference for selecting one of a multiple number of filed routes. The user can specify a factor which represents an air delay to ground delay tradeoff for a particular airline. This factor is multiplied by each alternate route's extra air time to yield the delay thresholds in the database. The GUI for this function is shown in Figure B.2. In practice, each carrier would likely set a delay threshold for each of its flights individually. This can be done in ERAP by modifying the database manually.

4.1.3 Define Traffic Classes

The traffic class definition GUI, shown in Figure B.3, allows a user to assign flights to traffic classes, as described in Chapter 3. The user can create rules that query the database and add flights to traffic class sets (represented as colors in ERAP).

4.2 Resource Allocation

The resource allocation function is the heart of ERAP. The GUI for this function is shown in Figure B.4. In resource allocation, the user can specify the timing parameters to support a resource allocation program. The user can also define the available resources and the traffic class goals for those resources. Figure B.5 shows that traffic class goals are set as percentage goals for the existing resources and traffic classes. Finally, the user can build resource allocation schemes as shown in Figure B.6.

4.3 Results Analysis

ERAP records a large amount of pertinent data for each flight and each resource as it executes a resource allocation program. The ERAP functions described in this section make use of data from a resource allocation to provide insight into the effectiveness of a resource allocation.

4.3.1 View Statistics and Graphs

The GUI shown in Figure B.7 allows a user to analyze flight assignment data in a number of ways. A user can quickly view a delay statistics report as shown in Figure B.8. ERAP also supports the creation of histograms, delay vs. assignment time graphs, and delay by airport, airline, and traffic class category graphs. These visual methods of describing the results of a resource allocation program appear in Figures B.9, B.10, and B.11.

The data set used for statistical and graphical analysis can be filtered by airport, airline, traffic class, or a percentage of the most delayed flights in order to investigate

specific groups of interest. Finally, if the default ERAP functions do not meet a user's needs, the resource allocation data can be output to a text file for spreadsheet analysis.

4.3.2 View Flight Tracks

ERAP can visually display flight tracks. The GUI in Figure B.12 permits a user to view flight tracks over a time period before and after resource allocation. ERAP permits the user to filter the data by flight identification number, airline, airport, and traffic class in order to limit the number of flight tracks shown. ERAP also provides the option to visualize traffic class sets using color.

4.3.3 View Resource Utilization

The "View Resource Utilization" function displays resource characteristics that result from running a resource allocation program. This information includes the measure of utilization for each resource as well as the actual assignments of flights to resource time slots. This information is displayed textually by ERAP as shown in Figure B.13.

4.3.4 View Traffic Class Goal Deviation

This results analysis function displays the outcome of a resource allocation as it pertains to traffic classes. This information is displayed textually as shown in Figure B.14. The final traffic class deviations are displayed for every hour time period of every resource.

Chapter 5.

Experimental Results

In this chapter, we describe a set of experiments carried out using ERAP. The results of these experiments provide some preliminary insight into the design of en route resource rationing priority schemes. In addition, they demonstrate the variety of analyses that can be carried out using ERAP. ERAP can be used in many ways to yield resource rationing insight and to point out strengths and weaknesses of various rationing algorithms. All experiments compare different priority algorithms for a controlled scenario.

The RBS priority algorithm has been successfully implemented as an equitable solution for airport arrival slot rationing in GDPs. Chapter 3 pointed out differences between RBS and accrued delay, and we believe that a TOAD priority function is the more suitable method for rationing en route airspace resources. In addition to the individual purposes of the experiments in this chapter, each one includes a comparison of the results of RBS and TOAD rationing solutions. It is shown that these two priority algorithms are very similar for the class of allocation problems considered.

ERAP provides several ways to analyze the results of a resource allocation. In the following experiments, resource utilization, average total delay per flight, and the percentage of flights rerouted are used to measure allocation efficiency. Utilization is

calculated at a resource as the fraction of the number of flights assigned to the total number of time slots available between the first and last flight to use the resource. Total delay is the difference between a flight's scheduled and controlled time of arrival.

Measures of allocation equity include the average total delay of the twenty percent most delayed flights, maximum delay, and delay standard deviation. Allocated delay, defined as total delay minus preprogram delay, can be another useful measure in some scenarios. One common practice in enforcing equity is to distribute equal parts of a limited resource to all interested parties [22]. The total delay standard deviation is an indication of divergence from this measure of equity. ERAP also supports visual methods of investigation, such as histograms and graphs of delay assignments over time, that are useful in comparing different resource rationing algorithms.

All experiments apply the scenario that is discussed in Chapter 2, and each experiment runs a resource allocation program from 15:00 to 17:00 Zulu time. This time period targets 147 flights for program inclusion. The current time field for program execution (see Figure B.4: ERAP GUI for Allocating Resources) is set to 10:00 Zulu. This time is early enough that all flights are on the ground at the time of running the resource allocation program.

5.1 Experiment One: Accrued Delay, RBS, and the Leapfrog Principle

The goal of this experiment is to point out a pitfall in basing resource allocation priority solely on the most accrued delay function (a common problem of any similar algorithm that schedules too far ahead in its solution). Specifically, we show that the most accrued

delay function should be preceded by an implementation of the earliest demanded time slot function in order to force time ordering into the allocation solution. This experiment is also designed to point out that proper use of accrued delay, in the form of TOAD, very closely resembles RBS.

5.1.1 Scenario Description

In this experiment, a single resource, used by all of the flights' preferred routes, is available for allocation. Assuming that the resource would require about 75 time slots per hour to accommodate normal operations (147 flights over 2 hours), the resource is reduced to 66% capacity (this gives 50 time slots per hour). A single flight, COA1254, has one hour of ground delay due to mechanical problems. All other flights are running on schedule. The three priority algorithms used for resource allocation in this analysis are shown below.

- Priority Algorithm 1: Straight Accrued Delay
    - Level 1: Most Accrued Delay (Range 0 Minutes)

- Priority Algorithm 2: TOAD
    - Level 1: Earliest Demanded Time Slot (Range 5 Minutes)
    - Level 2: Most Accrued Delay (Range 0 Minutes)

- Priority Algorithm 3: RBS
    - Level 1: Earliest Scheduled Meter Time (Range 0 Minutes)

5.1.2 Results

The results of using the different priority algorithms to ration airspace for the controlled scenario in this experiment are listed in Table 5.1. All three algorithms yield identical utilization and almost the same average total delay. Thus, these algorithms are effectively equivalent from an efficiency standpoint.

Table 5.1: Results from Experiment One

	Resource Allocation Priority Algorithm		
Measure of Effectiveness	**Straight Accrued Delay**	**TOAD**	**RBS**
Utilization	100%	100%	100%
Avg. Total Delay	30.32 min.	30.32 min.	30.31 min.
Avg. Total Delay of Top 20%	102.76 min.	50.11 min.	50.11 min.
Maximum Delay	141.35 min.	60.00 min.	60.00 min.
Standard Deviation	42.77 min.	14.36 min.	14.38 min.

On the other hand, the manner in which delay is distributed varies. This is best shown in the average total delay of the top twenty percent metric, where, from an equity point of view, the straight accrued delay algorithm is clearly inferior to the others. The metric shows that, when using straight accrued delay, the 29 flights with the most delay have over twice the average delay as those that are allocated using RBS or TOAD. The total delay standard deviation is another source of insight into the results. It shows that the allocation performed by the straight accrued delay allocation has much higher variability than the other two rationing schemes.

We refer to the cause for the poor performance of the straight accrued delay function as the Leapfrog Principle. The graph in Figure 5.1 illustrates this point. The straight accrued delay function actually achieves its objective. As a greedy algorithm based purely upon accrued delay, it assigns the COA1254 flight with 60 minutes of delay first. Due to that allocation, the only flights with accrued delay are the ones that

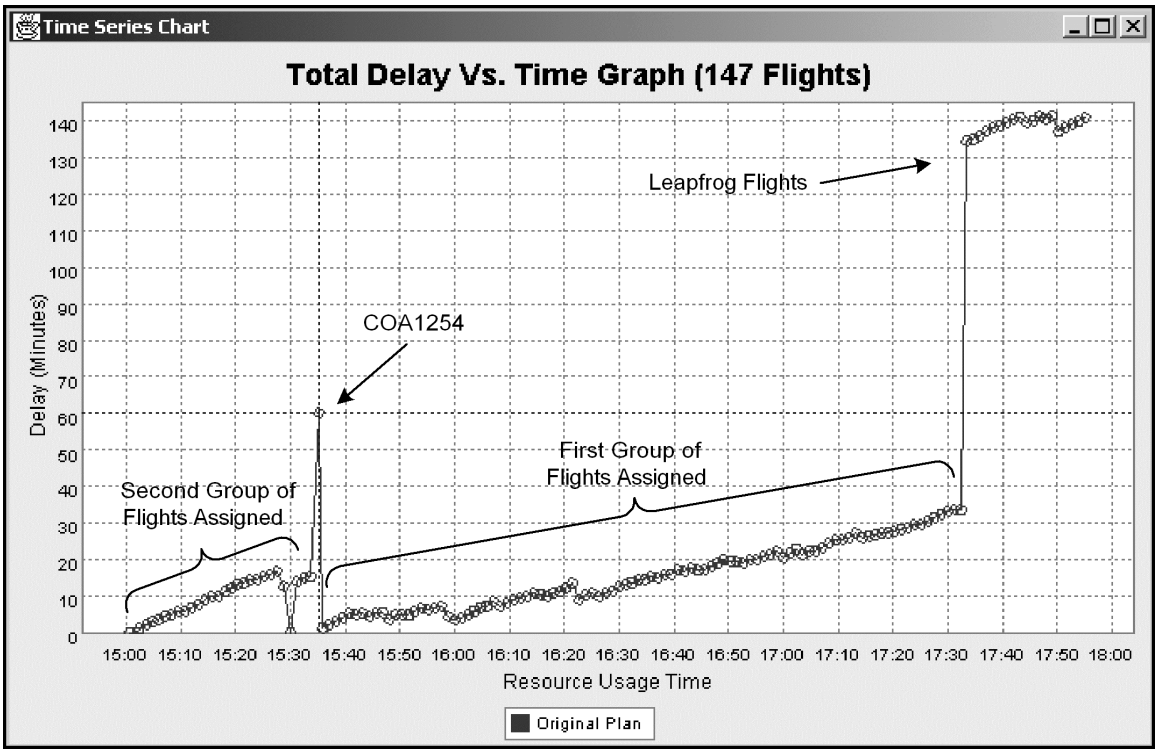

Figure 5.1: Delay vs. Time using Straight Accrued Delay in Experiment One

originally could have used the time slot given to COA1254. It follows that all of the flights with original best time slots that are at the same time as, or after, the one assigned to COA1254, are allocated by ERAP before any flights with ideal time slots that occur before COA1254's assigned slot. Since there are not enough slots existing before the COA1254 slot to accommodate the remaining number of flights, some flights receive exceptionally high amounts of delay by "leaping" over a large number of previously assigned slots. This inequitable distribution of delay can be averted by assigning flights in time order.

Figure 5.2 and Figure 5.3 show graphs of the total amount of delay vs. resource usage times for the TOAD and RBS rationing algorithms. Note that the graph for RBS is nearly indistinguishable from time-ordered accrued delay, and none of the priority functions in this experiment add any additional delay to COA1254.

