Selection of Datasets for NAS-Wide Simulation Validations

Authors:
Jimmy Krozel, Ph.D., Bob Hoffman, Ph.D., Steve Penny, and Taryn Butler

NASA Technical Monitor: Karlin Roth, Ph.D.
Metron Aviation Principal Investigator: Jimmy Krozel, Ph.D.

RTO 73 Task Team:
Taryn Butler
Goli Davidson
Bob Hoffman, Ph.D.
Jimmy Krozel, Ph.D.
Steve Penny
Anindya Roy, Ph.D.**
Travis White

October 2002

METRON AVIATION

Metron Aviation, Inc.
131 Elden St., Suite 200
Herndon, VA 20170

Contract with NASA Ames Research Center under subcontract with
TITAN Systems Corp., SRC Division.

** Prof. Anindya Roy is with the Dept. of Math and Statistics, University of Maryland at Baltimore County

TABLE OF CONTENTS

ABSTRACT

In this report, we investigate a method of selecting datasets for validating National Airspace System (NAS) simulations. First, a decomposition of the NAS is performed in order to identify the states, control actions, and performance measures, which best characterize the NAS. Data sets available for a historical statistical analysis of the NAS are reviewed. Then, a statistical analysis is performed for NAS states, control actions, and performance measures to generalize the statistics of the data over a period of 2000 to date (October, 2002). Feature vectors that describe the NAS are defined for the purpose of identifying several "types" of days in the NAS. A cluster analysis approach is invoked to identify the optimal feature vector for representing historical NAS data. Next, we demonstrate how to select the most "typical" day within each "type" of day in the NAS. Finally, we provide recommendations on how to use the "typical" days to validate simulations of the NAS.

List of Acronyms

AAR....................Airport Arrival Rate
ACARS...............Aircraft Communications Addressing and Reporting System
ACES..................Adaptation Controlled Environment System
ADR...................Airport Departure Rate
ANOVA.............Analysis of Variance
AOC...................Airline Operational Control
APREQ...............Approval Request
ARINC............Aeronautical Radio, Incorporated
ARTCC...............Air Route Traffic Control Center
ASPM.................Aviation System Performance Metrics
ASQP.................Airline Service Quality Performance
ATC....................Air Traffic Control
ATCSCC............Air Traffic Control System Command Center
ATCT..................Air Traffic Control Towers
BTS....................Bureau of Transportation Statistics
CAASD.............Center for Advanced Aviation System Development
CD&R................Conflict Detection and Resolution
CDR...................Coded Departure Route
CRS.................... Computerized Reservation System
CTA....................Controlled Time of Arrival
CTD....................Controlled Time of Departure
EDCT.................Expect Departure Clearance Time
ESP.....................En route Spacing Program
ETA....................Estimated Time of Arrival
ETE....................Estimated Time En Route
ETMS.................Enhanced Traffic Management System
FAA....................Federal Aviation Administration
FCA....................Flow Constrained Area
GA......................General Aviation
GDP...................Ground Delay Program
GS.......................Ground Stop
IFR.....................Instrument Flight Rule
INTL...................International (Airport)
LAADR.............Low Altitude Arrival and Departure Route
LIFR...................Low IFR
LOA....................Letter of Agreement
MC.....................Meteorological Condition
MIT.....................Miles In Trail
MSL...................Mean Sea Level
MVFR................Moderate VFR
NAS....................National Airspace System
NASA.................National Aeronautics and Space Administration
nmi......................Nautical Mile
NOAA................National Oceanic and Atmospheric Administration
OAG...................Official Airline Guide

OOOI.................Out, Off, On, In
OPSNET............Operations Network Database
PCA....................Principal Component Analysis
POET..................Post Operations Evaluation Tool
QA......................Quality Assurance
RVR....................Runway Visibility Range
SD......................Standard Deviation
SID.....................Standard Instrument Departure
SOP....................Standard Operating Procedures
STAR…………...Standard Terminal Arrival Route
STMP.................Special Traffic Management Program
SUA....................Special Use Airspace
TFM...................Traffic Flow Management
TMU...................Traffic Management Unit
TRACON...........Terminal Area Approach Control
US......................United States
UTC....................Coordinated Universal Time
VA......................Virginia, USA
VAMS................Virtual Airspace Modeling and Simulation
VAST…………..Virtual Airspace Simulation Technology
VFR....................Visual Flight Rules
Wx......................Weather
Z.........................Zulu Time

1 Introduction

Within NASA's Virtual Airspace Modeling and Simulation (VAMS) Project, there is a need to validate simulations of the National Airspace System (NAS). In the validation process, NAS simulation output data is compared – in some statistical sense – to the real world data characterizing the natural operation of the NAS. Within VAMS as well as other NAS simulation efforts, one must have NAS datasets available to validate low, medium, and high fidelity NAS simulations. This project specifies datasets for this purpose.

The NAS is very complex. It has a vast amount of states, controls, and performance metrics associated with it. Because of this, validation data sets have the potential to become quite large. Hence, an objective of this report is to minimize the amount of information that is required to validate a NAS simulation. To this end, we define an optimal NAS feature vector that is composed of a minimal amount of variables that describe the NAS. When used in the validation of a NAS simulation, the optimal NAS feature vector provides the minimum amount of validation data.

A useful point of departure for discussing NAS data is to consider the number of domestic enplanements as shown in **Figure 1**. Clearly, the NAS shows long term periods of growth which has been greatly disrupted by the Sept. 11, 2001 tragedy. This event complicates our analysis since the time period from Sept. 2001 to the present strays from the natural progression of the NAS. The recovery rate from Sept. 2001 to date exceeds the historical growth rate over the last decade. Because of this, we restrict our main focus to a period from Jan. 2000 to Sept. 2001. However, the most current data available is also analyzed for the sake of comparison.

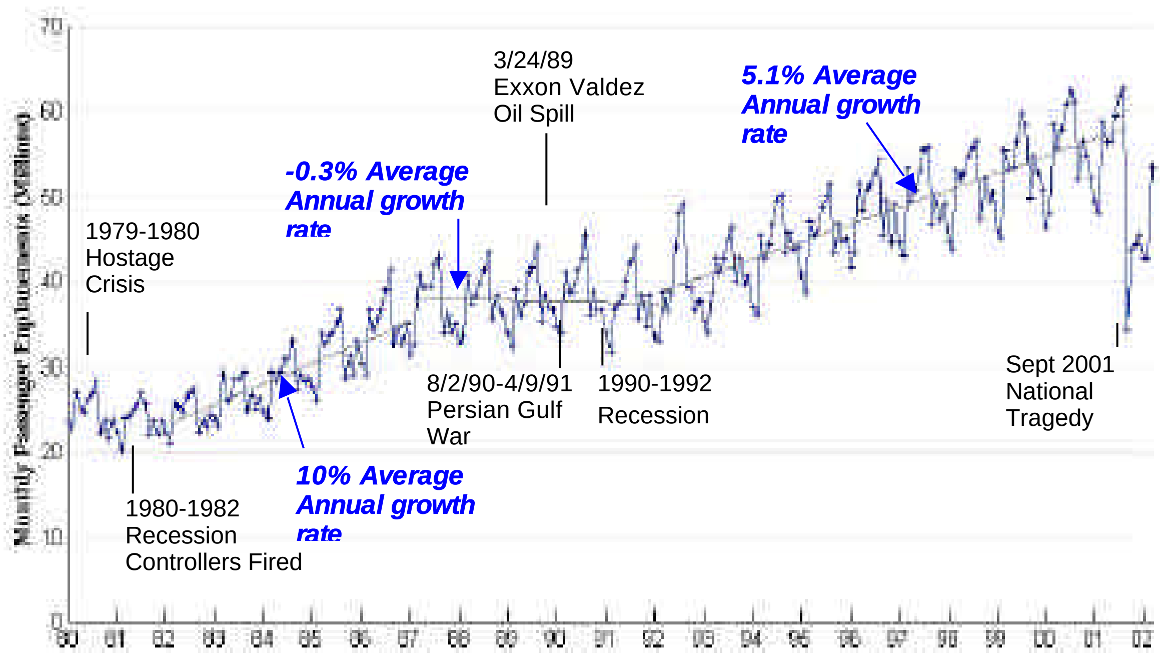

Figure 1. Historical timeline for Domestic enplanements.

Given the time period between Jan. 2000 and Sept. 2001, we perform a scientific evaluation of NAS data to determine standard "types" of days in the NAS and a "typical" day of each type. The purpose of defining the types of days in the NAS is to facilitate NAS simulation validations. By knowing the fundamental types of days in the NAS, one might validate a NAS simulation by focusing on each different type first. By spanning the set of fundamental types of days in the NAS during the validation effort, simulation builders can feel confident that a sufficient number of test days are chosen that represent the fundamental modes of operation of the NAS. By contrast, randomly choosing days or choosing days with some other criteria could fail to adequately test a NAS simulation. In this report, we identify the most typical day for each type of day in the NAS. In doing so, NAS validations can focus on validating the set of most-typical days in the NAS.

1.1 Objectives

The objectives of this analysis are:

- Define and quantify statistical properties of the NAS over a one-to-two year timeframe
- Identify standard types of days in the NAS, and the most typical day of each type
- Deliver a dataset useful for low, medium, and high fidelity NAS simulation validations

1.2 Technical Approach

The technical approach to this project is outlined below.

1.2.1 NAS Data Requirements

NAS simulation data requirements are determined for state variables, controls, and performance metrics. State variables include such things as track data, weather, demand, etc. Controls include Miles-In-Trail (MIT) restrictions, Ground Stops (GSs), Ground Delay Programs (GDPs), holding, cancellations, and the like. Performance metrics include all types of delays. The complete list is presented in **Chapter 2**.

1.2.2 Recommendation of Some Basic NAS Feature Vectors

We define a NAS feature vector, a vector that characterizes the NAS state variables, controls, and performance metrics. Each component of the vector represents one NAS statistic for a day. After considering several candidate NAS feature vectors, we proceed in our research to identify the optimal NAS feature vector – the vector with minimum size. The optimal NAS feature vector is later used to identify the varying degrees of "typical-ness” of NAS behavior.

1.2.3 Statistical Analysis of the NAS

We investigate the statistical properties of the NAS state variables, control variables, and performance metrics over time periods spanning as much as one-to-two years of data. These variables are generalized into daily statistics and grouped into sets of data spanning each year to support a subsequent cluster analysis. Where appropriate, we identify anomalies, outliers, minima, maxima, averages, standard deviations, as well as daily, weekly, monthly, or seasonal trends. Where appropriate, we identify transformations for data abstraction, cleansing, filtering, and averaging. Histogram statistics, comparison plots, and geographic plots are used to visually display the statistics for the NAS.

1.2.4 Establish an Optimal NAS Feature Vector

Based on the characterization of many potential NAS feature vectors, we analyze the statistical properties of the NAS data in order to reduce the size of the NAS feature vector to a minimum NAS feature vector size. This “optimal” NAS feature vector retains the most salient features of the historical data of the NAS. This optimal NAS feature vector is then used to identify the different types of days of the NAS.

1.2.5 Establish N Types of Days in the NAS and the most Typical Day for each Type

Utilizing the optimal NAS feature vector, we partition NAS daily statistics from Jan. 2000 to Sept. 2001 into *N* categories, each corresponding to a set of days with similar statistical behavior. The simple statistical center of each category may not actually fall near any particular day in the NAS, so we identify a distance metric that allows us to identify the closest day to the center. From any one of the *N* categories, one (or more) days can be selected as the “typical” representative of that type of day, based on the proximity to the statistical center. Finally, we describe how the most typical days for each type of day can be used in NAS simulation validations.

1.2.6 Prepare Datasets

From each of the categories of NAS behavior, representative datasets are constructed and delivered to NASA. In addition, we identify a method of storing and retrieving data from the datasets in a way that is useful to a wide variety of low, medium, and high fidelity NAS simulations. Key filters for data cleansing, smoothing, and removal of erroneous data are described to support the data retrieving process.

1.3 Report Organization

As illustrated in **Figure 2**, this technical report is organized as follows:

- Chapter 1 introduces the problem statement and approach to the problem
- Chapter 2 identifies general data requirements for NAS simulations
- Chapter 3 specifies the sources of data used in the study
- Chapter 4 performs a statistical analysis for the state, controls, and performance variables
- Chapter 5 invokes a cluster analysis approach to identify dependencies between variables and reduce the variable set to an optimal NAS feature vector size; furthermore, Chapter 5 identifies in a second cluster analysis the N-types of days in the NAS, and
- Chapter 6 investigates “special” types of days in the NAS.
- Chapter 7 states our conclusions and recommendations from this study.

Appendices contain additional information supporting this study:

- Appendix A specifies the notation used in graphs
- Appendix B specifies fleet mix data
- Appendix C illustrates the relationship between cancellations and weather
- Appendix D outlines the variable bundling process
- Appendix E outlines the type-of-day cluster analysis process
- Appendix F specifies airport abbreviations
- Appendix G specifies hub airports
- Appendix H presents the final rankings for the different types of days in the NAS

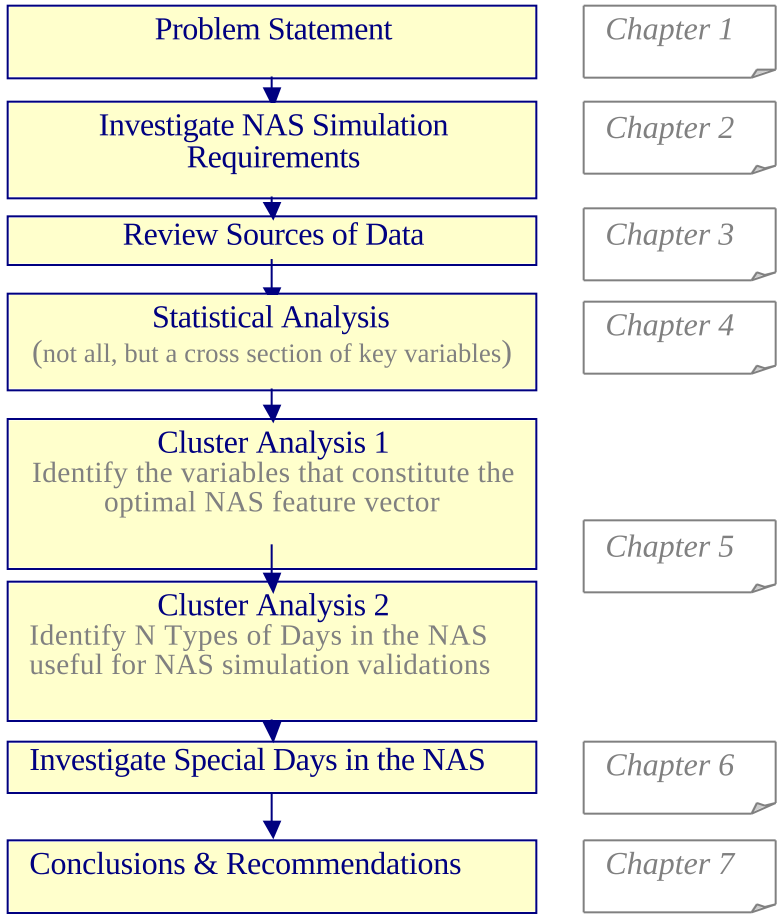

Figure 2. The organization of this report.