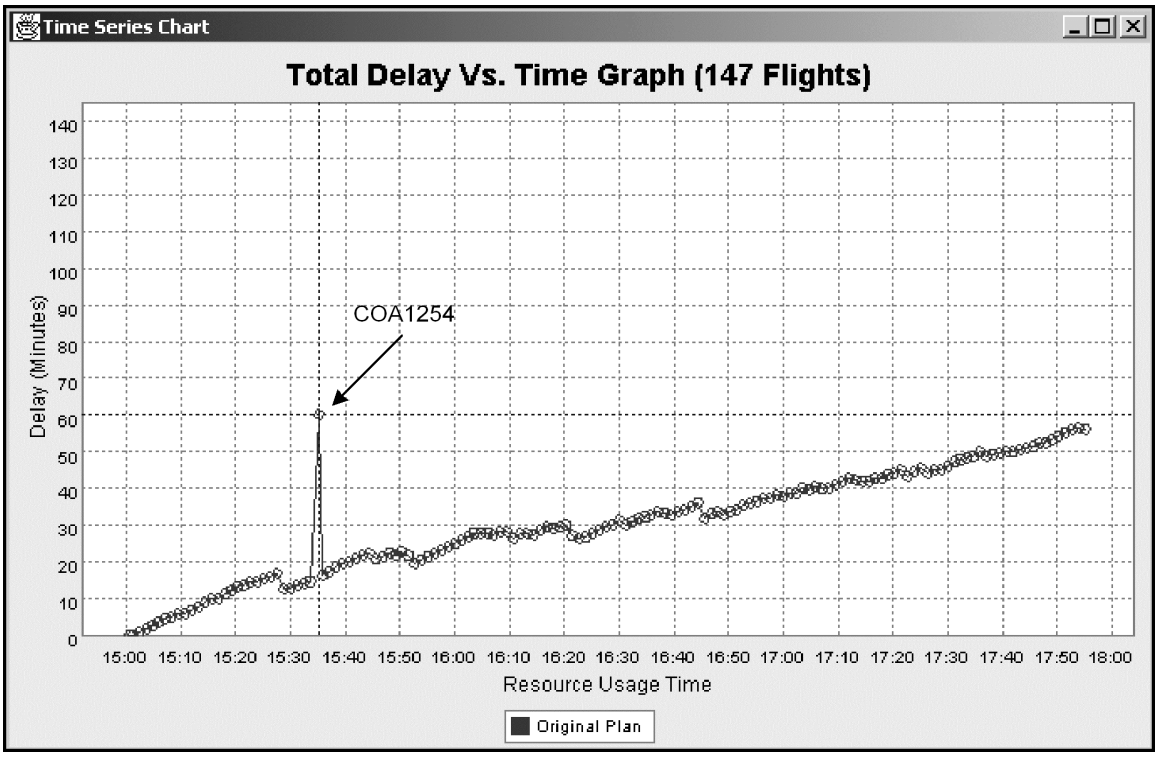

Figure 5.2: Delay vs. Time using TOAD in Experiment One

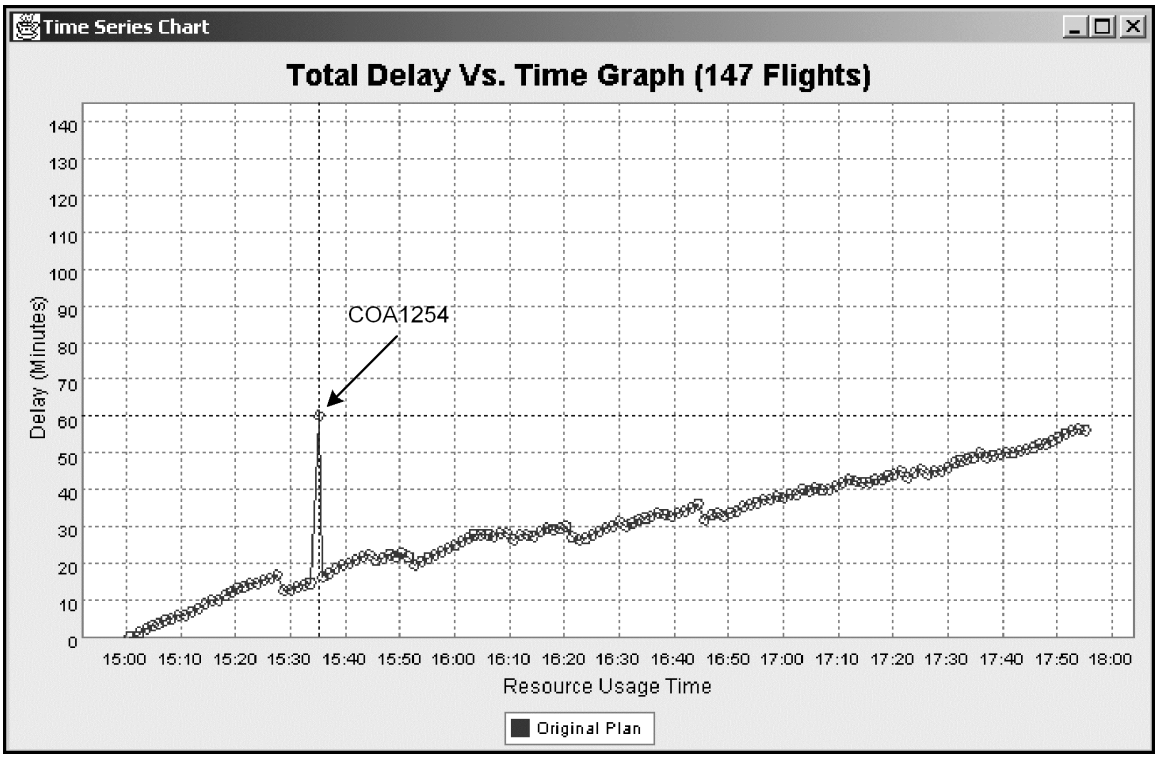

Figure 5.3: Delay vs. Time using RBS in Experiment One

The measures of effectiveness and the graphs show that TOAD is nearly identical to RBS when rationing a single resource. However, RBS and TOAD do yield a slightly different average total delay. This result follows from the ERAP assignment process and is illustrated in Figure 5.4.

Using TOAD, all flights (A, B, C, and D) in Figure 5.4 initially have no accrued delay (a first assignment is required for any delay to accrue). For the first assignment,

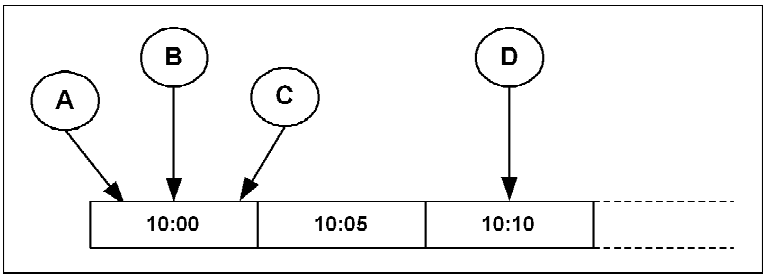

Figure 5.4: Comparing TOAD to RBS for the First Flight Assignment

TOAD chooses to assign either flight A, B, or C because they demand the first time slot. The algorithm cannot differentiate among these flights because no time slots have been used and no delay has accrued. RBS would assign Flight A first because it has the earliest scheduled time of resource usage. However, after a few assignments in TOAD, delay accrues and the algorithm converges to RBS. An investigation into the results found that three of the first five assignments from Experiment One are given different amounts of delay in TOAD than under RBS. All other assignments are the same.

5.2 Experiment Two: Traffic Class "Equity"

In this experiment, traffic classes are used to redistribute delay. The experiment exhibits TOAD and RBS results both with and without traffic class deviation priority. It is again shown that TOAD and RBS achieve similar rationing solutions.

5.2.1 Scenario Description

This experiment uses the same scenario as Experiment One. The only difference is that simple traffic class goals are enacted. The arbitrarily selected goal of this experiment is to restrict the number of time slots allotted to GA flights per hour. This could be used to address inequities that result because air carriers must publish schedules well in advance and GA flights can announce departure times on the day of operation. Of the 147 targeted flights, 22.4% are general aviation. For this experiment, a Blue traffic class that includes all GAs is defined, and it is given a goal of 10% for usage of the single resource. Another goal of 90% is set for the Other traffic class. The resource priority algorithms for this experiment are shown below:

- Priority Algorithm 1: TOAD Before Traffic Class Priority
    - Level 1: Earliest Demanded Time Slot (Range 5 Minutes)
    - Level 2: Most Accrued Delay (Range 0 Minutes)

- Priority Algorithm 2: RBS Before Traffic Class Priority
    - Level 1: Earliest Scheduled Meter Time (Range 0 Minutes)

- Priority Algorithm 3: TOAD After Traffic Class Priority
    - Level 1: Earliest Demanded Time Slot (Range 5 Minutes)
    - Level 2: Highest Traffic Class Deviation (Range 0%)
    - Level 3: Most Accrued Delay (Range 0 Minutes)

- Priority Algorithm 4: RBS After Traffic Class Priority
    - Level 1: Highest Traffic Class Deviation (Range 0%)
    - Level 2: Earliest Scheduled Meter Time (Range 0 Minutes)

5.2.2 Results

The aggregate results in Table 5.2 show that incorporating traffic class priority into the rationing algorithms does not significantly affect efficiency, as the utilization and average total delay results mirror the results without traffic classes. These results do, however, differ in the average delay of the top twenty percent, maximum delay, and total delay standard deviation measures of equity. This is expected, as the objective of this experiment is to honor an equity principle that adds delay to the GA flights.

Table 5.2: Results from Experiment Two

	Resource Allocation Priority Algorithm			
Measure of Effectiveness	**TOAD Before Class Priority**	**RBS Before Class Priority**	**TOAD After Class Priority**	**RBS After Class Priority**
Utilization	100%	100%	100%	100%
Avg. Total Delay	30.32 min.	30.31 min.	30.32 min.	30.32 min.
Avg. Total Delay of Top 20%	50.11 min.	50.11 min.	87.87 min.	87.87 min.
Maximum Delay	60.00 min.	60.00 min.	119.08 min.	119.08 min.
Standard Deviation	14.36 min.	14.38 min.	30.71 min.	30.72 min.

Figure 5.5 is a good example of the results of using traffic classes to redistribute delay. It shows that delay is moved from the Other traffic class to the GA (Blue) traffic class while the average total delay across all flights remains constant. This Figure also shows that TOAD and RBS are virtually identical.

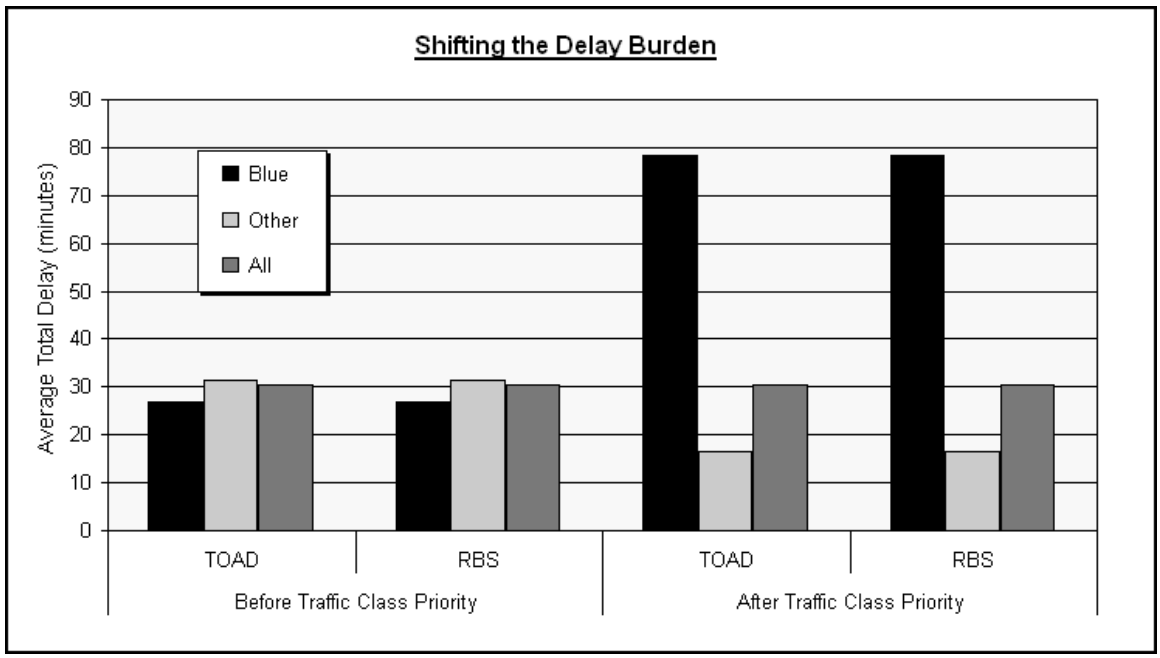

Figure 5.5: Using Traffic Classes to Reallocate Delay Burden in Experiment Two

5.3 Experiment Three: Double Penalty

This third ERAP experiment is used to illustrate that a "double penalty" pitfall, a well-documented concept proven to occur in the Grover-Jack algorithm in pre-CDM implementations of GDPs [20], can also arise when rationing en route resources. This example demonstrates ERAP capabilities and also shows that the double penalty pitfall can be avoided by using TOAD (in a similar fashion to RBS).

5.3.1 Scenario Description

This experiment uses the same single resource as Experiments One and Two. However, instead of having a single flight with pre-allocation delay, two GDPs are simulated – one at Boston's Logan airport and one at New York's LaGuardia airport. The GDP simulations affect 23 of the 147 flights. These 23 flights have delays that range from 64

to 170 minutes, and their average preprogram ground delay is 125 minutes. The three priority algorithms selected for comparison appear below:

- Priority Algorithm 1: Grover-Jack
    - Level 1: Earliest Expected Meter Time (Range 0 Minutes)

- Priority Algorithm 2: TOAD
    - Level 1: Earliest Demanded Time Slot (Range 5 Minutes)
    - Level 2: Most Accrued Delay (Range 0 Minutes)

- Priority Algorithm 3: RBS
    - Level 1: Earliest Scheduled Meter Time (Range 0 Minutes)

5.3.2 Results

The results listed in Table 5.3 show a distinct difference in equity (maximum delay, delay standard deviation, and assigned delay distribution) for allocations that yield similar efficiency (utilization and average total delay per flight). All measures of equity shown in show worse results for Grover-Jack than for either TOAD or RBS. In fact, Grover-Jack allocates an average of over 25% more delay to the top twenty percent highest delayed flights than either TOAD or RBS. The reason for the equity imbalance is that Grover-Jack does not credit the GDP flights for the delay they have accumulated prior to the resource allocation. In this experiment, Grover-Jack assigns an average of 55.12 minutes of extra delay to the GDP flights while TOAD and RBS assign an average

Table 5.3: Results from Experiment Three

	Resource Allocation Priority Algorithm		
Measure of Effectiveness	**Grover-Jack**	**TOAD**	**RBS**
Utilization	100%	100%	100%
Avg. Total Delay	49.52 min.	49.59 min.	49.58 min.
Avg. Total Delay of Top 20%	139.31 min.	111.03 min.	111.03 min.
Maximum Delay	225.80 min.	170.00 min.	170.00 min.
Standard Deviation	52.73 min.	38.28 min.	38.42 min.
Avg. Assigned Delay to GDP Flights	55.12 min.	0.14 min.	0.14 min.

of 0.14 minutes. The extra delay that Grover-Jack gives to the GDP flights is commonly referred to as a "double penalty" in ATM.

The delay inequity can easily be seen in the ERAP-generated histograms shown in Figure 5.6, Figure 5.7, and Figure 5.8 below. Note that Grover-Jack increases the delays given to the GDP flights, and in the TOAD and RBS distributions, the delays are much more tightly bound to the average than in the Grover-Jack solution.

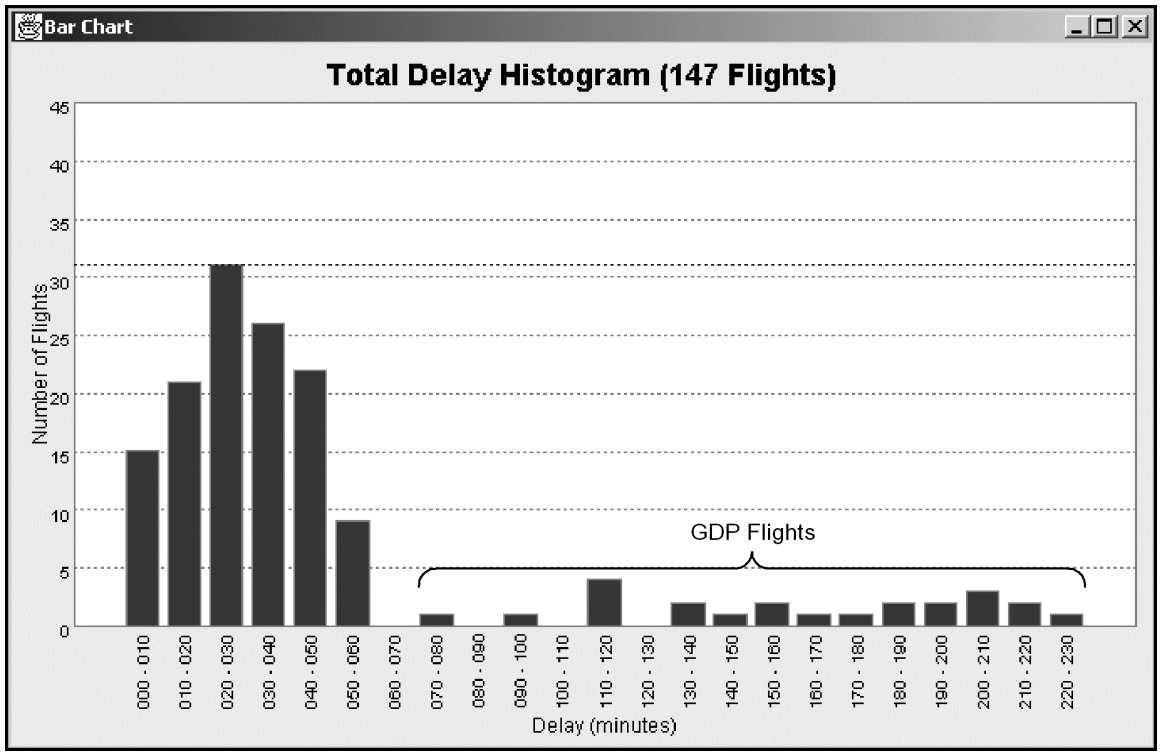

Figure 5.6: Histogram for Grover-Jack Allocation in Experiment Three

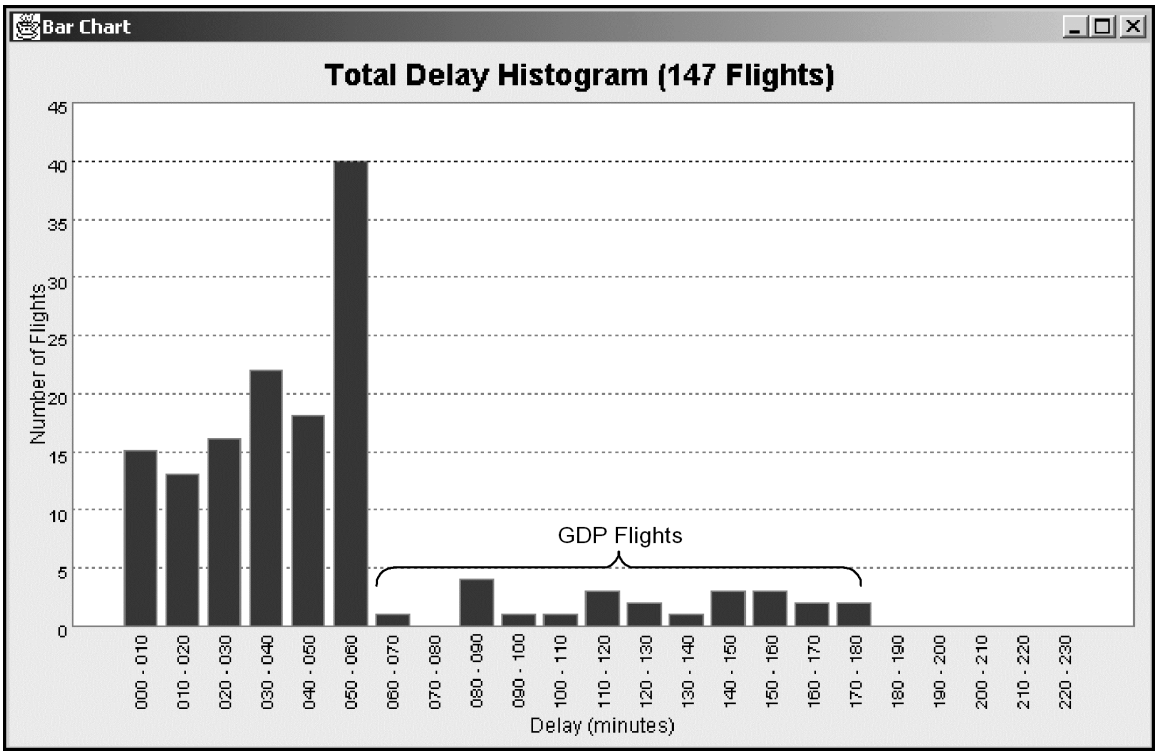

Figure 5.7: Histogram for TOAD Allocation in Experiment Three

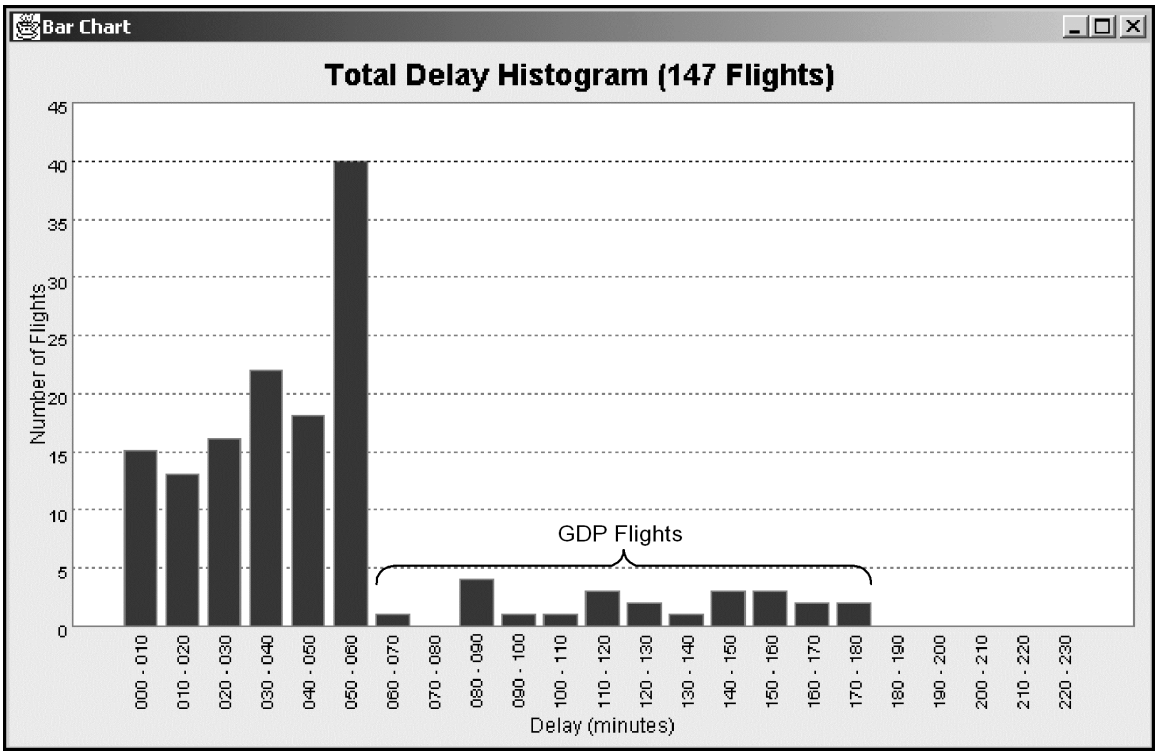

Figure 5.8: Histogram for RBS Allocation in Experiment Three

5.4 Experiment Four: Alternate Routes and Tuning TOAD

This fourth and final experiment uses alternate resources to demonstrate two important ERAP lessons. It is shown that using alternate resources can increase the variability of the results between TOAD and RBS. Also, the range used in TOAD's earliest demanded time slot priority function can significantly affect the results.