2 NAS Simulation Data Requirements

In this chapter, we describe data requirements for NAS simulations. NAS state variables, controls, and performance metrics are identified. We first start with a general decomposition of the NAS to enumerate variables. Then, in particular, VAMS data requirements are considered.

2.1 NAS Decomposition

In this section, we identify classifications of states and control actions that describe the NAS. We have identified six initial categories for the decomposition as presented in **Figure 3**. These categories are described next.

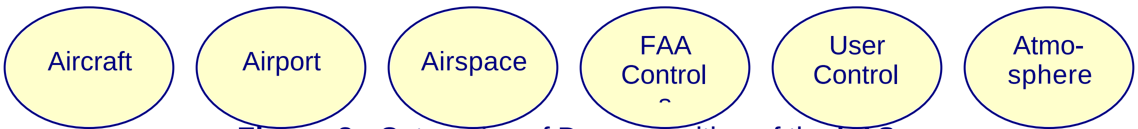

Figure 3. Categories of Decomposition of the NAS.

Note: In our decomposition, we focus on an enumeration of candidate variables. This does not imply that all these variables are necessary to investigate a macroscopic model of the NAS.

2.1.1 Aircraft States

NAS state variables that describe aircraft include, but are not limited to:

- Aircraft type
- Arrival/Departure and Alternate Airports
- Aircraft Class – (Small, Large, B757, Heavy)
- Call Sign
- Airline Operator
- Flight Plan
- Flight Plan Amendment
- Aircraft Track
- Flight Rule Category (Instrument Flight Rules (IFR) or Visual Flight Rules (VFR))

2.1.2 Airport/Surface State

Airport/surface states that influence the NAS include, but are not limited to:

- Airport Arrival Rate (AAR)
- Airport Departure Rate (ADR)
- Runway Configuration
- Airport Location (Latitude, Longitude, Mean Sea Level (MSL))
- Surface Traffic Facilities
- De-icing
- Runway Visibility Range (RVR)
- ILS Category (I, II, IIIa, IIIb, IIIc)
- Airport Category
- Terminal Forecast (LIFR, IFR, MVFR, VFR) – Low IFR, IFR, Marginal VFR, VFR
- Gate Availability

2.1.3 Airspace Infrastructure

FAA airspace infrastructure includes:

- Regions
- Centers – 20 Centers in the NAS
- Sectors – Each Center has from 20 to 80 Sectors

Figure 4 illustrates the Centers and Sectors in the NAS.

- Airways
- Fixes
- Navaids
- Standard Instrument Departures (SIDs)
- Standard Terminal Arrival Route (STARs)
- Controlled Airspace Classes (Class A, B, C, D, E, and G)
    - Class A – Positive Control Area
    - Class B – Terminal Control Areas
    - Class C – Airport Radar Service Areas
    - Class D – Control Zone/Airport Traffic Areas
    - Class E – Control Zone/Non-Towered airports
    - Class G – Uncontrolled Airspace
- Special Use Airspace (SUA)
- Flow Constrained Areas (FCAs)

In any airspace other than Class G, aircraft may be subject to Air Traffic Control (ATC). Furthermore, in Class A airspace, flights are normally operated under IFR.

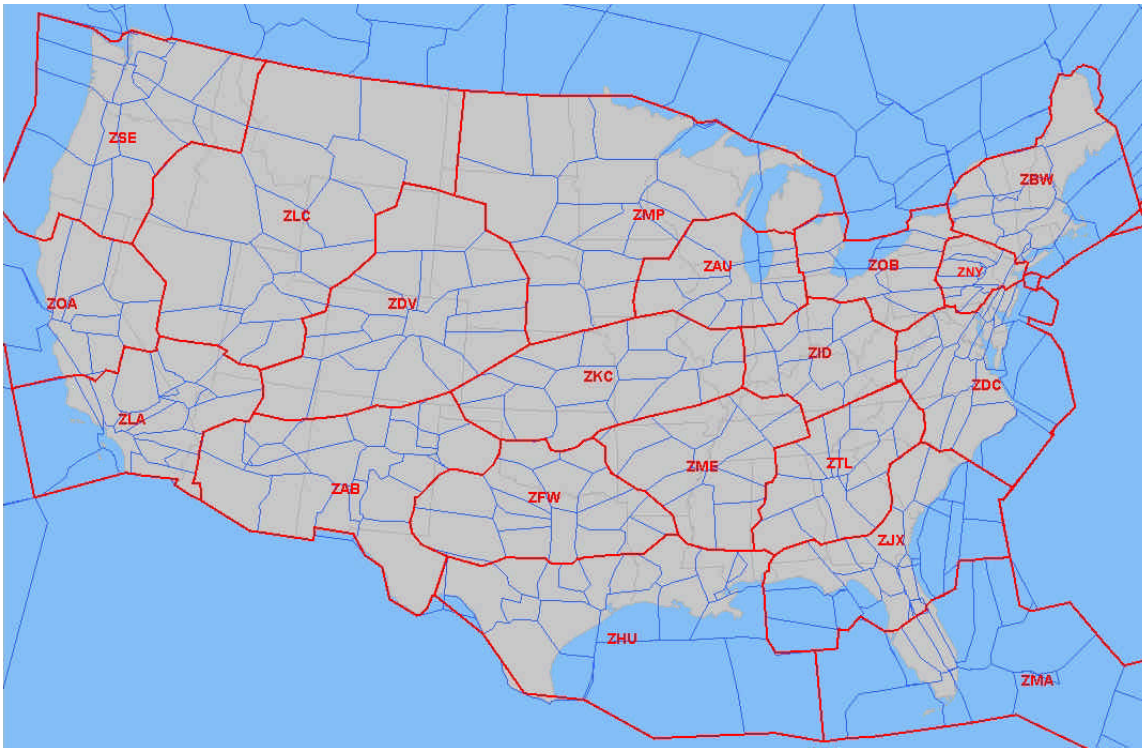

Figure 4. The centers (labeled with 3-letter identifiers) and high altitude sectors of the NAS.

2.1.4 FAA Controls

FAA control facilities include:

- Air Traffic Control System Command Center (ATCSCC)
- Air Route Traffic Control Center (ARTCCs) - (Traffic Management Unit (TMU), Local Controllers)
- Terminal Radar Approach Control (TRACONs)
- Air Traffic Control Towers (ATCTs)

FAA procedures include:

- Standard Operating Procedures (SOPs)

- Letters of Agreement (LOAs)
- Missed Approach Procedures
- En route Spacing Program (ESP)- (Center, Destination Airport)
- Special Traffic Management Programs (STMPs)

FAA control actions include:

- Ground Delay Programs (GDPs) (including Expect Departure Clearance Times (EDCTs))
- Ground Stops (GSs) – (Start time, Stop time)
- Miles-In-Trail (MIT) restrictions – (Start time, Stop time, Miles)
- Approval Requests (APREQs) – EDCT
- Direct To Route – (Direct to Fix location) for ESP
- Vector (Speed, Altitude, Heading) for ESP
- Vector (Speed, Altitude, Heading) for Conflict Detection and Resolution (CD&R)
- Vector (Speed, Altitude, Heading) for LOA compliance
- Airborne Circular Holding (Holding Pattern Fix Location)
- Airborne Path-Stretch for Holding
- Playbook Plays (Re-routes; Coded Departure Routes (CDRs))
- Low Altitude Arrival and Departure Routes (LAADRs)
- Special Use Airspace (SUA) activation and de-activation

2.1.5 User Controls/Preferences

User controls (Pilot and/or Airline Operational Control (AOC)) include:

- Flight Cancellations
- Altitude Changes for Turbulence Avoidance
- Direct To Route – (Direct to Fix location) for User Route Preference
- Vector (Speed, Altitude, Heading) for Separation Assurance
- Vector (Speed, Altitude, Heading) for Weather Avoidance
- Vector (Speed, Altitude, Heading) for Terrain Avoidance

User preferences include:

- AOC policies (crew resource management)
- AOC flight schedules
- Ramp Tower Taxi Preferences
- Military AOC Preferences

2.1.6 Atmospheric State

There are several atmospheric states that influence the NAS; they include, but are not limited to:

- Winds aloft
- Precipitation
- Convective Activity (hail, microbursts, tornados, wind-shear, turbulence, icing, lightning, and reduced visibility)
- Cloud Tops
- Ceiling
- Visibility
- Temperature
- Pressure
- Humidity
- Turbulence (Light, Moderate, Severe, Extreme)
- Icing (Rime, Clear)
- Volcanic Ash*

---

* Note: While not actually a weather phenomenon, volcanic ash is weather-related because it is transported within the atmosphere (i.e., by winds aloft).

2.1.7 Notes on States and Controls

Note that this section has identified a wide variety of states and controls that describe the NAS. In addition to these states and controls, there are also performance metrics that describe the NAS. These performance metrics are described next in **Section 2.2**. Additionally, note that the amount of states surveyed in this section of the report far exceeds the amount of data that is currently available in historical databases for describing the NAS (**Chapter 3** reviews databases that describe useful datasets for this study). Furthermore, each set of historical data spans a different start-to-finish time period. Some of these data are available up to the current time, while other datasets have a time delay between when the data are collected and when they are made available for analysis purposes. Given these data availability complications and noting the effect of the September 2001 tragedy, the most complete set of data was for a time period of January 2000 to September 2001.

2.2 NAS Data Requirements based on VAST Simulations

Virtual Airspace Simulation Technology (VAST) simulations determine data requirements [Ra02a, Ra02b] to validate the NAS simulation results with actual NAS data and statistics. A simulation may include a mixture of low, medium, and/or high fidelity models for a number of components:

1. Traffic Control
2. Aircraft
3. Airline
4. Communications
5. Navigation
6. Surveillance
7. Meteorological Condition
8. Traffic Demand
9. Airspace

Since a mixture of different fidelity models can potentially be included in a simulation, the focus on validation of a simulation is placed on the properties of the simulation states and performance metric outputs. The emphasis of validation is placed on the following performance metrics:

- Flight Event Times
- Delays
- Fuel
- Controller Workload

These performance metrics and how they are validated are discussed next.

2.2.1 Flight Event Times

The times of certain simulated flight events are recorded and used in the validation process. The flight events include:

- Gate Departure Time
- Taxi Out Time
- Take Off Time
- Airborne Time
- Landing Time
- Taxi In Time
- Gate Arrival Time
- Block Time

There are three distinct values of these times:

1. The Actual Times – The times when the events actually occurred.
2. The Measured Times – These represent measurements made of actual events.
3. Simulated Times – The times of occurrences of certain events within a simulation.

In general, data from real-world measurements establishes the measured times, and the NAS simulation determines the simulated times. The actual times are only used for discussion. In this report, these events are not recorded or analyzed explicitly. Instead, we investigate the relative information determined by delays associated with these events, and we study them in aggregate statistics only. The next section of this report explains this further in terms of delays.

2.2.2 Delays

Delay metrics important to VAST simulations include:

- Gate Departure Delay – This represents the effect of ground holds, ground stops, or other delays incurred at the gate.
- Taxi Out Delay – This represents queuing delays that might occur due to restrictions in the airport departure capabilities.
- Take Off Delay – This is the Gate Departure Delay plus the Taxi Out Delay.
- Airborne Delay – The airborne delay is the difference between the "scheduled" airborne time and the measured airborne time. The "scheduled" airborne time is taken to be the Estimated Time En Route (ETE) that is filed by the airline that is contained in the Enhanced Traffic Management System (ETMS) FZ message.
- Landing Delay – This is the take off delay plus the simulated airborne delay.
- Taxi In Delay – This represents queuing delays that might occur due to restrictions in the airport arrival capabilities.
- Gate Arrival Delay – This is the difference between the scheduled arrival time and the actual arrival time.
- Block Time Delay – Block time is the time between gate arrival and gate departure. Block time delay is the gate arrival delay minus the gate departure delay.

There are three distinct values of delays:

1. Actual Delays – these are the delays that actually occurred.
2. Measured or Estimated Delays – Delays are not directly available from ETMS but can be approximated from ETMS and OAG schedule data. Delays by individual flight cannot be obtained from ASPM. The ASPM system uses these times internally to estimate some of the delay quantities it measures, but such times are not currently available for the validation efforts.
3. Simulated Delays – These are the delays as computed within the simulation.

In this report, all of these delays are included in the statistical analysis with the exception of landing delays.

2.2.3 Fuel

In VAST NAS simulations, the fuel consumed by individual aircraft is calculated. However, the fuel-consumed metric is not validated against real-world data since the actual fuel-consumed data are not available for analysis. While BTS provides aggregate fuel cost and consumption data on a yearly basis, they do not provide it on a daily basis. In this report, we also do not include fuel consumed in the analysis due to the lack of data.

2.2.4 Controller Workload

Workload is derived from the number and types of controller actions, which include:

- Count of air/ground communications
- Vectors issued by the controller
- Reroutes by the controller

In this report, we do not investigate any of these specific controller workload variables because there is a lack of data to describe these variables. Instead, we suggest that controller workload is reflected in the terminal holding, GS, and MIT statistics investigated in this report.