5.4.1 Scenario Description

This scenario uses all five resources that are described in Chapter 2. All flights that do not either originate or terminate at airports located in the FCA have three flight plans: a preferred route through the FCA, one north of the FCA, and one south of the FCA. Flights using airports located in the FCA have only a single flight plan because it is impossible for them to reroute around the FCA. All flight delay thresholds are set to zero so that a flight will always prefer the fastest route to the destination. The FCA resource is reduced in capacity from Experiment Three to 30 time slots per hour, and the four alternate resources each have a capacity of 7.5 time slots per hour. Thus, there is a total potential capacity of 60 flights per hour, which is 20% greater than the total capacity in Experiment Three. This experiment uses the same GDP as Experiment Three, with 23 flights having pre-allocation delay. The algorithms selected for comparison are listed below:

- Priority Algorithm 1: TOAD Implementation One
    - Level 1: Earliest Demanded Time Slot (Range 0 Minutes)
    - Level 2: Most Accrued Delay (Range 0 Minutes)

- Priority Algorithm 2: TOAD Implementation Two
    - Level 1: Earliest Demanded Time Slot (Range 32 Minutes)
    - Level 2: Most Accrued Delay (Range 0 Minutes)

- Priority Algorithm 3: RBS
    - Level 1: Earliest Scheduled Meter Time (Range 0 Minutes)

5.4.2 Results

Figure 5.9 shows the visual results of using alternate resources to alleviate congestion. The image shows the level of air traffic through the FCA before and after a resource allocation performed with TOAD Implementation Two for the 15:00 to 17:00 Zulu time period.

The process used to select the two variations of TOAD used in this comparison requires justification. One purpose of this experiment is to show the impact of altering the range used in TOAD's earliest demanded time slot priority function. This range variable is important when simultaneously allocating to multiple resources. If set correctly, it promotes competition among the various resources.

As an example of how the earliest demanded time slot range affects resource competition in this experiment, assume that the range is set to zero minutes. Also assume

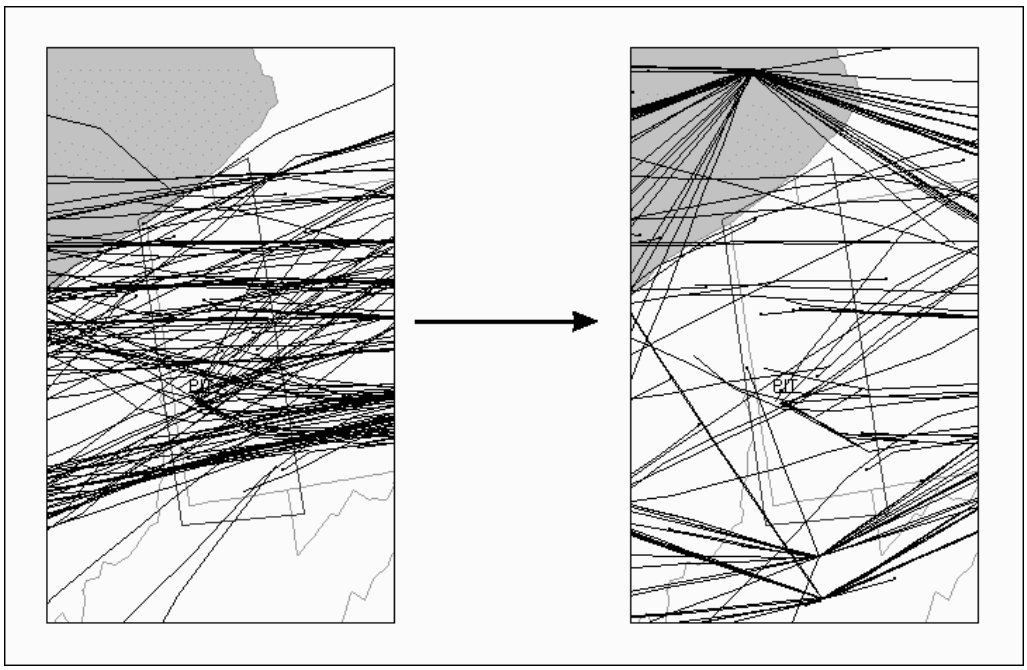

Figure 5.9: Using Alternate Routes to Alleviate Congestion in Experiment Four

that all resources' first time slot starts at 15:00 and that one slot from each resource has been allocated. As the four alternate resources have time slots that are 8 minutes in length and the FCA resource time slots are 2 minutes long, the resulting earliest time slot for all of the alternate resources is 15:08. The earliest time slot of the FCA resource is 15:02. If, at the next point in assignment, any single flight demands the 15:02, 15:04, or 15:06 time slots of the FCA resource, then it is impossible for a flight to be assigned to an alternate resource, regardless of its accrued delay, because only the FCA resource will be considered.

The range values used in this experiment are chosen by investigating the average total delay and average total delay of the top twenty percent measures that result from increasing the range value by 8-minute increments. The 8-minute increment size used in this experiment is chosen to equal the size of the time slots in the most sparse resource (the alternate resources) so that each increment will extend the range of comparison by a

further time slot unit. The results of this selection process are shown below in Table 5.4. Mainly due to the average total delay of the top twenty percent metric, range values of 0 and 32 minutes are selected to represent the respective bad and good implementations of TOAD for this scenario.

Table 5.4: Selecting the Range Values for TOAD In Experiment Four

Range	Avg. Total Delay	Avg. Total Delay of Top 20%
0	37.12	110.53
8	36.27	109.23
16	36.42	108.71
24	35.42	108.61
32	35.58	108.04
40	35.72	108.10
48	35.80	109.16

The results from the resource allocations in Experiment Four are listed in Table 5.5. The utilization, average total delay, and percent rerouted measures indicate slight differences in system efficiency among the three algorithms. This means that the range value used in the earliest demanded time slot function in TOAD can have a slight impact upon allocation efficiency.

The standard deviation, maximum delay, and average total delay of the top twenty percent metrics appear to show very similar equity results for the algorithms. However, the highly delayed GDP flights, which comprise 16% of the total flights, dilute these results. An investigation of the average assigned delay to the top twenty percent metric produces a glaring weakness in the first implementation of TOAD, as it assigns an average of 29% (44.7 minutes) more delay to the most delayed flights than the second implementation of TOAD (34.63 minutes).

The graph in Figure 5.10 exhibits the assigned delay distributions. It shows that the first implementation of TOAD gives less delay than the other two algorithms for a

Table 5.5: Results from Experiment Four

	Resource Allocation Priority Algorithm		
Measure of Effectiveness	**TOAD 0 Minute Range**	**TOAD 32 Minute Range**	**RBS**
FCA Resource Utilization	100% (79 flights)	100% (74 flights)	100% (75 flights)
North Westerly Resource Utilization	76.47% (13 flights)	94.11% (16 flights)	88.24% (15 flights)
North Easterly Resource Utilization	100% (19 flights)	100% (20 flights)	100% (20 flights)
South Westerly Resource Utilization	100% (17 flights)	100% (18 flights)	100% (18 flights)
South Easterly Resource Utilization	100% (19 flights)	100% (19 flights)	100% (19 flights)
Avg. Total Delay	37.12 min.	35.58 min.	35.85 min.
Percent Rerouted	46.26%	49.66%	48.98%
Avg. Total Delay of Top 20%	110.53 min.	108.04 min.	108.11 min.
Maximum Delay	170.00 min.	170.00 min.	170.00 min.
Standard Deviation	42.67 min.	42.13 min.	42.03 min.
Avg. Assigned Delay of Top 20%	44.70 min.	34.63 min.	34.98 min.

number of the lower delayed flights, but the higher delays are much more severe. This graph also shows that the second TOAD implementation is very similar to RBS in terms of assigned delay distribution.

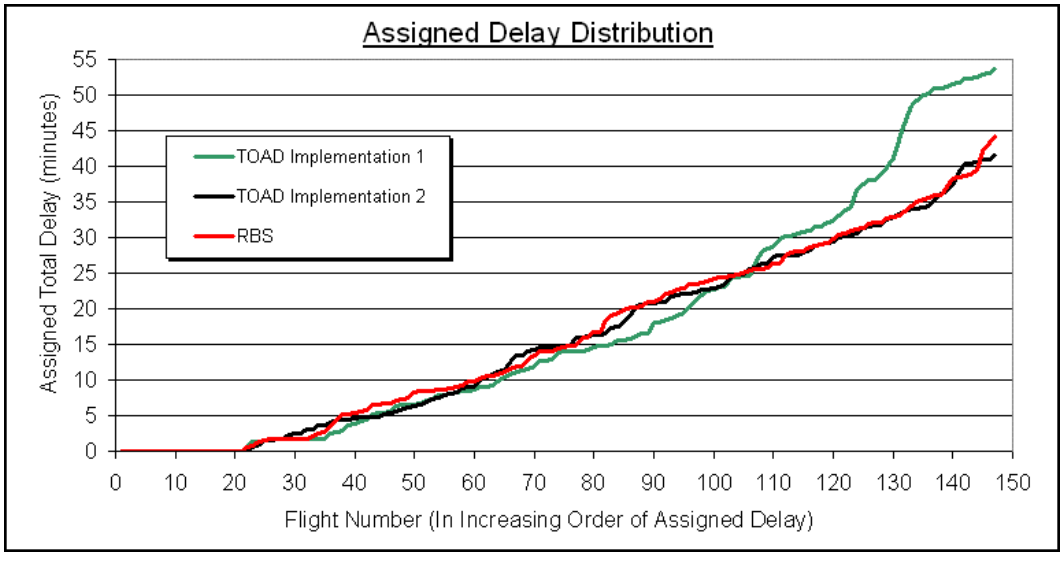

Figure 5.10: Assigned Delay Distributions in Experiment Four

5.5 Summary of Experiments

The experiments in this chapter have demonstrated the power of ERAP for building and testing en route resource rationing schemes. It is shown that time ordering is advisable when implementing an accrued delay rationing scheme because, otherwise, the schedule can result in a high level of variance in the distribution that is probably unpalatable. However, the TOAD implementation achieves results are very similar to RBS. This chapter has also demonstrated that Grover-Jack performs poorly from an equity standpoint as a result of a double penalty phenomenon, and traffic classes can be used to redistribute delay among meaningful sets of flights.

Chapter 6.

Conclusions

Since this thesis provides an initial implementation of a number of alternative concepts defined by the Long-Term CR group, it furthers the goal of CR to develop efficient and equitable procedures for rationing congested en route airspace. A feasible en route rationing implementation (ERAP) is developed to serve as a baseline for the future comparison and final industry acceptance of an ideal CR rationing solution. ERAP implements a number of proposed CR concepts so that the arena of comparison of these ideas can move from the theoretical domain to the practical. It is shown that alternate route filing by flights gives airlines more control over business operations, and priority-based rationing can be employed to enforce user-defined equity principles.

In this thesis, a strong case is made for accrued delay as a suitable priority mechanism for assigning en route resources, and several lessons are derived from initial ERAP experimentation. When rationing a single resource within the ERAP framework, it is shown that accrued delay priority, when properly implemented in the form of TOAD, mirrors the results of RBS – the current equity solution for GDPs. Also, experiments in this thesis demonstrate that efficiency-equivalent resource rationing solutions can differ significantly in measures of equity.

We feel that the ERAP framework, algorithm implementations, data requirements, and measures of effectiveness will provide a basis for further analysis of the en route resource rationing problem. It is also anticipated that some of the concepts presented in this thesis will find their way into the en route rationing solutions of the future.

6.1 Recommendations for Future Work

The outcome of this thesis led to a software prototype (ERAP) and provided an initial set of experimental results for evaluating en route resource rationing concepts. ERAP is designed as a transparent tool to facilitate the elucidation of broad en route resource rationing principles and goals. It is not a comprehensive representation of the entire en route problem. Therefore, we propose several avenues for building upon the work of this thesis.

We first recommend the evaluation of further scenarios. A larger number of resources, different kinds of resources (such as sectors), and more accurate data could be used to further differentiate between rationing algorithms. Also, the peripheral (non-target) flights ignored in this analysis could be addressed in future ERAP revisions through the use of traffic classes or a separate priority function.

Another ERAP improvement would be to add the capability to handle stochastic events. There are a number of stochastic elements in the en route airspace domain (such as weather and flight times) that could be accounted for in ERAP resource allocations to produce more meaningful results.

ERAP can handle parallel resources, but other rationing scenarios are likely to be comprised of interdependent resource networks. Determining how to measure priorities,

allocate resources, and maintain system efficiency in this environment represents a complex but necessary challenge for improving ERAP.

Finally, the most important step in ERAP's future is to be subjected to evaluation by representatives from the FAA, industry, and academia. It is by this collaborative process that the final requirements for an en route resource allocation system can be elicited.

APPENDIX A: ERAP DATABASE DEFINITIONS

ERAP uses a number of database tables to support its various functions. Table A.1 lists the data fields that are used during resource allocation. All flights that appear in this table are considered "target flights" and are included in a resource rationing program. These fields must be instantiated prior to using ERAP, and, where indicated, ERAP modifies some fields. ERAP can support any number of alternate routes.

Table A.2 shows the fields that are necessary to support traffic class definition in ERAP. All of these fields must exist for every target flight appearing in Table A.1 before using ERAP.

ERAP maintains a number of tables that describe flight trajectories that are used for displaying flight tracks. An example of the data needed for these tables is shown in Table A.3. A separate trajectory table is required for all target flights' preferred and alternate routes. These tables must exist before using ERAP. During resource allocation, ERAP builds a table of modified flight trajectories that accounts for flight delays and reroutes.

ERAP also maintains a table of descriptions for traffic classes as shown in Table A.4. ERAP modifies this table as a user defines traffic classes using the GUI shown in Figure B.3.

Table A.1: ERAP Database Table for Resource Allocation

Column Name	Data Type	Definition
FID	Integer	A unique number for each flight. Example: 49803536
ACID	Text	A flight's call sign. Example: AAL100
AIRLINE	Text	Abbreviation for a flight's operating airline. Example: AAL
DEPT_ARPT	Text	The code for a flight's departure airport. Example: EWR
ARR_ARPT	Text	The code for a flight's arrival airport. Example: EWR
TRAFFIC_CLASSES	Text	The list of traffic classes that this flight belongs to. ERAP modifies this field when making traffic classes. Example: BLUE RED
ORIG_DEPT_TIME	Date/Time	The original Zulu time when a flight is scheduled to depart. Example: 7/11/01 17:59:00
ORIG_ARR_TIME	Date/Time	The original Zulu time when a flight is scheduled to arrive at its destination (according to the preferred route). Example: 7/11/01 17:59:00
PRE_PROGRAM_GROUND_DELAY	Integer	The amount of time (in seconds) that a flight was ground delayed at the time when the user runs a program. ERAP modifies this field when simulating a ground delay program. Example: 350
PRE_PROGRAM_AIR_DELAY	Integer	The amount of time (in seconds) that a flight was air delayed at the time when the user runs a program. Example: 0

LABEL_PREF_ROUTE	Text	A flight's preferred resource name. Example: North Resource
METER_TIME_PREF_ROUTE	Date/Time	The Zulu time when a flight is predicted to use its preferred resource at the time when the user runs a program. Example: 7/11/01 17:59:00
SCHED_METER_TIME_PREF_ROUTE	Date/Time	The Zulu time when a flight is scheduled to use its preferred resource. ERAP modifies this field when simulating a ground delay program. Example: 7/11/01 17:59:00
NUMBER_ALT_ROUTES	Integer	The number of alternate routes that a flight has filed. ERAP can handle any number of alternate routes as long as the corresponding data fields (appearing below) exist. Example: 9
LABEL_ALT_ROUTE1	Text	The name of the resource used by a flight's first alternate route. Example: South West
METER_TIME_ALT_ROUTE1	Date/Time	The Zulu time when a flight is predicted to use its first alternate resource at the time when the user runs a program. Example: 7/11/01 17:59:00
SCHED_METER_TIME_ALT_ROUTE1	Date/Time	The Zulu time when a flight is scheduled to use its first alternate resource. ERAP modifies this field when a user simulates a ground delay program (see Figure B.1). Example: 7/11/01 17:59:00

DELAY_THRESHOLD_ALT_ROUTE1	Integer	A flight's specified delay threshold (in seconds) for its first alternate route. ERAP modifies this field when a user sets alternate route delay thresholds (see Figure B.2) Example: 300
EXTRA_FLIGHT_TIME_ALT_ROUTE1	Integer	The amount of extra time (in seconds) that it would take to fly a flight's first alternate route (over the preferred route) at the time when the user runs a program. Example: 300

Table A.2: ERAP Database Table of Traffic Class Criteria

Column Name	Data Type	Definition
FID	Integer	A unique number for each flight. Example: 49803536
ACT_DATE	Date/Time	The date (in Zulu time) that a flight is scheduled to operate. Example: 7/11/01 17:59:00
OPERATION	Text	Code for a region-specific operation: Arrival (A), Departure (D), Departure & Arrival (B), or Overflight (O). Example: A
DEPT_ARPT	Text	The code for a flight's departure airport. Example: EWR
ARR_ARPT	Text	The code for a flight's arrival airport. Example: EWR
PHYSICAL_CLASS	Text	This describes if a flight is a piston (P), jet (J), or turbo (T) craft. Example: P
USER_CLASS	Text	The flight's user class. It is either an Other (O), Air Taxi (T), Cargo (F), Commercial (C), General Aviation (G), or Military (M). Example: C
WEIGHT_CLASS	Text	This is the weight class of an aircraft, based upon wake vortices. Small (less than 41,000lb), Large (41,000-255,000lb), and Heavy (>255,000lb). Values can be S, L, or H. Example: S

Table A.3: ERAP Database Table of Flight Trajectories

Column Name	Data Type	Definition
FID	Integer	A unique number for each flight. Example: 49803536
ORIG_TIME	Date/Time	A timestamp for a point in a flight's trajectory. Example: 7/11/01 17:59:00
CUR_LAT	Integer	The latitude (in minutes) for a point in a flight's trajectory. Values are negative for the Southern Hemisphere (POET format). Example: 2073
CUR_LON	Integer	The longitude (in minutes) for a point in a flight's trajectory. Values are negative for the Eastern Hemisphere (POET format). Example: 7263

Table A.4: ERAP Database Table of Traffic Class Descriptions

Column Name	Data Type	Definition
COLOR	Text	A user-instantiated traffic class. Example: GREEN
DESCRIPTION	Text	An ERAP-generated description for a rule to add flights to a traffic class. Example: 7/11-7/11 RD-LGA, P=J, U=*, W=*
COUNT	Integer	The number of flights that meet this description. Example: 40