2.3 The Use of NAS Feature Vectors

Based on the NAS decomposition, the NAS simulation requirements, as well as the availability of NAS data, next we describe feature vectors that can be used to describe the behavior of the NAS. A feature vector attempts to describe the state of the NAS in terms of a minimal set of state, control, and performance variables. One feature vector is established for each day in the NAS. The components of the feature vector are aggregate statistics collected for that day, as illustrated in **Figure 5**. For instance, the first component would describe weather status, the second would describe hub traffic volume, and so on.

Wx State	***NAS HUB Traffic Volume***	***NAS Traffic Mix***	***MITs***	***GSs and GDPs***	***Cancell-ations***	***Holds***	***Delays***

Figure 5. Example feature vector for the NAS.

These feature vector elements support characterizing statements about the NAS that can be either general or specific. Specific statements might take the form such as "on a typical day in the NAS, x number of the ASPM-50 airports are operating at reduced capacity, y number of MIT restrictions have been in place and there are z flight cancellations."

In order to support general statements about the NAS, we identify a small finite number of linguistic descriptors based on the statistical properties of these states, control variables, and performance metrics. These descriptors are:

- **High**, **Medium**, and **Low** are determined using the standard deviation (σ) and mean (μ) of a Gaussian (or Poisson) distribution for a state, control, or performance metric. Thus, High is above one standard deviation, Medium is within one standard deviation, and Low is below one standard deviation from the mean.
- **Normal** vs. **Extreme** is also determined using the standard deviation and mean. Normal is used to describe when the state is within two standard deviations from the mean, and Extreme is used when a variable is outside two standard deviations from the mean.
- **Rare** is determined by the maximum, minimum, standard deviation, and mean. A data point that defines the maximum or minimum is labeled as rare unless the maximum or minimum is within two standard deviations (2σ) from the mean.

Because we have collected a wide variety of NAS statistics, the potential number of components to make up the feature vector is quite large. Moreover, variables that might make up the components are often heavily correlated. In **Chapter 5**, we reduce the number of variables to a less redundant, more manageable optimal collection, which can be more easily interpreted. Cluster analysis (defined later) is used to gather the variables by similar statistical behavior. The result is M clusters. From each of the clusters, a representative variable is selected, and the others are discarded. This variable is the one with the strongest "presence" in the cluster.

Once the M dimensions of the optimal NAS feature vector are identified, one optimal feature vector is created for each day's worth of data. The entries are filled with the appropriate statistics. A second cluster analysis is performed to group the daily vectors by similar behavior. Each cluster then represents a category of NAS behavior, ranging from least typical to most typical in the category. Each day's feature vector then belongs to only one of the clusters. Details of the cluster analysis are provided in **Chapter 5**.

One of the primary objectives of this project is to deliver datasets representing typical days in the NAS. The vectors within the most "typical" cluster are ranked from 1 to k, where k is the number of vectors in the cluster. The rank of one is given to the vector closest to the statistical center of the cluster, while k is given to the vector the farthest away. (Ties are broken arbitrarily.) If three typical days worth of NAS data were collected, they would correspond to the three vectors with the highest ranking. Similarly, representative days can be selected from any of the other clusters of vectors.

3 Data Sources

In this section, we briefly describe the data sources used in the analysis and the time periods that are applicable for each dataset. These data sources are organized in alphabetical order.

3.1 ACARS OOOI Data

The Bureau of Transportation Statistics (BTS) provides Aircraft Communications Addressing and Reporting System (ACARS) data for Out, Off, On, In (OOOI) times:

- "Out" is the time when the aircraft leaves the gate. A message is automatically sent when the parking brake is released.
- "Off" is the time when the aircraft becomes airborne. Sensors on the aircraft landing gear detect when the aircraft leaves the ground, triggering an "Off" message.
- "On" is the time when the aircraft touches the runway on a landing. Sensors on the landing gear detect when the aircraft is on the ground triggering an "On" message.
- "In" is the time when the aircraft parks at the gate, doors are opened, and the parking brake is set.

ACARS OOOI data were not explicitly collected for the study. ACARS data are implicitly used in this study since BTS and ASPM data incorporate these data.

These data are published monthly on the website www.bts.gov and is generally available 40 days after the end of the reporting month. OOOI data are provided only for carriers that participate in the program[1]:

- American Airlines
- Air Canada
- Continental Airlines
- Delta Airlines
- Fedex
- Northwest Airlines
- United Airlines
- UPS
- US Air
- Alaska Airlines
- American Eagle
- America West
- Southwest Airlines

3.2 ASPM Data

The FAA provides Aviation System Performance Metrics (ASPM) data for flight metrics at 50 major airports (see **Figure 6**), referred to as the ASPM-50 airports. ASPM data were collected from January 2000, to date. ASPM provides data as described in **Table 1.** ASPM metrics for each individual flight can be correlated to other flight characteristics including:

- **Season** - data are reported on a monthly basis and in four seasonal blocks: winter, spring, summer, and fall.
- **Carrier** - data for all air carriers.
- **Time** - data by day of month and flights identified by time of departure and arrival.
- **Weather** - fields for weather conditions including ceiling, visibility, temperature, wind angle, and wind speed at both departure and arrival airports.
- **Runway Configuration** – data includes the particular combination of runways that is in use at an airport.

[1] Note: There are no General Aviation (GA) flights included in OOOI data.

ASPM integrates data from two primary sources: ETMS and OOOI data. The Official Airline Guide (OAG) planned times form the basis for comparing actual to scheduled departure times. Note that the OAG does not record information on non-scheduled flights, cargo flights, GA flights, or military operations. Also, Airline Service Quality Performance (ASQP) data are integrated into ASPM on a monthly basis. Cancellations are obtained from ETMS RZ messages, and then updated with ASQP data. Weather data are obtained from the National Oceanic and Atmospheric Administration (NOAA). Note that ASPM-reported metrics are estimates and not actual measures.

Table 1. ASPM Record fields and descriptions.

Field	Description
1	Airport ID
2,3,4	Year, Month, Day; Hour (0 to 23); Quarter Hour (1 to 4)
5,6	Scheduled Departures; Arrivals
7,8	ASPM Departures; Arrivals
9,10	Cancelled Departures; Arrivals
11,12	Count of OAG-Based Gate Delays; Percent OAG-Based on Time Gate Departs
13,14	Count of OAG-Based Airport Departure Delays; Percent OAG-Based On Time Departures
15,16	Count of OAG-Based Arrival Delays; Percent OAG-Based On Time Arrivals
17,18	Total OAG-based Gate Delay; Average OAG-Based Gate Delay
19,20	Total Taxi Out Delay; Average Taxi Out Delay
21,22	Total OAG-Based Airport Departure Delay; Average OAG-Based Airport Departure Delay
23,24	Total Airborne Delay; Average Airborne Delay
25,26	Total Taxi In Delay; Average Taxi In Delay
27,28,29	Count of Flights with Block Delay; Total Block Delay; Average Block Delay
30,31	Total OAG-Based Arrival Delay; Average OAG-Based Arrival Delay
32,33	Departure Count; Arrival Count (used for Score Card Calculation)
34,35,36	Meteorological Conditions Flag (I – Instrument); Ceiling; Visibility
37,38,39	Temperature; Wind Angle; Wind Speed
40	Airport supplied runway configuration
41,42	Number of Aircraft waiting to Depart; Number of Aircraft waiting to Arrive
43, 44	Airport supplied Departure Rate; Airport supplied Arrival Rate
45,46	Measure of Departures on Overall Score; Departure Utilization Score
47, 48	Measure of Arrivals on Overall Score; Arrival Utilization Score
49	Total Utilization Score

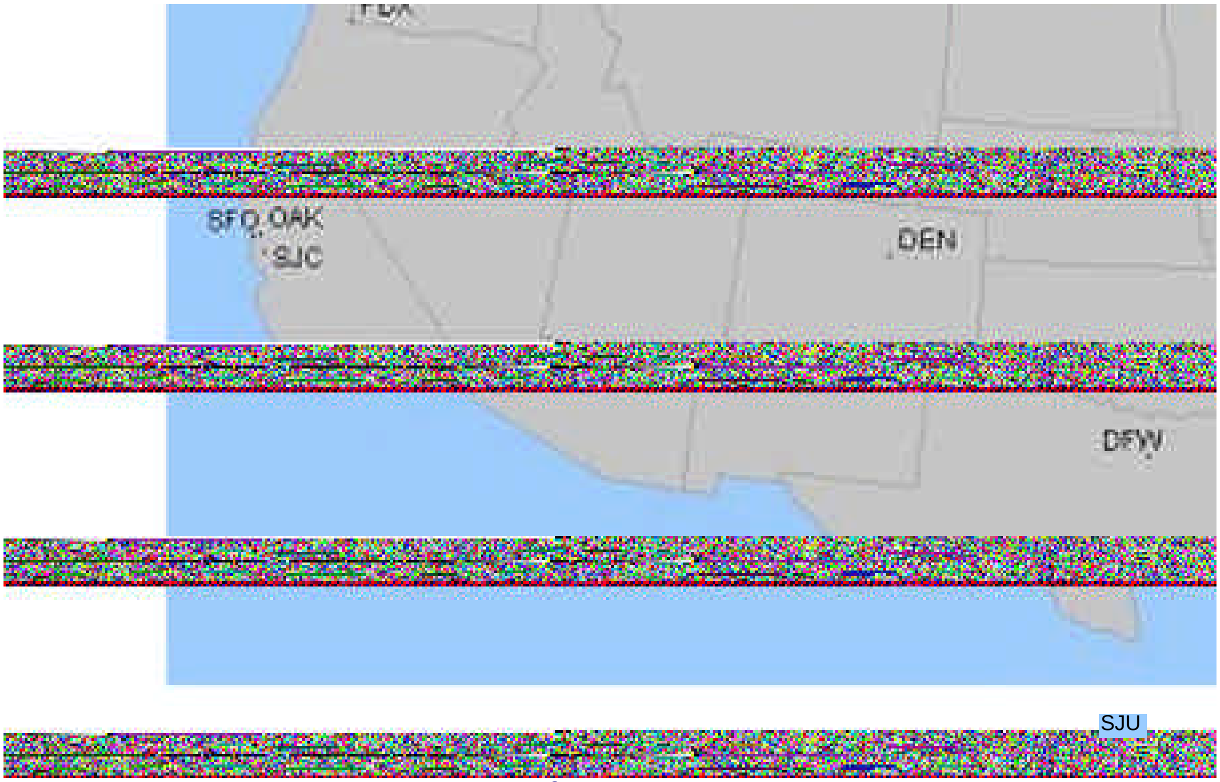

Figure 6. The 50 airports[2] used by ASPM for NAS statistics.

ASPM data posed a few problems during analysis. The ceiling, Meteorological Condition (MC), and visibility fields were not formatted for easy analysis. The ceiling data contained null records for approximately 2% of the data, and was therefore omitted. MC data was nonnumeric but numeric values were assigned for analyses purposes. IFR data was assigned a value of zero (0) and VFR was assigned a value of one (1). Visibility data contained fractional numbers. A Perl script was written that searched the file for fractional numbers and converted them to decimal numbers. Also, the fractional data was not written in the same format. Visibility data also included nonnumeric data. Nonnumeric data constituted approximately 2% of the visibility data. Definitions of the nonnumeric data could not be located; therefore they were deleted. Less than 1% of the data were outliers and approximately 2% were null records. These were also excluded from the analysis. **Table 2** summarizes the data quality issues related to ASPM data.

Table 2. ASPM Record fields and descriptions.

Type	% Ceiling Records Deleted	% MC Records Deleted	% Visibility Records Deleted
Nonnumeric	0%	0%	~2%
Outliers	0%	0%	<1%
Null records	~2%	0%	~2%

3.3 ATCSCC National Log Data

Data from ATCSCC logs were collected for the following periods:

- January 2000 to date for MIT restrictions
- April 2000 to date for GS data

[2] Appendix F specifies the mapping between 3-letter identifiers and geographical location of airports.

- September 1998 to date for GDP data

Staff position logs provide advisory information and comments on discussions/coordination with the facilities. Data are added to these files only when a facility contacts the ATCSCC to report the restriction[3]. There are different types of position logs. The Severe Weather Area Log provides a summary of activity and lists the specific events by time. East and West area logs provide (1) weather data; (2) staffing and outage information concerning centers; (3) GDP code, outage, arrival and departure runways configurations, and arrival/departure rates; (4) detailed listing of events including MITs, GSs, and GDPs; and (5) a short narrative summarizing the shift.

Sample MIT restrictions from ATCSCC Logs are as follows. MIT restrictions are formatted (for the most part) as:

> time stamp (Zulu) || Center From/To ... # MIT, issuing constraint, time periods, reason for MIT and/or additional constraints.

For example: "1640 | | ZOB/ZID...20 MIT, O/MIZAR DTW, 1700-1745, COMPACTED DEMAND. " means that at 1640 ZOB told ZID "starting at 1700 lasting until 1745 give 20 MIT for flights arriving DTW over MIZAR because of compacted demand." Inconsistencies may exist since this is not an automated process (there is no standard template). Example restrictions from 1/29/02 (miscellaneous messages have been filtered out) are shown in **Figure 7.**

```
1106 | | ZOB/ZNY...20 MIT, ORD, 1145-1230, TERM VOL. ZAU/ZOB 10/30 MIT.
1106 | | ZOB/ZBW...20 MIT, ORD, 1115-1230, TERM VOL. ZAU/ZOB 10/30 MIT.
1106 | | ZOB/ZBW...20 MIT, MDW, 1200-1300, TERM VOL.
1106 | | ZOB/ZNY...20 MIT, MDW, 1200-1300, TERM VOL.
1110 | | ZOB/ZBW...25 JETS MIT, CLE, 1230-1315, ZOB30.
1110 | | ZOB/ZDC...25 JETS MIT, CLE, 1230-1315, ZOB30.
1110 | | ZOB/ZNY...25 JETS MIT, CLE, 1230-1315, ZOB30.
1157 | | ZOB/ZBW...30 MIT, ORD, 1245-1345, TERM VOL. ZAU/ZOB 15 MIT V/PMM.
1157 | | ZOB/ZNY...30 MIT, ORD, 1245-1345, TERM VOL. ZAU/ZOB 15 MIT V/PMM.
```

Figure 7. Example MIT restrictions recorded in the ATCSCC National Logs.

A number of data quality issues exist with ATCSCC data. MIT data contained many typographic errors. This is due in part to the fact that each MIT restriction is hand-typed at the time of recording the MIT for the ATCSCC logs. A large number of obviously mistaken entries were either corrected or deleted. GDP and GS data were collected from the GDPE database. GS data are recorded in two different locations, depending on whether they were or were not part of a GDP. On many occasions, a single GS is recorded numerous times; this usually corresponds to updates or revisions of the GS. However, merely considering the first and last records at each airport per day is not sufficient, as it is possible that multiple GSs did occur at an airport. In this case, a closer study is necessary in order to determine which records correspond to which specific GSs.

3.4 BTS Data

Bureau of Transportation Statistics (BTS) data were collected for the period from 1980 to the present. BTS data provides a wide variety of NAS data for domestic and international air travel, from the major airlines:

Alaska Airlines	Northwest Airlines
America West Airlines	Southwest Airlines
American Airlines	Trans World Airlines
Continental Airlines	United Airlines

[3] When restrictions are imposed internal to a center without affecting other neighboring centers, it is not required to report such restrictions to the ATCSCC.

Delta Airlines	US Airways

These airlines generate over 90% of domestic operating revenues. Because each of these airlines earns at least 1% of the total domestic scheduled passenger revenue, FAA regulations require them to report on-time performance data to and from the 27 largest airports. Additional data are reported on a voluntary basis and airline data are uploaded from its Computerized Reservation System (CRS). In particular, this includes ACARS data.

There were no data quality issues encountered with the BTS data.

3.5 ETMS Data

ETMS data for all flights in the NAS were collected from January 2000 to the present date. ETMS records flight information for all IFR flights, including air carrier, cargo, air taxi/commuter, GA, and military operations. ETMS also has flight information for arrivals of international flights. ETMS receives the NAS state messages as shown in **Table 3**. ETMS track data (TZ messages) are currently received at one-minute intervals. Track data are presented to the truncated minute of latitude and longitude. Thus, ETMS track data has limited accuracy.

Table 3. ETMS data identifiers, descriptions, and purpose.

Identifier	Description	Purpose
FS	Scheduled Flight Plan	Scheduled flight plan ahead of the filed flight plan
RS	Scheduled Flight Cancellation	Cancels a scheduled flight previously fed into ETMS
FZ	Flight Plan	Flight Plan as filed with the NAS
AF	Flight Plan Amendment	Amendment to flight plan as filed with the NAS
RZ	Cancellation	Cancels a flight plan previously filed with the NAS
DZ	Departure	Signifies the activation of a proposed flight
UZ	Center Boundary Crossing	Current flight plan data as sent from ARTCC from which flight is leaving to the ARTCC which the flight is entering
TZ	Position Update	Current position, altitude, and speed of a flight as tracked by the NAS
AZ	Arrival	Signifies the termination of an active flight

There are several ETMS data quality issues to note. When ARTCC computers are particularly busy they may temporarily cease to generate some ETMS messages. Furthermore, it is common to find ETMS DZ messages missing, for which, an ETMS TZ message must be used to infer the take off time of an aircraft. ETMS AZ messages are also missing in archived data sets, for which ETMS TZ messages may also be used to infer the landing time of an aircraft.

3.6 OPSNET Data

Air Traffic Operations Network (OPSNET) data for all flights in the NAS was collected for January 1, 2000 to the present date. OPSNET data are described in comparison to the ASPM data described in a previous section. Importance will be placed on the differences between the two databases.

The FAA maintains data on air traffic activity at ARTCC and preliminary airport traffic counts, instrument operations and instrument approaches, as well as delays in the OPSNET database. Air traffic activity is the total of the number of operations at FAA and contractor controlled airports, instrument operations at FAA and contractor controlled airports, and aircraft handled at ARTCCs. A flight is under FAA control from the time the aircraft leaves the departure gate to when the flight arrives at the arrival gate. All OPSNET data are aggregate and not flight-specific.

Flight delay is separated into reportable, non-reportable, and international delay. Reported delays are measured with a 15-minute threshold, that is, when the elapsed flight time exceeds the flight plan times filed with the FAA by 15 minutes. Delays are recorded for the time the aircraft is at the gate, on a taxiway, or holding en route. Non-reportable delay is not recorded. Non-reportable delays are delays caused by pilot initiated en route deviations around adverse weather, cancelled flights, and delays due to aircraft mechanical problems. Also, taxi times spent under airport or airline ramp control (which are non-FAA entities) are also not part of OPSNET data. International delays are due to delay caused by initiatives imposed by facilities outside the United States (US). International delays are recorded in the OPSNET database and are not separated.

Delays reported in OPSNET are also categorized as terminal or en route delays. Terminal delays are incurred as a result of conditions at the departure or arrival airport. En route delays occur when aircraft incur airborne delays of 15 minutes or more as a result of an initiative imposed by a facility to manage traffic. En route delay statistics do not distinguish between circular airborne holding and other types of airborne holding (e.g., path stretching). OPSNET produces reports of delays by category, class, and cause; it also reports ground delays as ground stops and EDCTs (GDPs). The total number of ground delays, total number of delay minutes, and the average amount of delay minutes are reported. In contrast, ASPM does not provide causal delay data; it only provides the overall total number of delays.

Unlike ASPM, OPSNET allows queries by tower, instrument, or center, in addition to airport searches. ASPM may be queried for airports only. OPSNET data may be gathered based on the location of control towers: airport, state, or region. This data includes the type and number of operations for each area. A specific facility may be chosen or a summary of all tower operations is also available. Filtering is available based on period, comparison or ranking, and FAA or contract operations. When looking at total operations by state, the data can be filtered on the individual towers located in the state. Only monthly and yearly reports are available for regional data, whereas daily data are available for the other capabilities. ASPM provides delay data down to the hour and quarter-hour. OPSNET is much more coarse; it simply gives the total for the day.

Several query filters are available in OPSNET. Instrument operations in OPSNET are divided into primary, secondary, and over flight categories. Filtering by center is available also. This can be done by choosing individual centers or by choosing a state or region and receiving data for all centers in that state or region. Summary reports are available for all centers also. The reports include the total number of aircraft handled by each center or in each state or region and the number of domestic and oceanic aircraft, if selected.

In comparison, ASPM uses more data sources than OPSNET and has a more stringent data quality control. Recall that OPSNET only records IFR flights, while ASPM records IFR and VFR. Data in ASPM may be changed or updated at the last minute. This allows for up-to-date information. OPSNET data may not be changed or updated after a certain number of days. The data are frozen after the cutoff point, which can leave the data incomplete or incorrect in some cases. Therefore, total counts from each database may not be equivalent when comparing the two. **Table 4** provides a summary of the data available through OPSNET with a comparison to ASPM.

Table 4. Comparison of OPSNET and ASPM.

	OPSNET Capabilities	ASPM Capabilities
Query Types	Delays, Towers, Instrument, Centers	Delay Weather analyses, Taxi times, Individual flight information, Cancelled flight information
Data Sources	Airports, ARTCC	ETMS, ARINC OOOI data, OAG, ASQP
Airports	All airports, centers, towers, regions	50 airports (ASPM 50)
Time	Quarter Hour, Hourly, Daily, Monthly, Yearly	Quarter Hour, Hourly, Daily, Monthly, Yearly
Weather	General – Delays were either caused by weather or they were not.	Ceiling, Visibility, Temperature, Wind (Angle and Speed)
Delay	Reported by category, class, and cause	General – Flights were delayed or not
Other	Rankings and comparisons are available.	Rankings and comparisons are available.

4 Statistical Analysis of the NAS

The analysis of aggregate statistical properties of the NAS is presented in this chapter. We establish the aggregate properties of the NAS state variables, control actions, and performance measures by investigating the statistical properties of these data over time periods spanning as much as one-to-two years. When possible, each variable is investigated from Jan. 2000 to the present date. This chapter is meant to be a survey of variables that describe the NAS; not all the variables that are investigated later in **Chapter 5** are reviewed here. Some of the variables in this chapter are also left out of the subsequent analysis in **Chapter 5** due to the lack of a full data set. Notes are provided in the sections that cover such variables.

Note: **Appendix A** describes the notation for the majority of the plots in this chapter.

4.1 Statistics of NAS State Variables

Next, we establish the aggregate properties of the following NAS state variables:

- Passenger Enplanements
  - Yearly Trends
  - Weekly Trends
- Scheduled Arrivals
- Scheduled Departures
- Actual Arrivals
- Actual Departures
- Total Operations
- Average Airport Arrival Rate
- IFR Traffic
- VFR Traffic
- Airport Ceiling Condition
- Airport Visibility Condition
- Changes in Runway Configuration

These NAS states span the passenger demand, airport states, and weather states.

4.1.1 Seasonal Trends of Enplanements

BTS provides statistics of the monthly total number of enplanements for both domestic and international air travel. We investigate domestic travel only. Domestic revenue passenger enplanements record the total number of passengers boarding an aircraft in the NAS. These data are reported monthly for the years 1980-2002 in **Figure 8**.

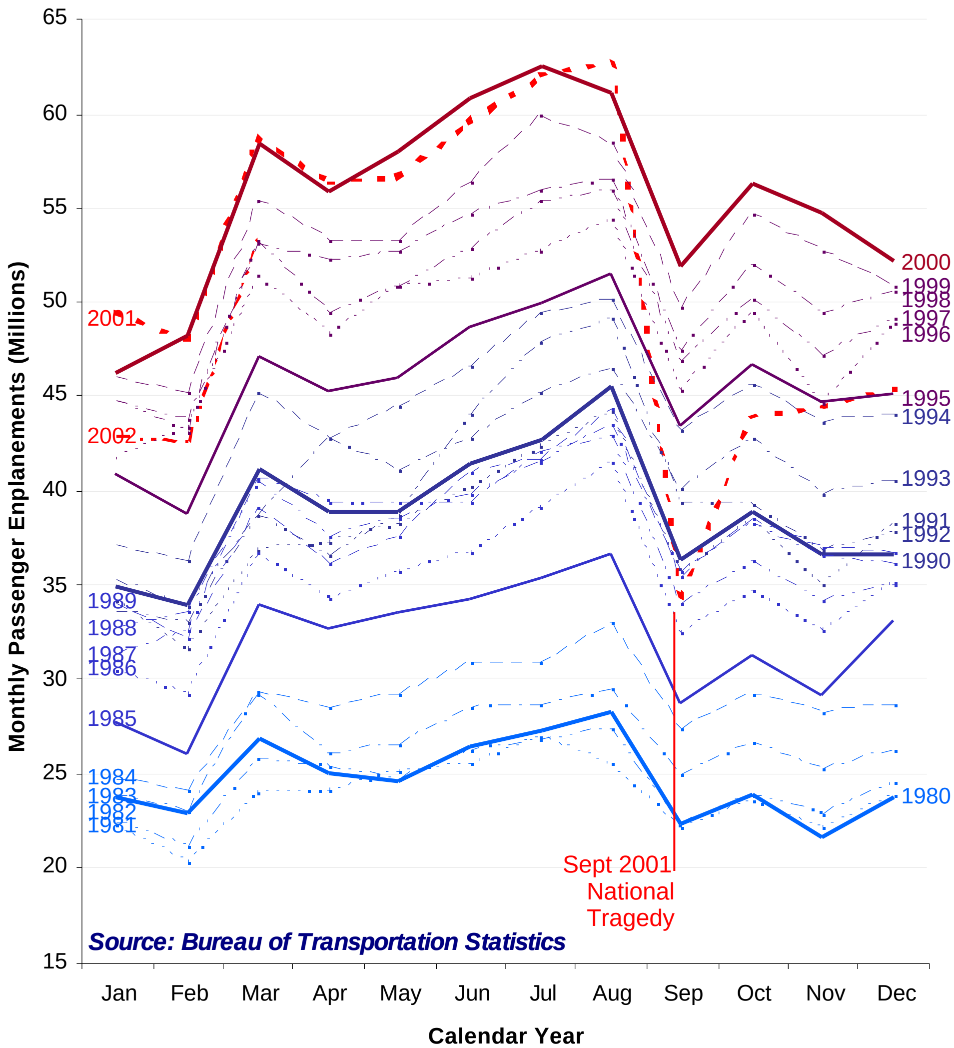

Figure 8. Comparison of yearly data for Domestic enplanements identifies seasonal trends.

Enplanement data provides evidence of the seasonal trend in demand. BTS provides the statistics of the monthly total number of domestic enplanements for this analysis. The BTS notes that the traffic data are reported to the BTS by Large Certificated Air Carriers to include carrier groups: Majors, Nationals, Large Regionals, and Medium Regionals. Traffic statistics for Small Certificated Air Carriers and Commuter Air Carriers are not included. For further information about fleet mix statistics, see **Appendix B**.

4.1.2 Weekly Trends

It is common knowledge within the air traffic community that air travel is heavier during the week than on weekends. This is easily verified by collecting departure and arrival counts by day of week

(Monday, Tuesday, Wednesday, etc.). **Figure 9** and **Figure 10** show that average departure and arrival counts are significantly less on Saturdays and Sundays than on weekdays. Moreover, traffic volume is fairly uniform across the days of the week.