APPENDIX B: SCREENSHOTS OF ERAP GRAPHICAL USER INTERFACES

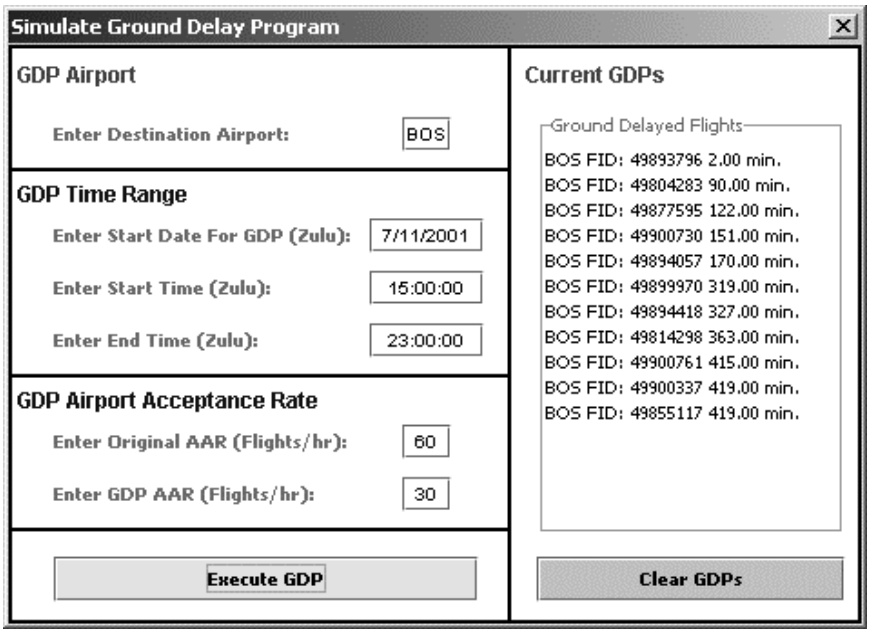

Figure B.1: ERAP GUI for Simulating a Ground Delay Program

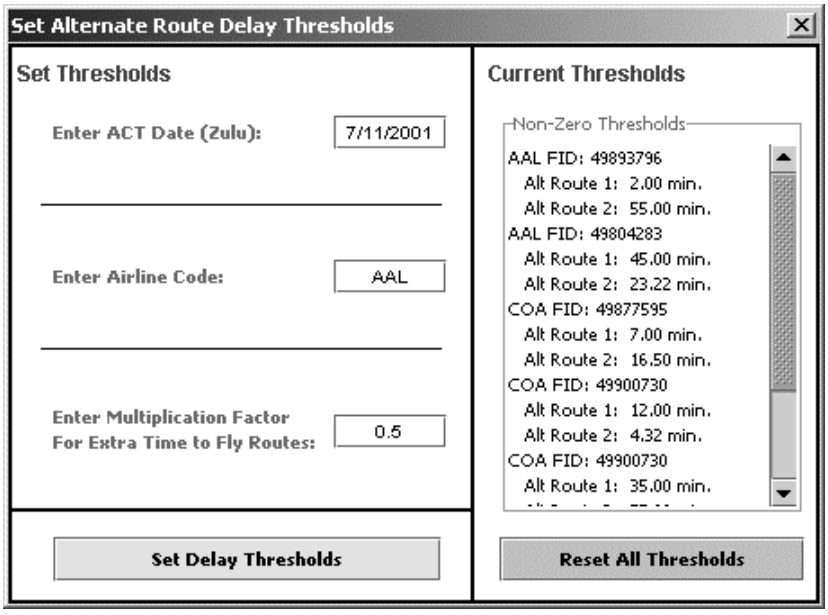

Figure B.2: ERAP GUI for Setting Alternate Route Delay Thresholds

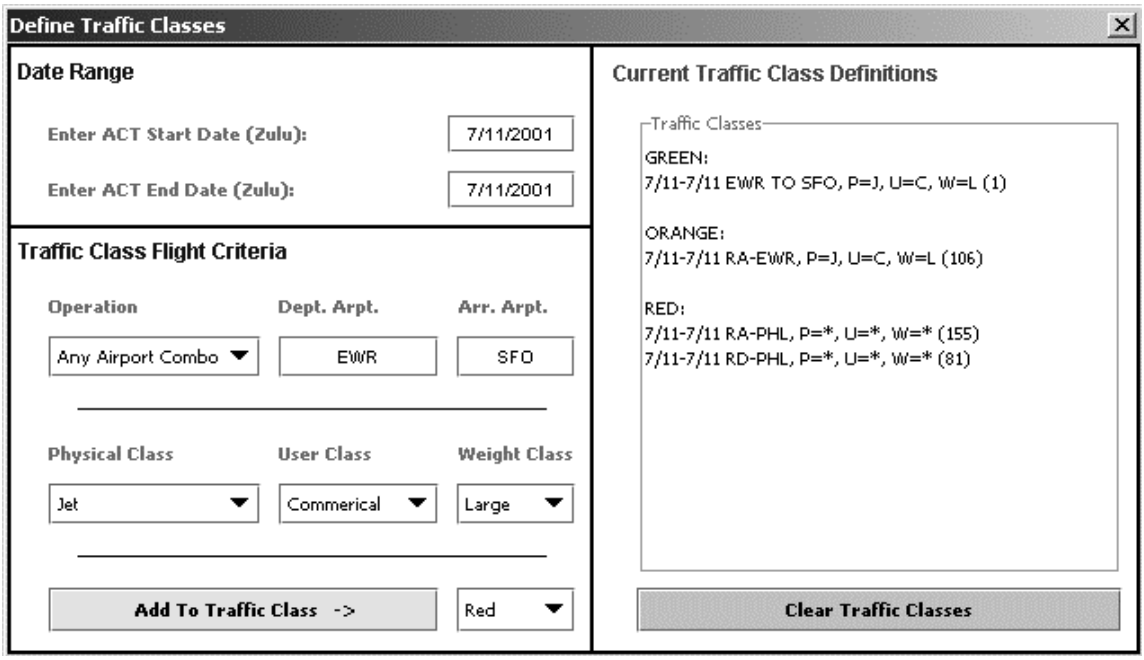

Figure B.3: ERAP GUI for Defining Traffic Classes

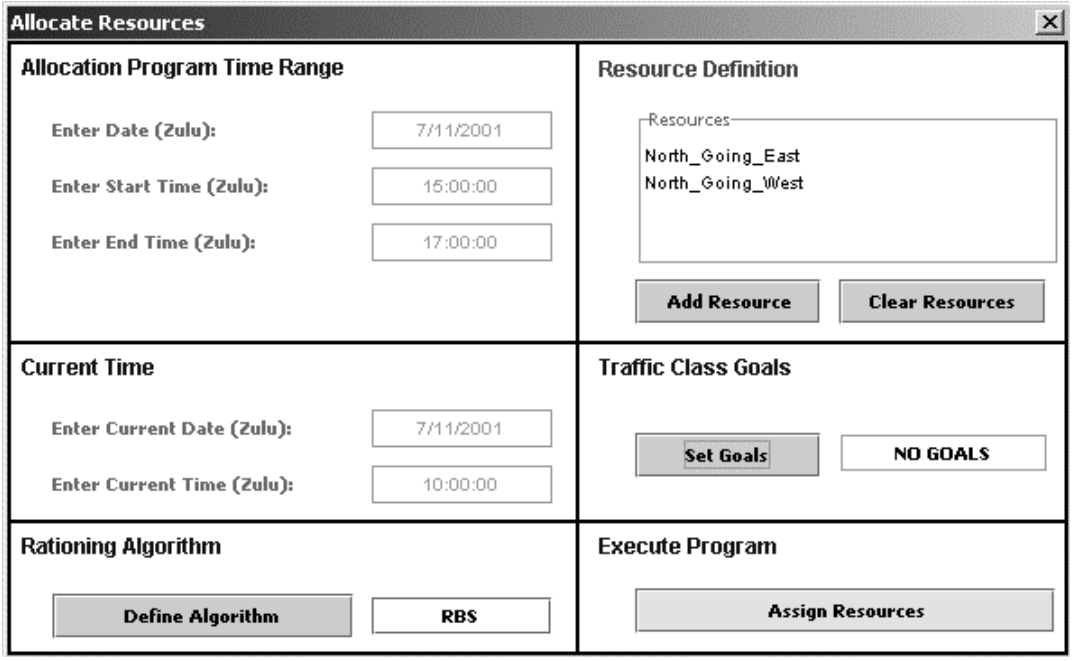

Figure B.4: ERAP GUI for Allocating Resources

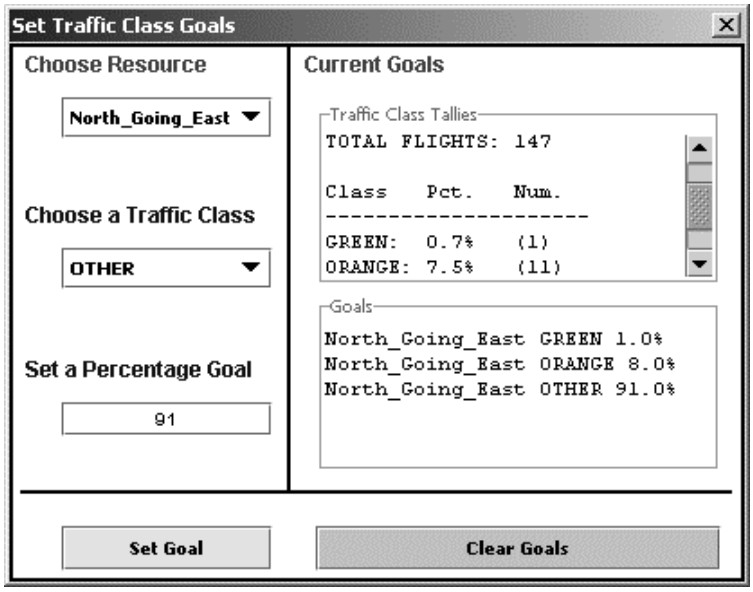

Figure B.5: ERAP GUI for Setting Traffic Class Goals

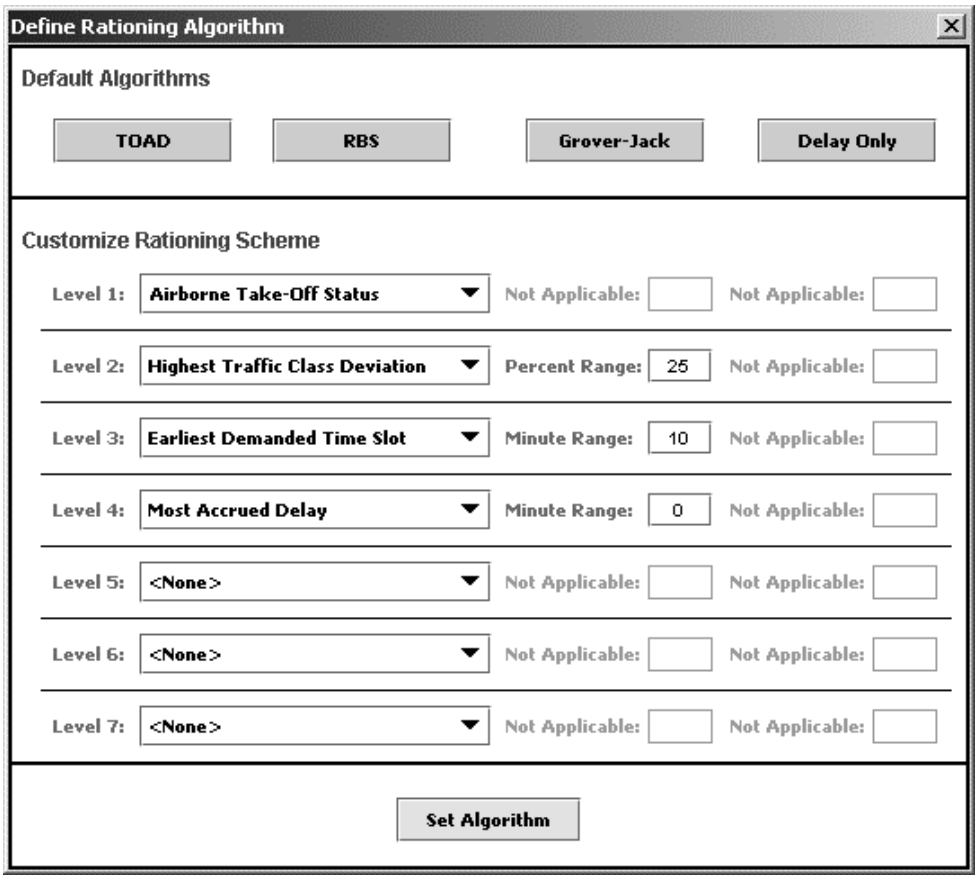

Figure B.6: ERAP GUI for Defining Rationing Schemes

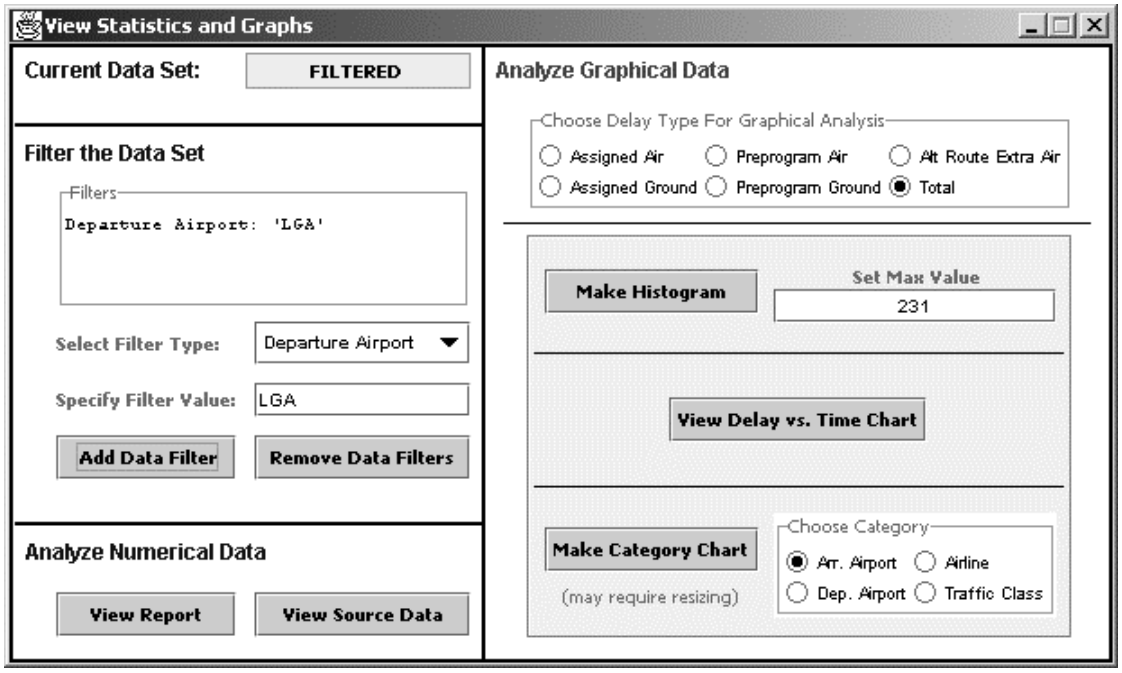

Figure B.7: ERAP GUI for Viewing Statistics and Graphs

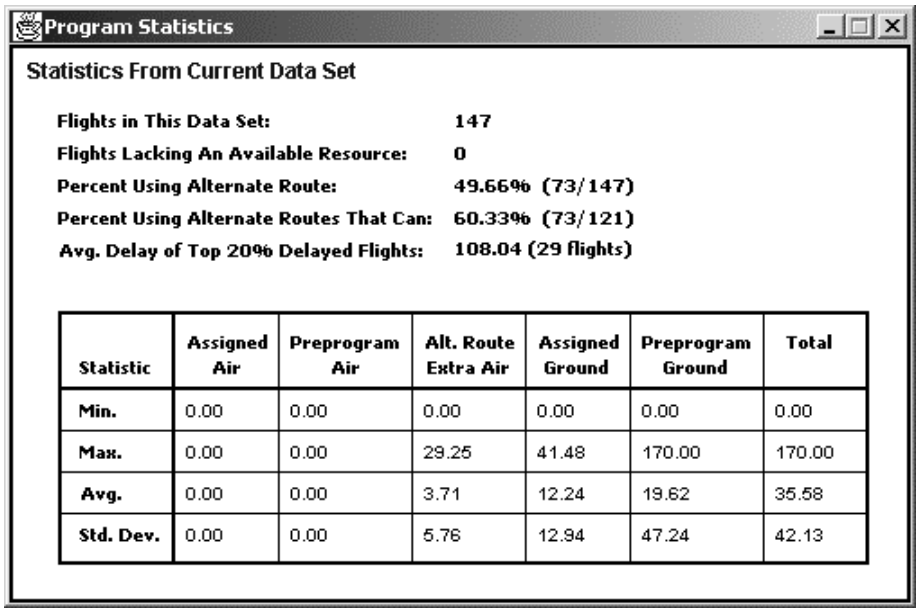

Statistic	Assigned Air	Preprogram Air	Alt. Route Extra Air	Assigned Ground	Preprogram Ground	Total
Min.	0.00	0.00	0.00	0.00	0.00	0.00
Max.	0.00	0.00	29.25	41.48	170.00	170.00
Avg.	0.00	0.00	3.71	12.24	19.62	35.58
Std. Dev.	0.00	0.00	5.76	12.94	47.24	42.13