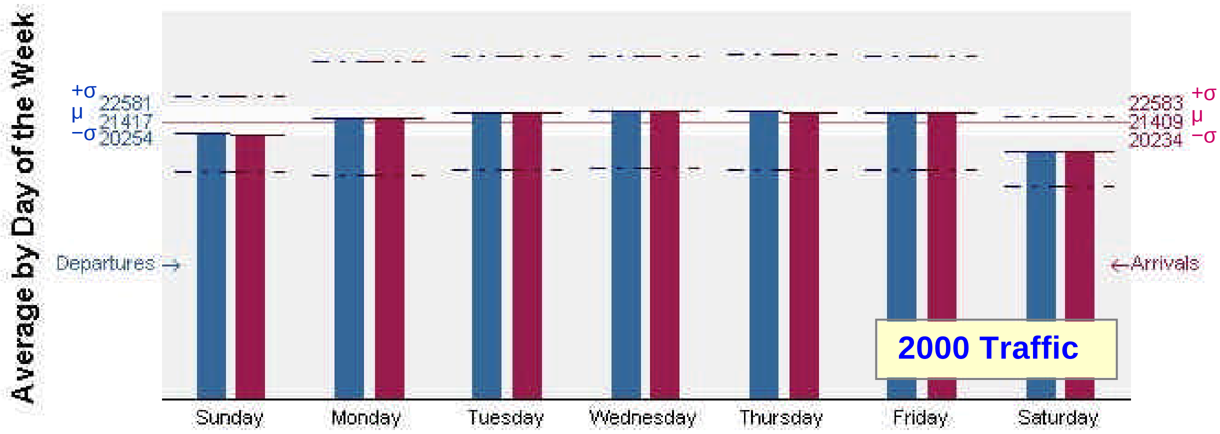

Figure 9. Comparison of traffic based on day of week for 2000.

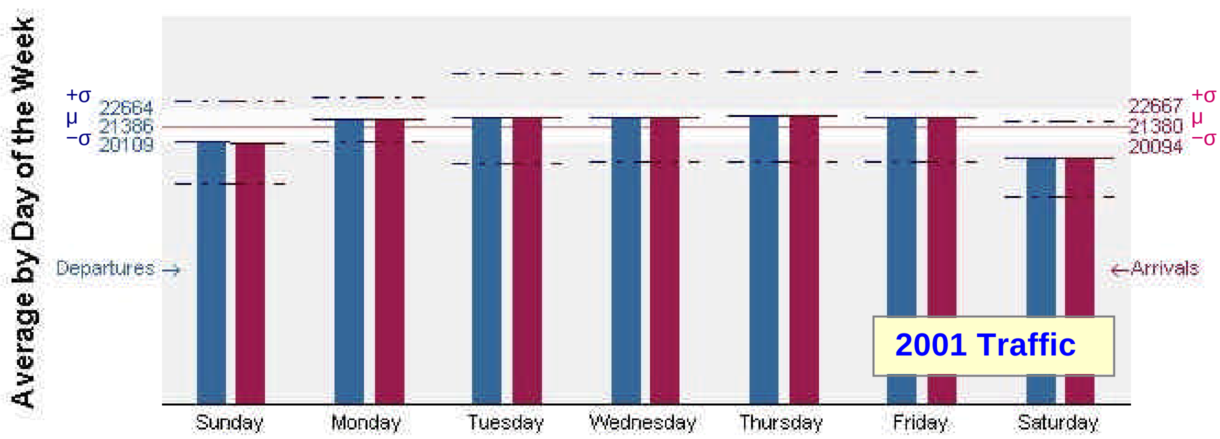

Figure 10. Comparison of traffic based on day of week for 2001.

Weekly trends are further illuminated by **Figure 11** through **Figure 14**. **Figure 11** shows scheduled arrivals over time for the year 2000. This demonstrates that scheduled traffic volumes tend to hover around three values: Weekday, Saturday, and Sunday. For this reason, weekend and weekday statistics generally separate into separate clusters.

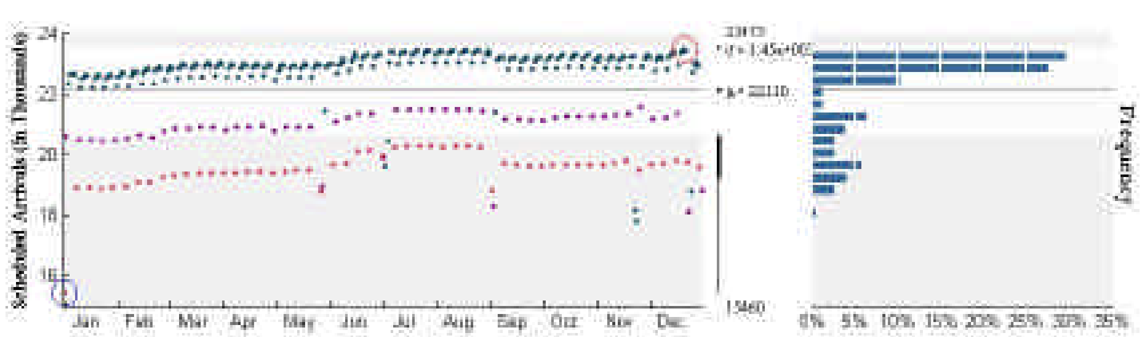

Figure 11. Total traffic for scheduled arrivals in 2000.

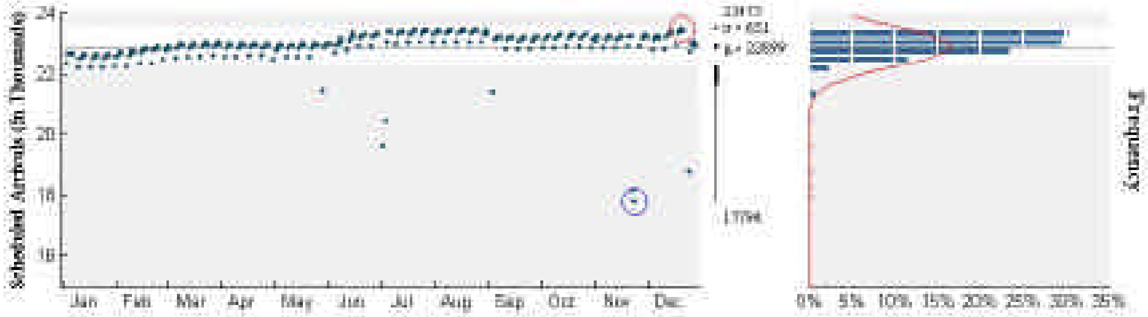

Figure 12. Business day traffic for scheduled arrivals in 2000.

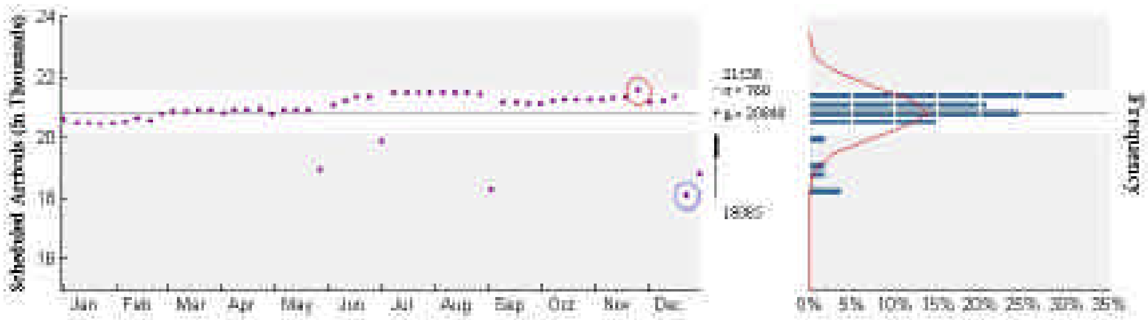

Figure 13. Sunday traffic for scheduled arrivals in 2000.

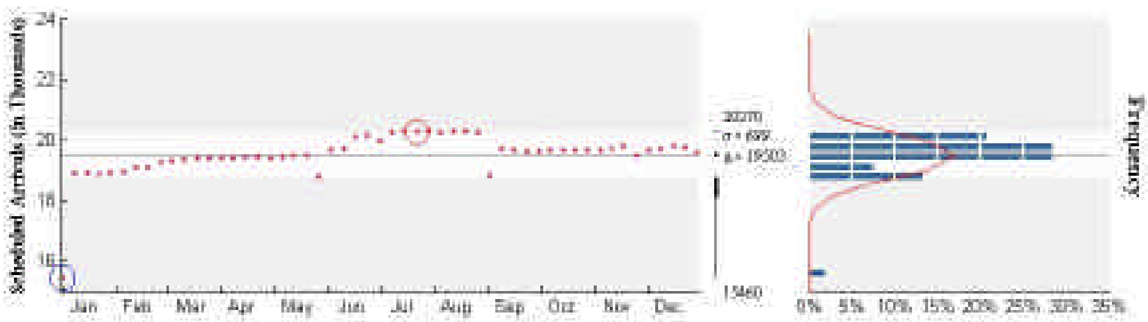

Figure 14. Saturday traffic for scheduled arrivals in 2000.

4.1.3 Scheduled Arrivals and Departures

Scheduled arrivals and departures are defined by OAG data from ASPM. Arrival and Departure data generally demonstrate the tri-modal distribution of Saturday, Sunday, and Weekday traffic, as illustrated in **Figure 15** through **Figure 22**.

2000 OAG Arrivals/Departures

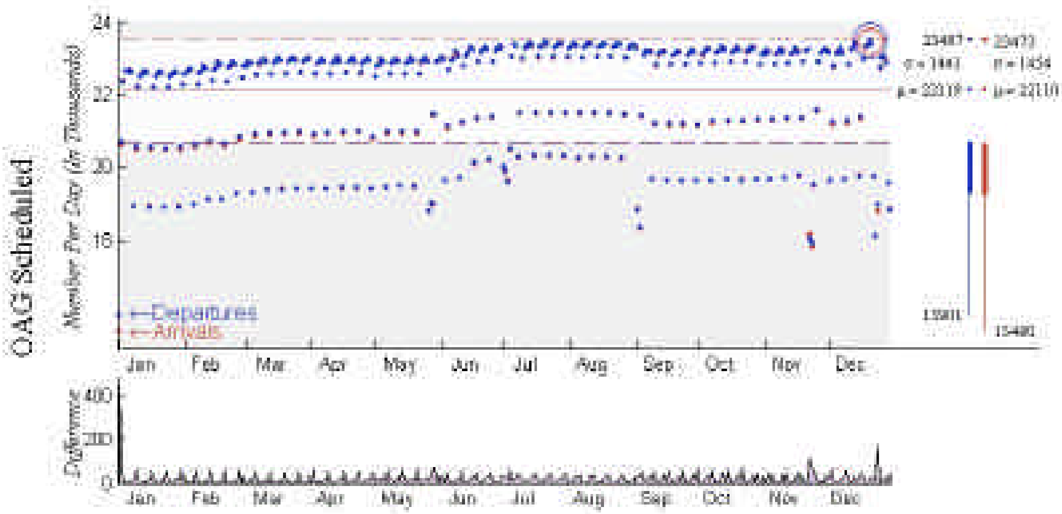

Figure 15. Scheduled arrivals / departures for domestic flights in 2000.

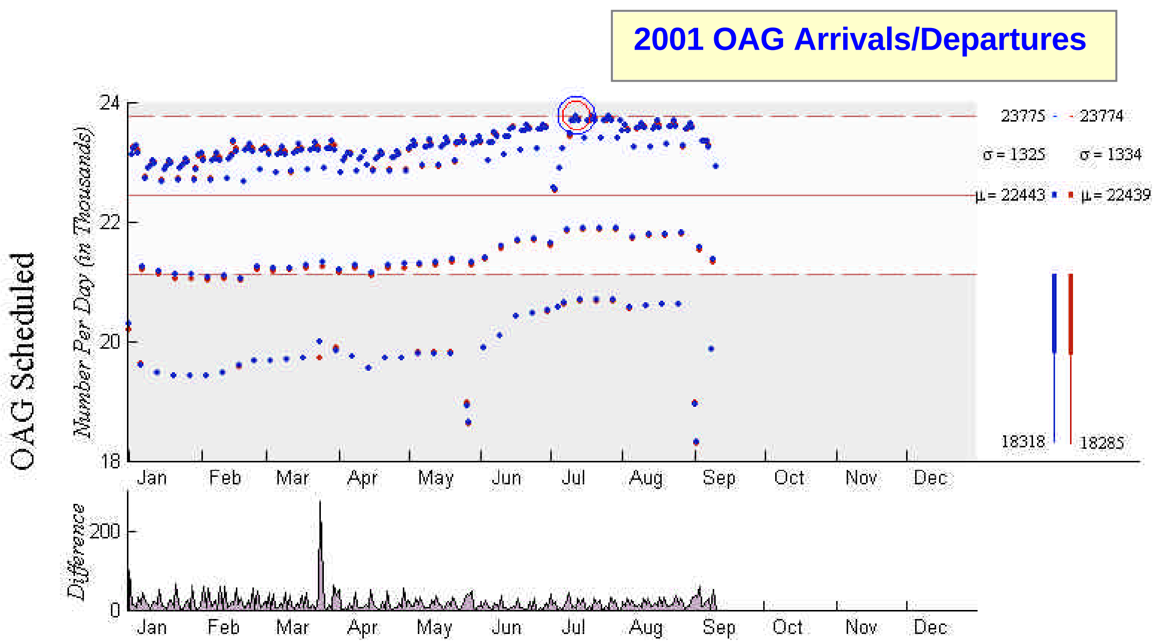

Figure 16. Scheduled arrivals / departures for domestic flights up to Sept 2001.

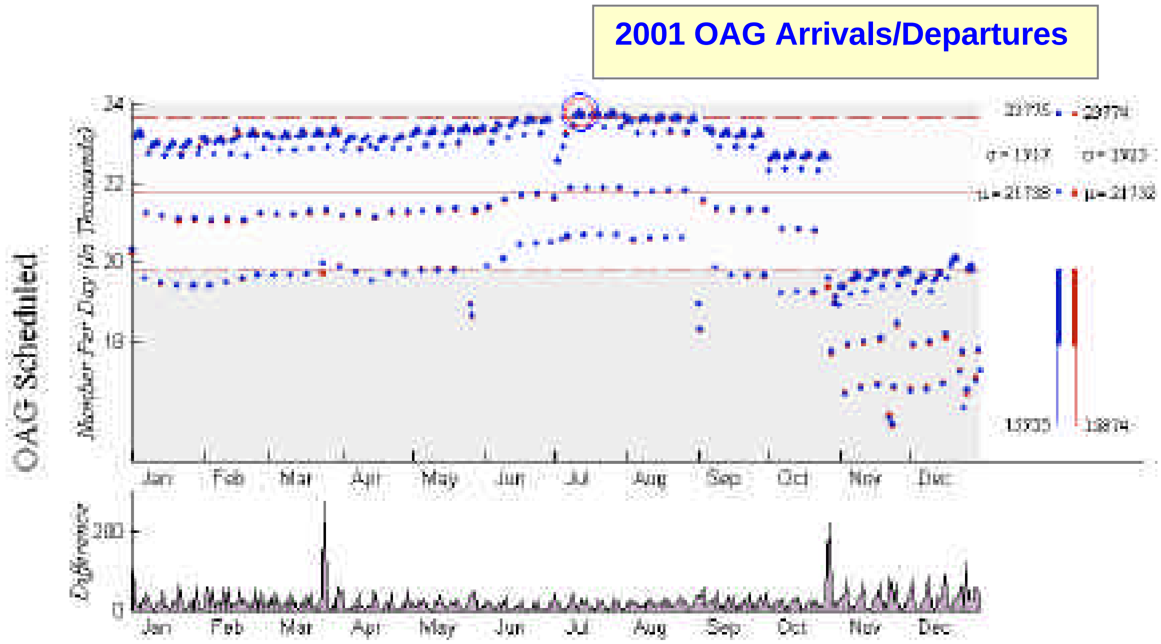

Figure 17. Scheduled arrivals / departures for domestic flights in 2001.

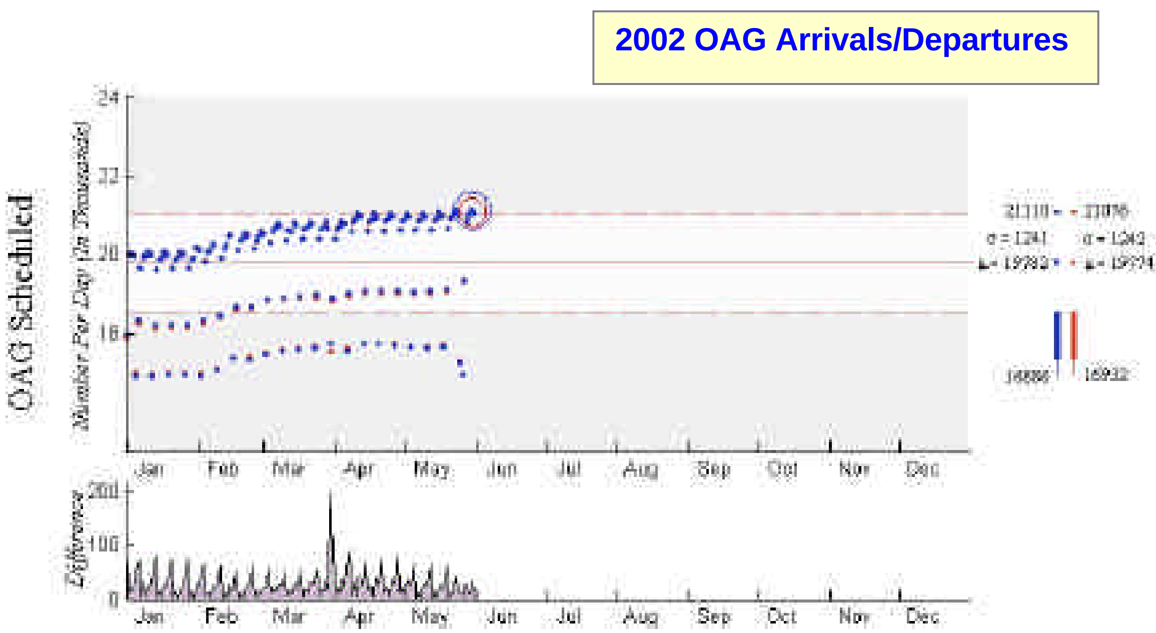

Figure 18. Scheduled arrivals / departures for domestic flights in 2002.

The difference in arrivals and departures can be primarily attributed to flights departing from ASPM airports with non-ASPM airport destinations and arrivals at ASPM airports from non-ASPM airport origins.

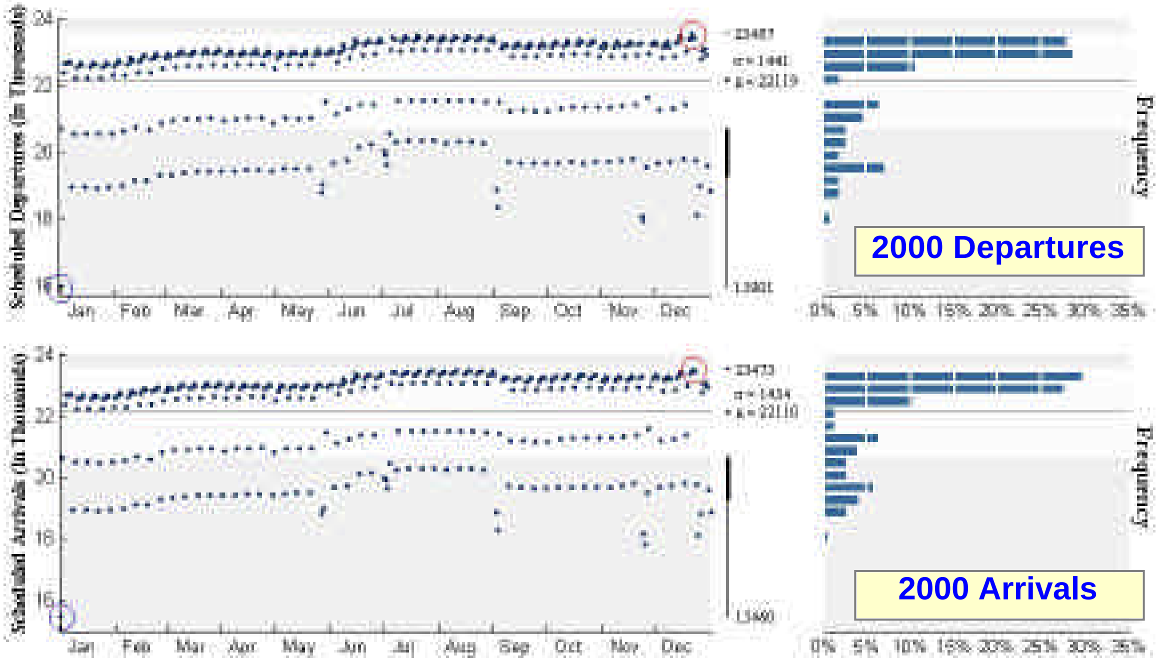

Figure 19. Distributions for scheduled arrivals/departures for domestic flights in 2000.

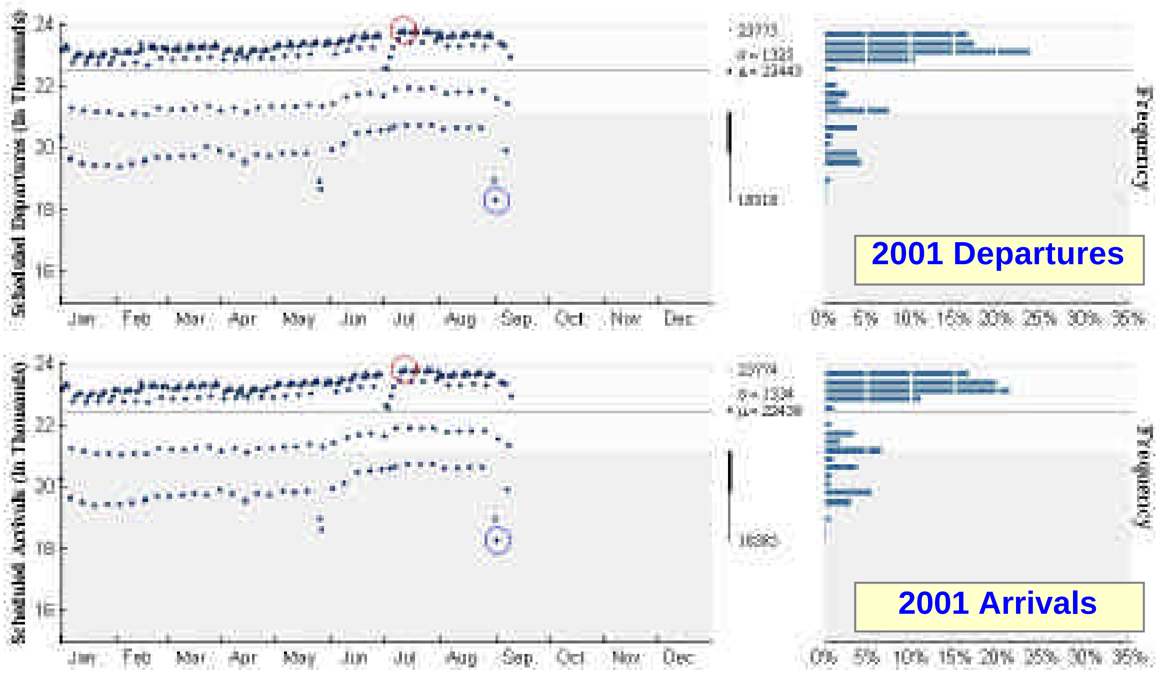

Figure 20. Distributions for scheduled arrivals/departures for domestic flights up to Sept., 2001.

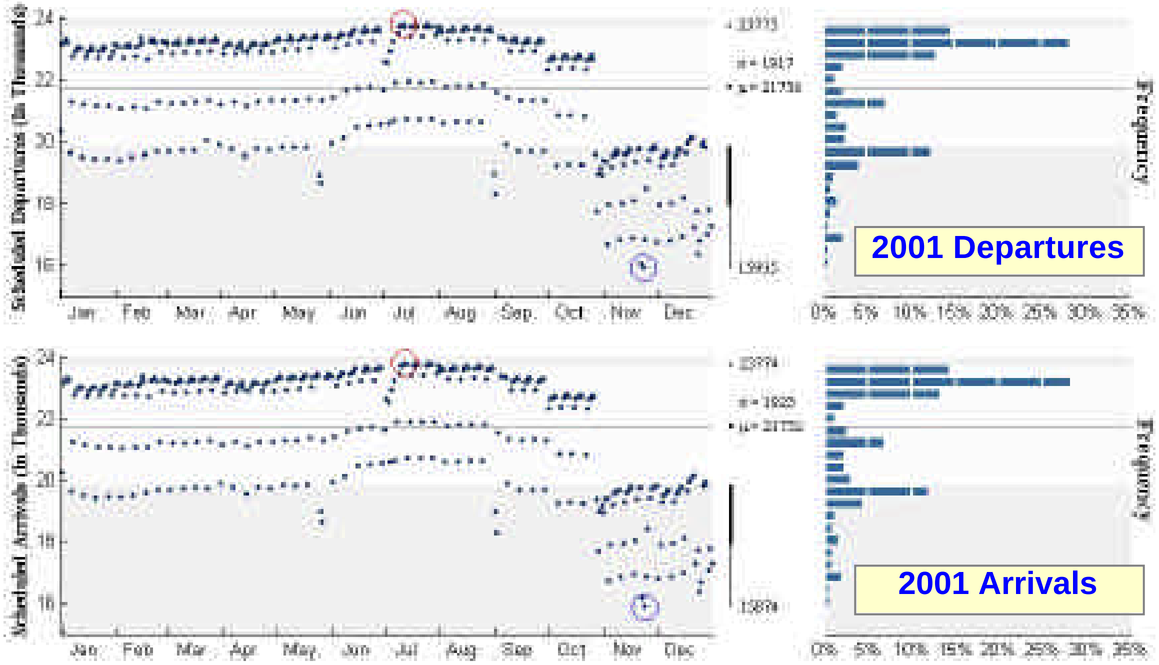

Figure 21. Distributions for scheduled arrivals/departures for domestic flights for 2001.

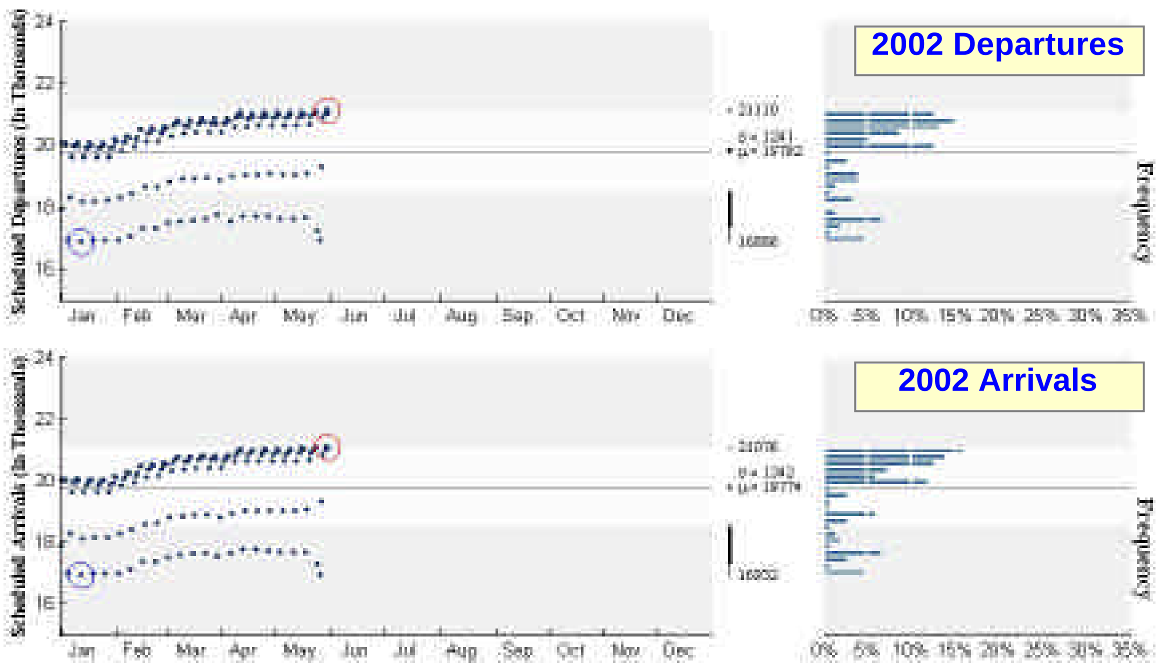

Figure 22. Distributions for scheduled arrivals/departures for domestic flights for 2002.