Figure B.8: ERAP GUI for Displaying Delay Statistics

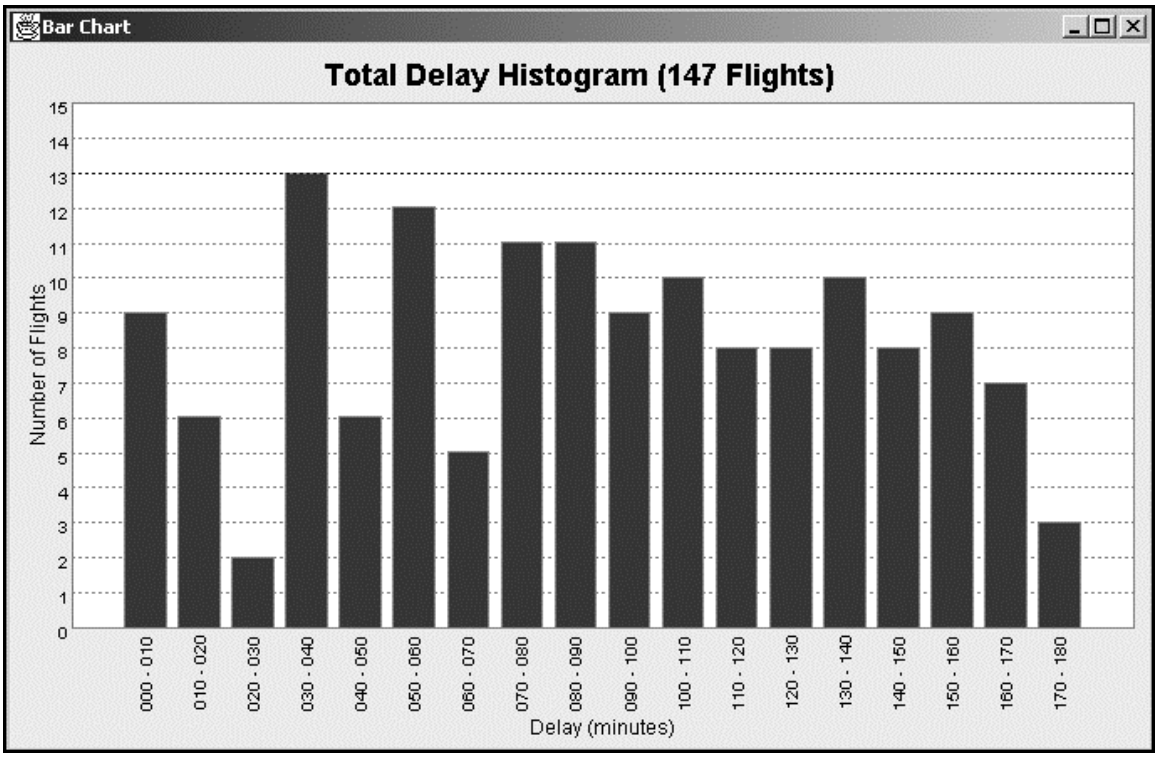

Figure B.9: Example of ERAP Delay Histogram

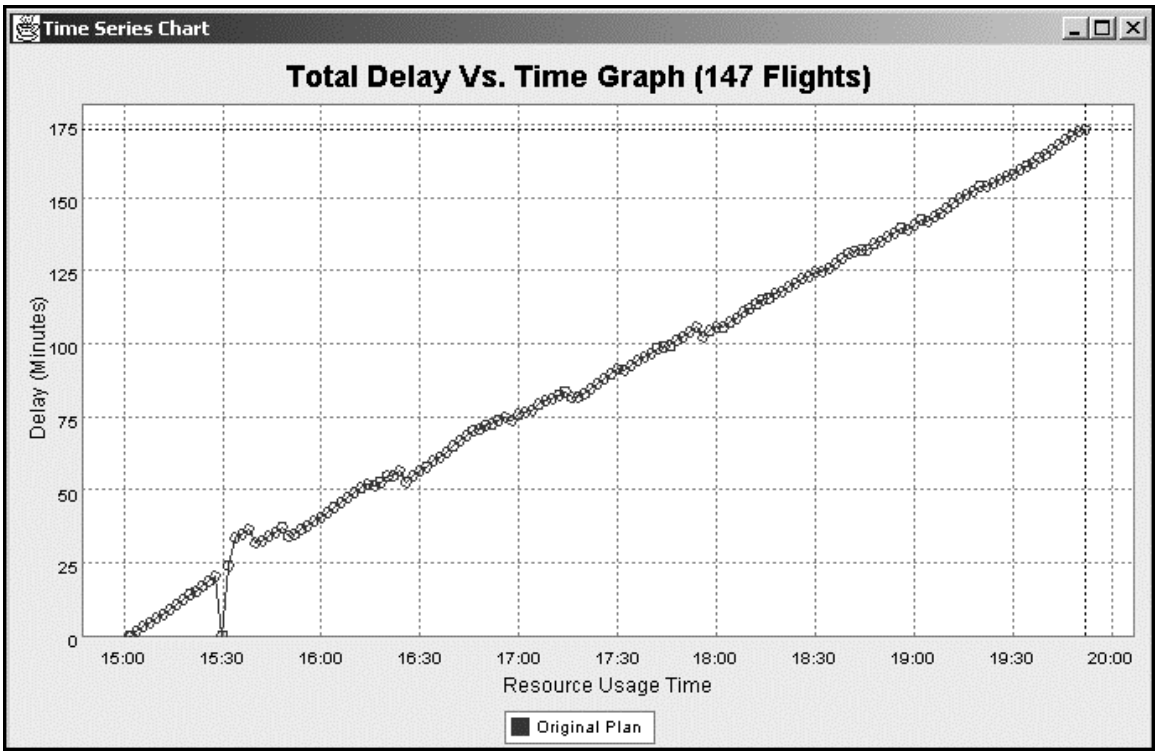

Figure B.10: Example of ERAP Delay vs. Time Graph

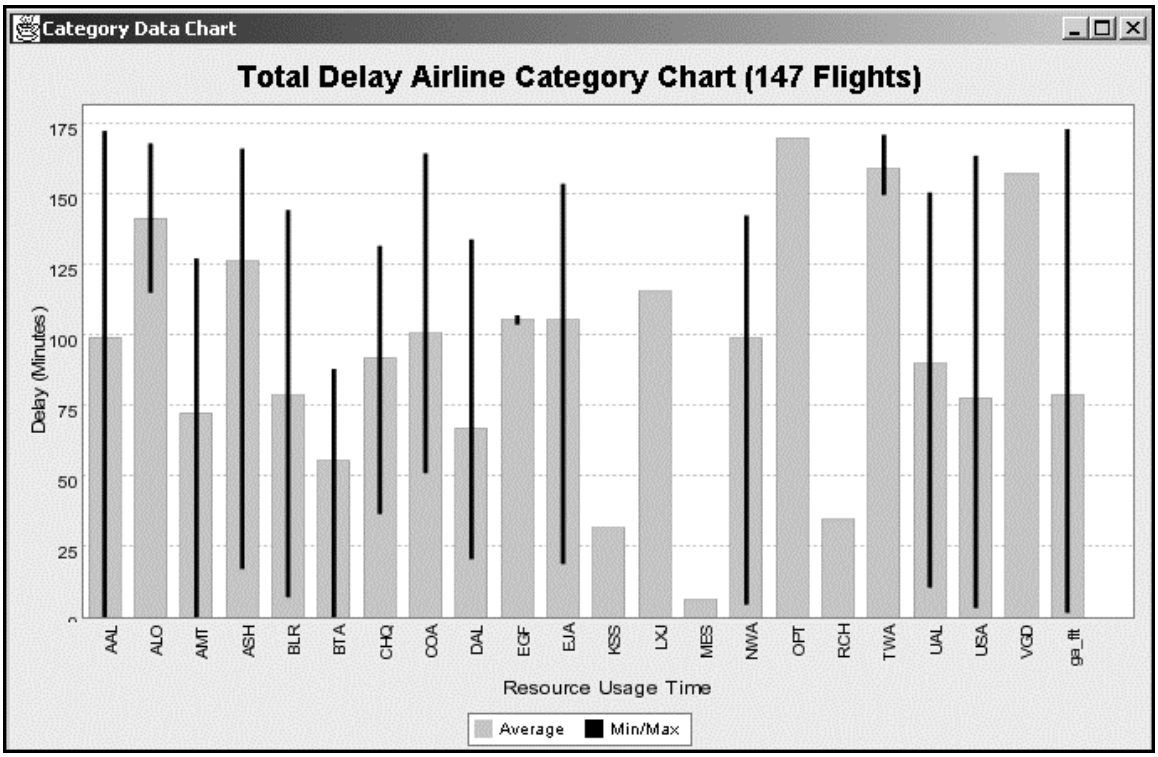

Figure B.11: Example of ERAP Category Chart

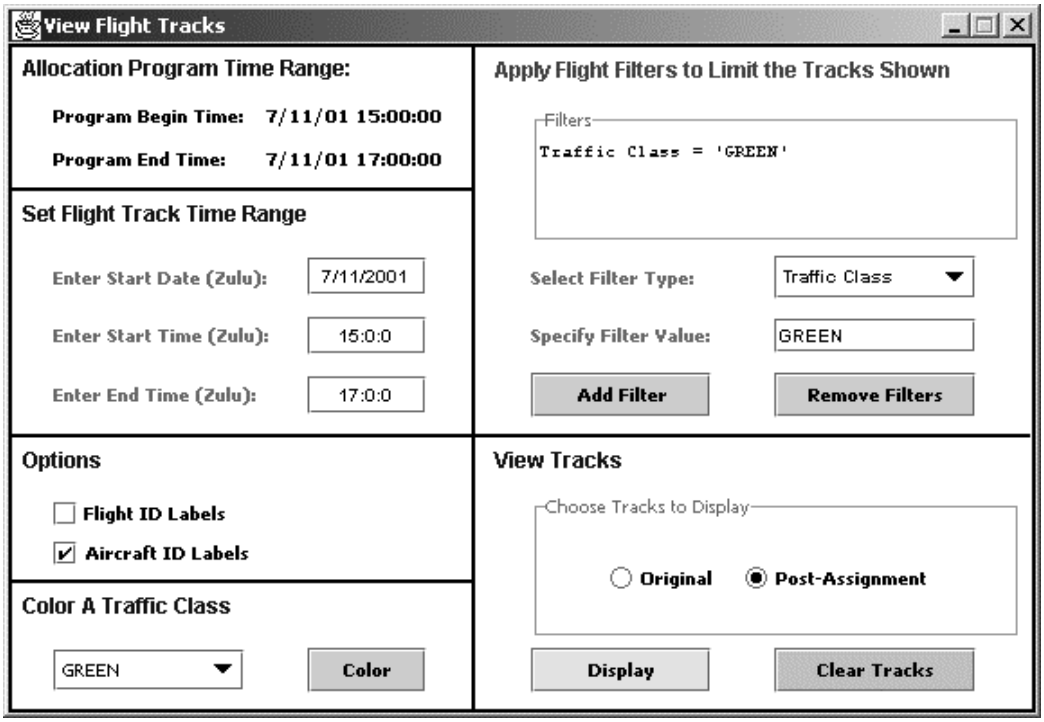

Figure B.12: ERAP GUI for Viewing Flight Tracks

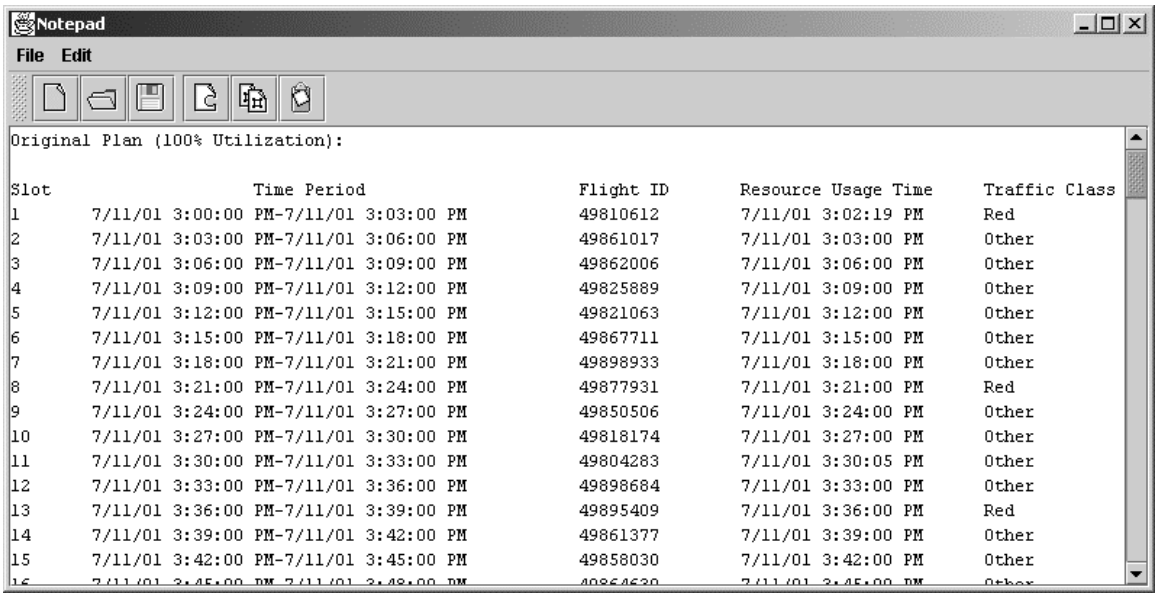

Figure B.13: ERAP GUI for Viewing Resource Utilization

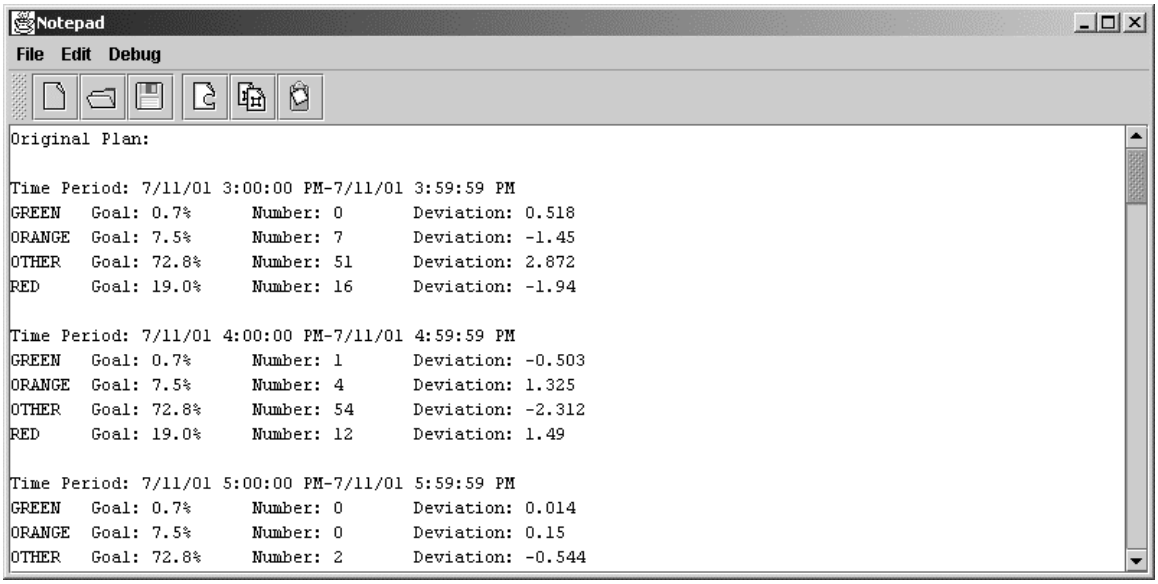

Figure B.14: ERAP GUI for Viewing Traffic Class Deviation

REFERENCES

[1] Air Transport Association, Office of Economics. (2001, August 30). *Aviation Infrastructure & Capacity*. Retrieved August 6, 2002, from the Air Transport Association Web site: http://www.airlines.org/public/industry/display1.asp?nid=1201

[2] Archut, A. (2001, June 28). Private communication.

[3] Ball, M., Hoffman, R., Hall, W. and Muharremoglu, A. (1998, August). *Collaborative Decision Making in Air Traffic Management: A Preliminary Assessment* (Technical Report RR-99-3). NEXTOR, UC Berkeley.

[4] Ball, M. O., Hoffman, R. L., Knorr, D., Wetherly, J. and Wambsganss, M. (2000). Assessing The Benefits Of Collaborative Decision Making In Air Traffic Management. In: *Proceedings of 3rd USA/Europe Air Traffic Management R&D Seminar*, 239-250.

[5] Ball, M. O., Hoffman, R. L. and Vossen, T. (2002, September). An Analysis of Resource Rationing Methods for Collaborative Decision Making. To be published in: *Proceedings of ATM 2002; Capri, September 2002*.

[6] Barnier, N., Brisset, P. and Riviere, T. (2001, December). Slot Allocation with Constraint Programming: Models and Results. In: *Proceedings of 4th USA/Europe ATM R&D Seminar*, Santa Fe, NM. Retrieved August 6, 2002, from

the Air Traffic Management R&D Seminar Web site: http://atm2001.eurocontrol.fr/finalpapers/pap119.pdf

[7] Bertsimas, D. and Patterson S. S. (2000, August). The Traffic Flow Management Rerouting Problem in Air Traffic Control: A Dynamic Network Flow Approach. *Transportation Science, 34*, 239-255.

[8] Collaborative Routing Sub-Group. (2000). *A Framework for Collaborative Routing.* Internal CDM Report. Retrieved August 6, 2002, from the Collaborative Decision Making Web site: http://vivaldi.metsci.com/cdm/cr/2002report&appendices/Appendix28aFrameworkForCollaborative.pdf

[9] Collaborative Routing Sub-Group, RBS-in-the-sky Tiger Team. (2001, November 26). *Rationing Schemes for En-route Air Traffic Management.* Internal CDM memorandum.

[10] Federal Aviation Administration, Air Traffic Control System Command Center. (2002, June 13). *National Severe Weather Playbook.* Washington, D.C.: Author. Retrieved August 6, 2002 from the Air Traffic Control System Command Center Web site: http://www.fly.faa.gov/Operations/playbook/current/current.pdf

[11] Federal Aviation Administration, Office of System Capacity. (2001, December). *2001 Aviation Capacity Enhancement Plan*. Washington, DC: Author. Retrieved August 6, 2002, from the Office of System Capacity Web site: http://www.faa.gov/ats/asc/publications/01_ACE/01_ACE.pdf