4.1.4 Actual Arrivals and Departures

Actual arrivals and departures are defined by OPSNET data. Arrival and Departure data generally demonstrate the tri-modal distribution of Saturday, Sunday, and Weekday traffic, as illustrated in **Figure 23** through **Figure 26**.

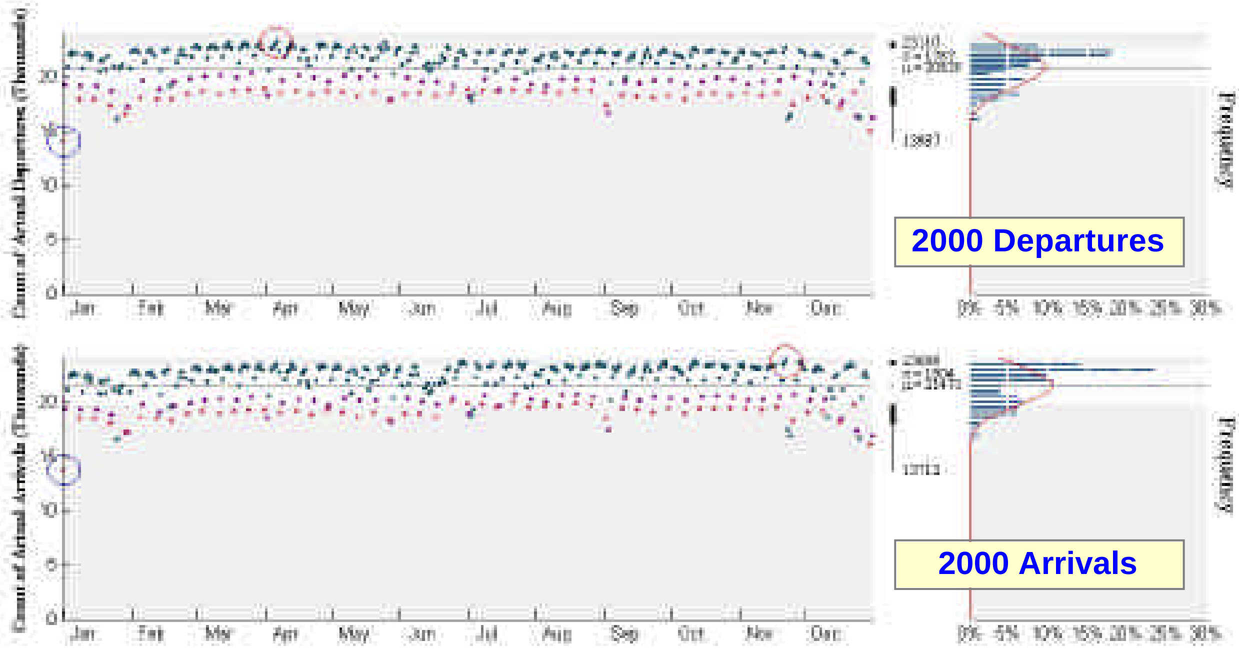

Figure 23. Distributions for actual arrivals/departures for domestic flights in 2000.

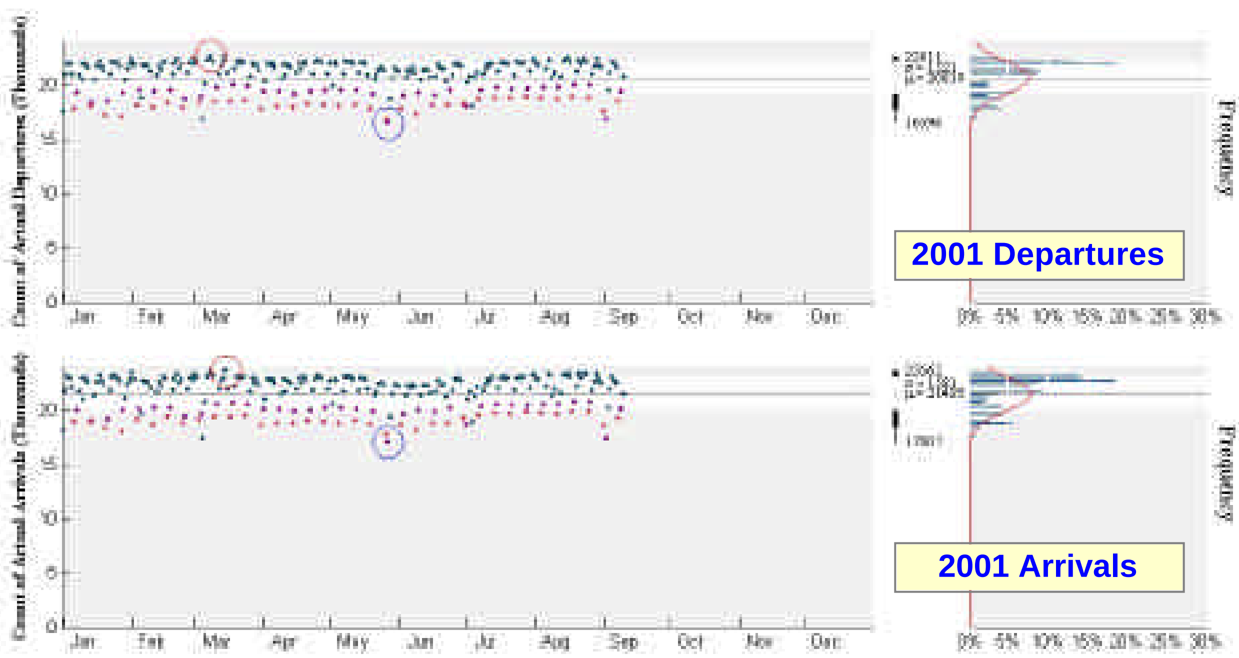

Figure 24. Distributions for actual arrivals/departures for domestic flights up to Sept., 2001.

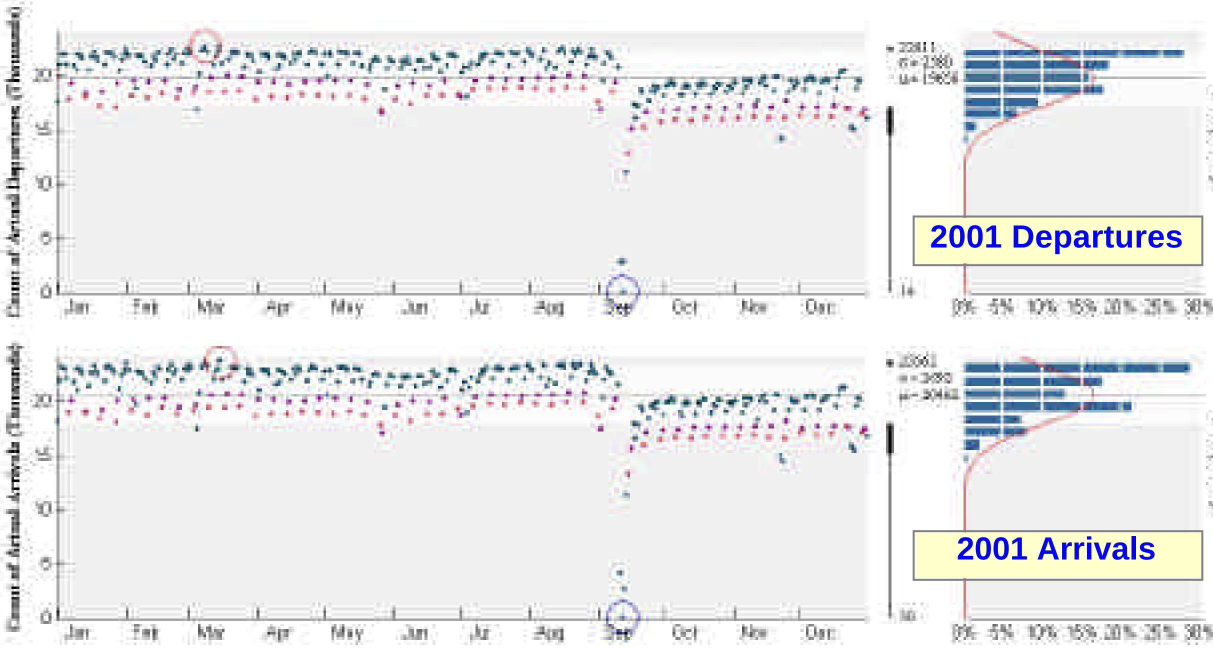

Figure 25. Distributions for actual arrivals/departures for domestic flights for 2001.

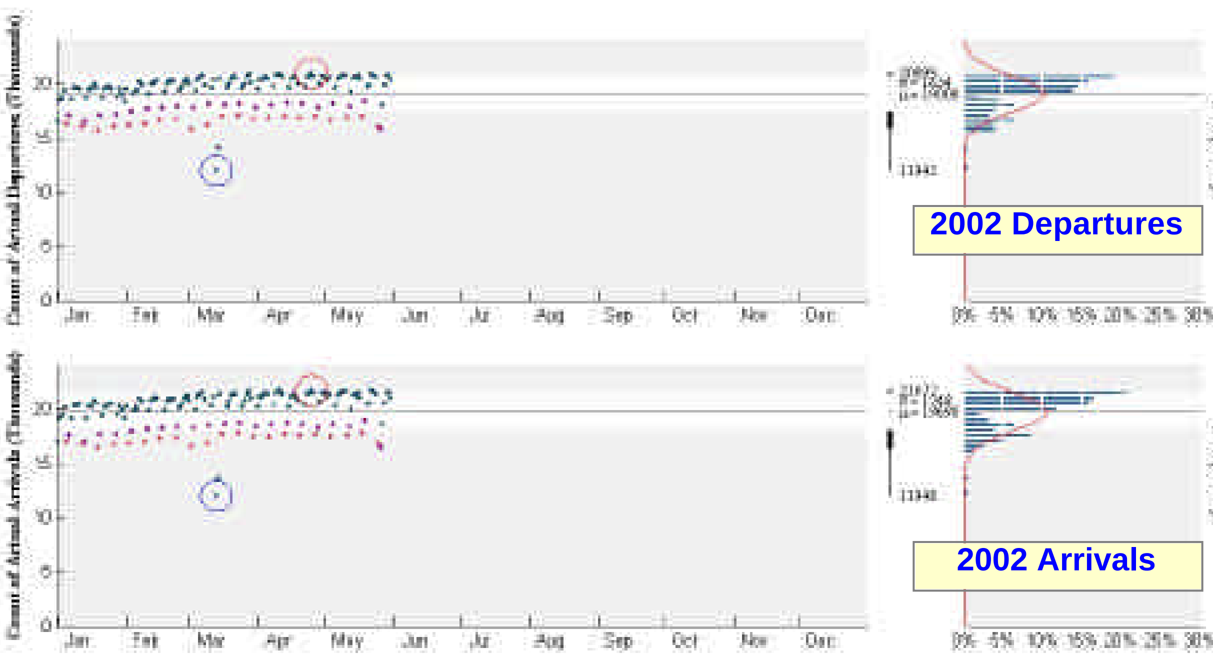

Figure 26. Distributions for actual arrivals/departures for domestic flights for 2002.

4.1.5 Total Operations

Total Operations are reported from OPSNET data. These data are described in **Figure 27** through **Figure 33**.

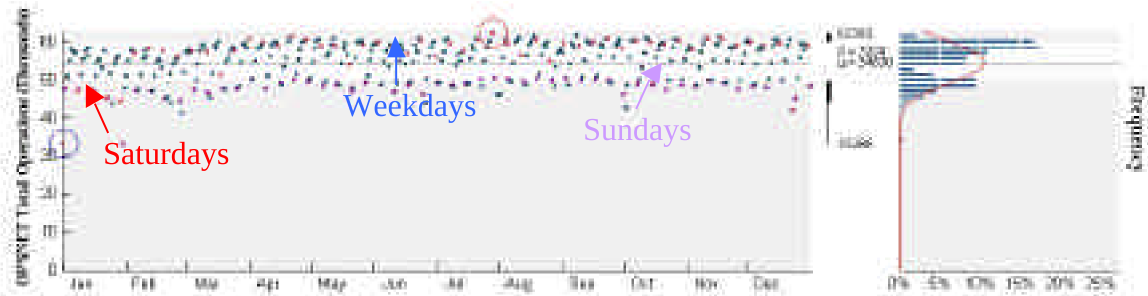

Figure 27. Total Operations in 2000.

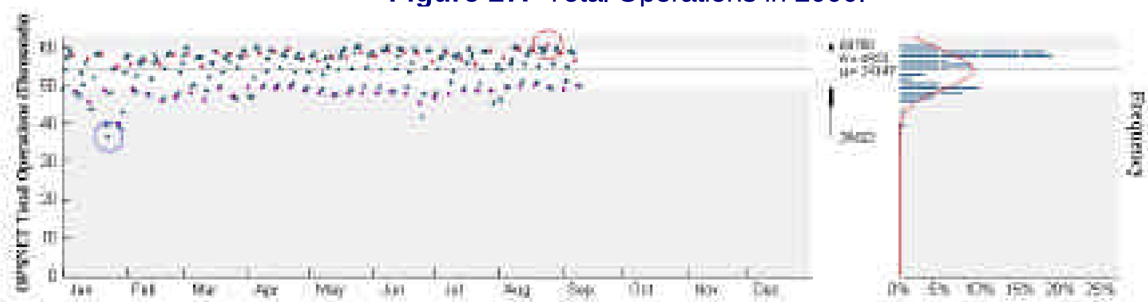

Figure 28. Total Operations in 2001 up to September 10.