[12] Federal Aviation Administration, Traffic Flow Management Product Development Team. (n.d.). *Traffic Flow Management Development: Congestion Management*. Retrieved August 6, 2002 from the Traffic Flow & Enterprise Management Integrated Product Team Web site: http://www.faa.gov/aua/aua700/products/traffic_flow/new_tech/congestion.shtml

[13] Histon, J. M., Aigoin, G., Delahye, D., Hansman, R. J. and Puechmorel, S. (2001, December). Introducing Structural Considerations into Complexity Metrics. In: *Proceedings of 4th USA/Europe ATM R&D Seminar*, Santa Fe, NM. Retrieved August 6, 2002, from the Air Traffic Management R&D Seminar Web site: http://atm2001.eurocontrol.fr/finalpapers/pap108.pdf

[14] Hoffman R. and Ball M. (2000, July 23). *A Proposal for Resource Rationing of Enroute Resources.* Internal CDM Report to Collaborative Routing Sub-Group.

[15] Knorr, D. and Wetherly, J. (2000, February). *An Operational Assessment of Collaborative Decision Making in Air Traffic Management: "Measuring User Impacts through Performance"* (Document Control Number: R90145-05). Washington, DC: Federal Aviation Administration. Retrieved August 6, 2002 from the Free Flight Web site: http://ffp1.faa.gov/approach/media/CDM_Reports_Final.htm

[16] Sherry, J. E., Ball, C.G. and Zobell, S.M. (2001, December). Traffic Flow Management (TFM) Weather Rerouting. In: *Proceedings of 4th USA/Europe ATM R&D Seminar*, Santa Fe, NM. Retrieved August 6, 2002, from the Air Traffic

Management R&D Seminar Web site:
http://atm2001.eurocontrol.fr/finalpapers/pap11.pdf

[17] United States. Congress. House of Representatives. Subcommittee on Aviation, Committee on Transportation and Infrastructure. (1999, October 14). *The Recent Increase in Air Traffic Control Delays.* Retrieved August 6, 2002 from the U.S. House of Representatives Web site: http://www.house.gov/transportation/ aviation/hearing/ 10-14-99/10-14-99memo.html

[18] United States. Congress. Senate. Commerce Committee. (2000, June 28). *Hearing on DOT-IG's Interim Report on Airline Voluntary Customer Service Plans.* (Testimony of Carol B. Hallett). Retrieved August 6, 2002 from the U.S. Senate Web site: http://www.senate.gov/~commerce/hearings/0628hal.pdf

[19] United States Department of Transportation. (2001, July 27). *AOC 1.4.* Internal DOT Report. Retrieved August 6, 2002 from the Airline Dispatchers Federation Web site: http://www.dispatcher.org/library/AOCFinalGPO.pdf

[20] Wambsganss, M. (1996). Collaborative Decision Making through Dynamic Information Transfer. *Air Traffic Control Quarterly, 4,* 107-123.

[21] Wambsganss, M. C. (1999). Collaborative Decision Making (CDM) in Air Traffic Management. In: *Proceedings of ATM '99, Workshop on Advanced Technologies and Their Impact on Air Traffic Management in the 21st Century,* 256-266.

[22] Young, H. P. (1994). *Equity In Theory and Practice*, Princeton, NJ: Princeton University Press.