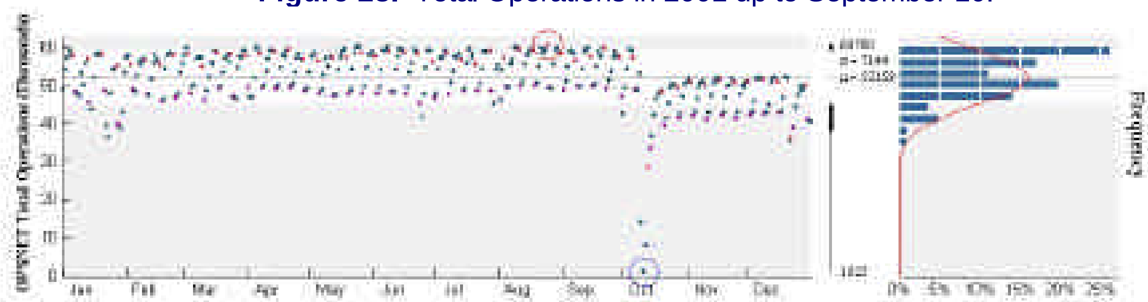

Figure 29. Total Operations in 2001.

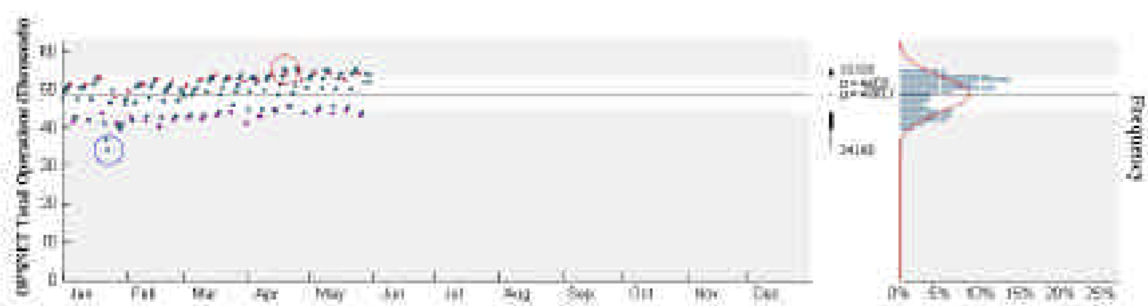

Figure 30. Total Operations in 2002.

4.1.6 Average Airport Arrival Rates (AARs)

AARs are determined from ASPM data. (See **Figure 31** through **Figure 34**). These plots indicate the AARs for the 50 ASPM airports for 2000 to date.

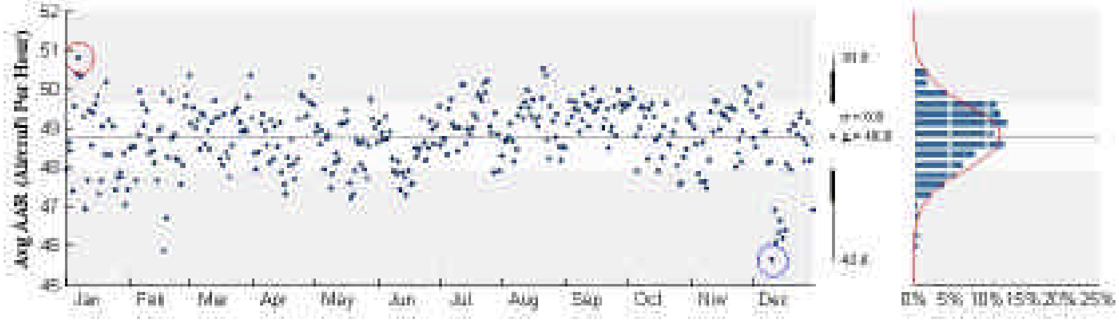

Figure 31. Distributions for airport reported arrival rates for domestic flights in 2000.

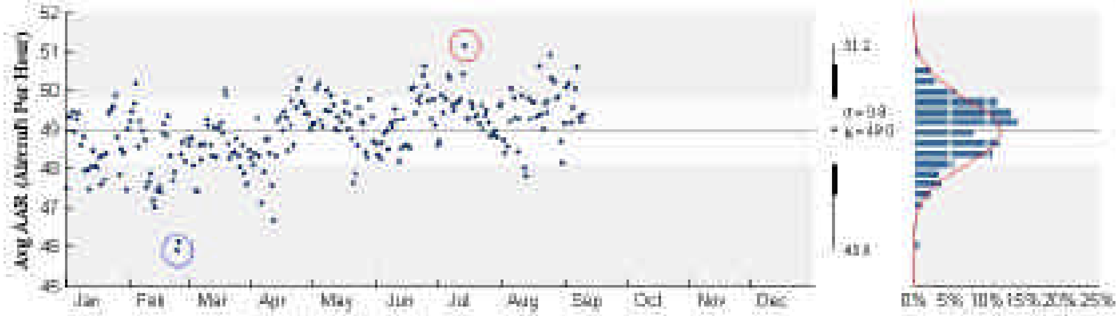

Figure 32. Distributions airport reported arrival rates for domestic flights up to Sept. 2001.

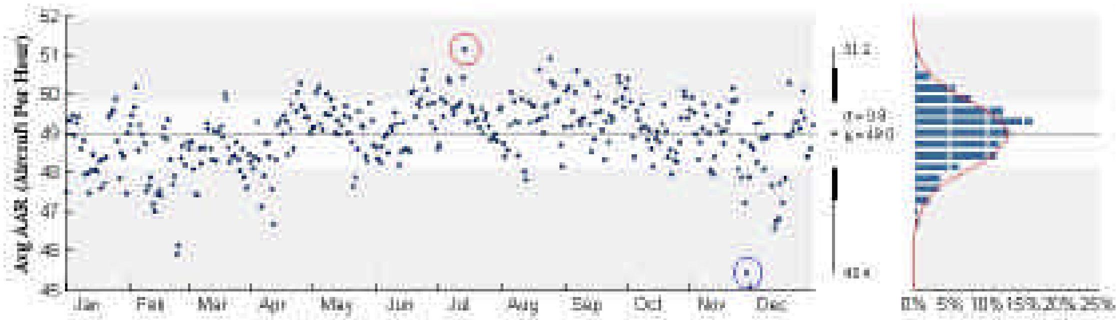

Figure 33. Distributions airport reported arrival rates for domestic flights in 2001.

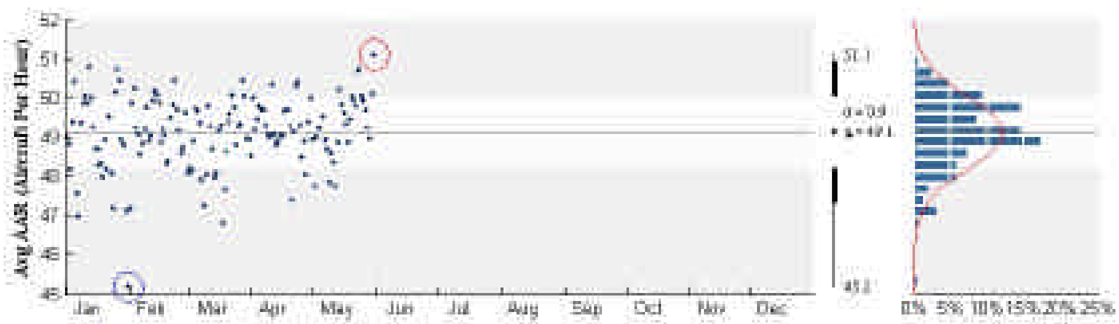

Figure 34. Distributions airport reported arrival rates for domestic flights in 2002.

4.1.7 Airport Approach Conditions (IFR vs. VFR)

ASPM provides data to describe the meteorological conditions at airports. VFR signifies a visual approach condition at the airport, while IFR indicates that there were instrument approach conditions at the airport. IFR and VFR conditions were recorded in ASPM for every quarter hour at each airport. The total number of 15-minute periods over all airports with IFR and VFR conditions recorded were counted and used to compute the percentage of IFR vs. VFR conditions for each day. **Figure 35** through **Figure 37** show what percentage of IFR and VFR conditions occurred every day over all ASPM-50 airports from January of 2000 to May 2002.

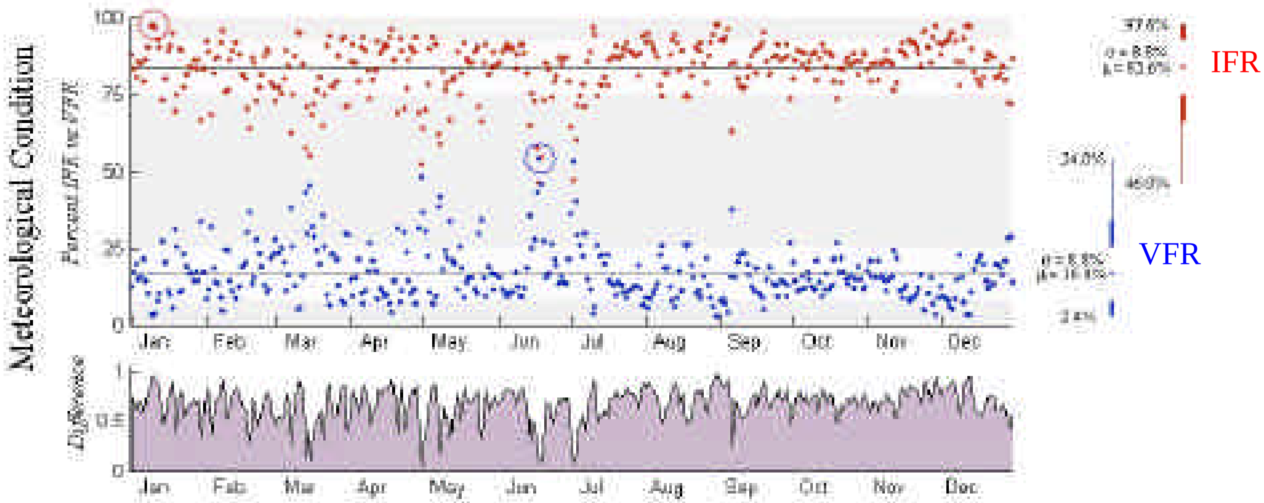

Figure 35. Average IFR vs. VFR conditions for 2000.

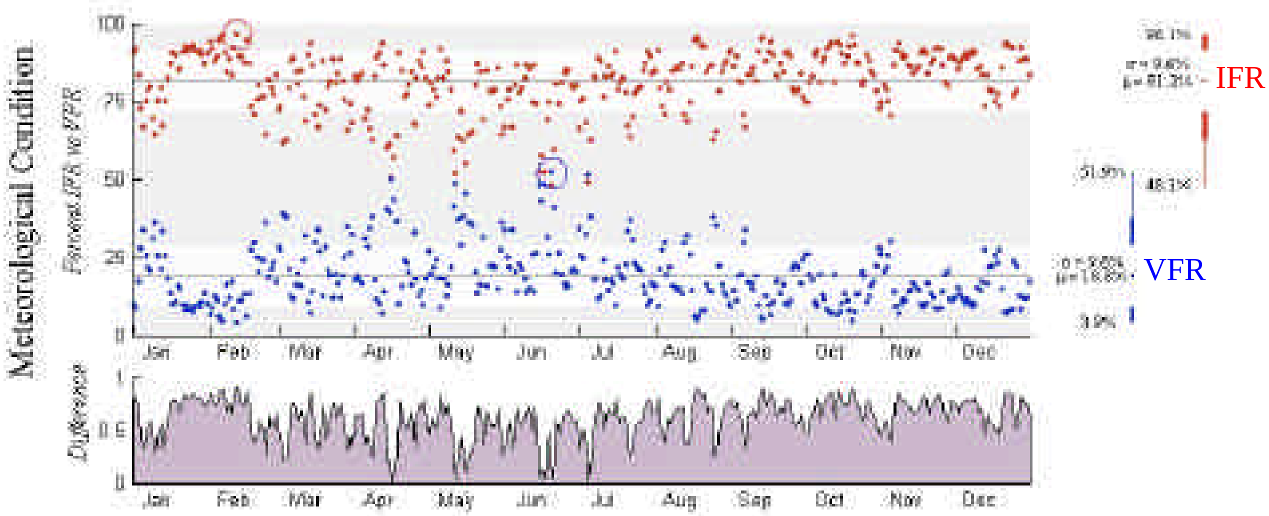

Figure 36. Average IFR vs. VFR conditions for 2001.

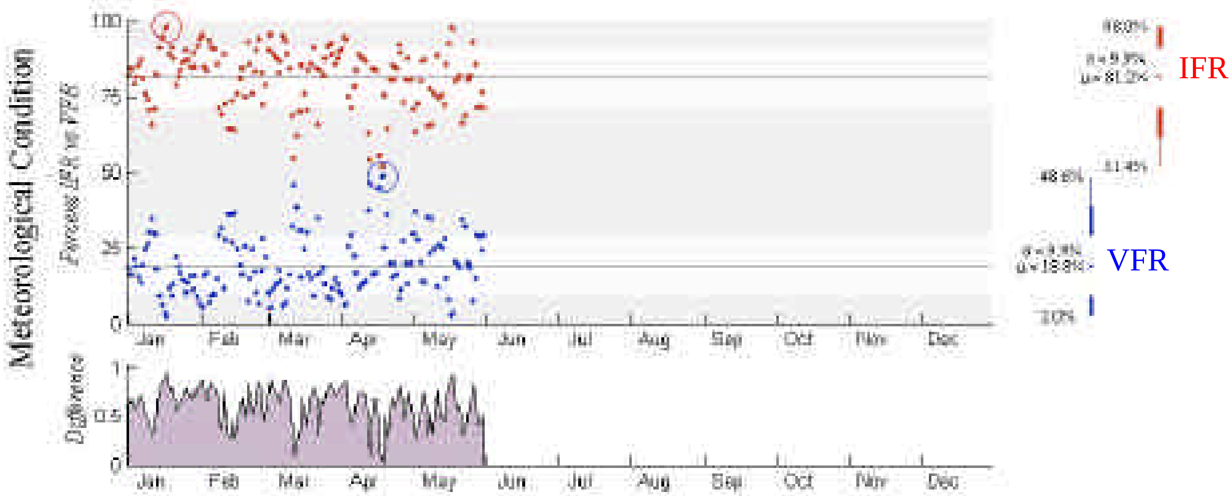

Figure 37. Average IFR vs. VFR conditions for 2002.

4.1.8 Airport Ceiling Conditions

ASPM provides data to describe the ceiling conditions at airports. **Figure 38** through **Figure 40** indicate the average ceiling conditions for the ASPM-50 airports for 2000 to the present date.

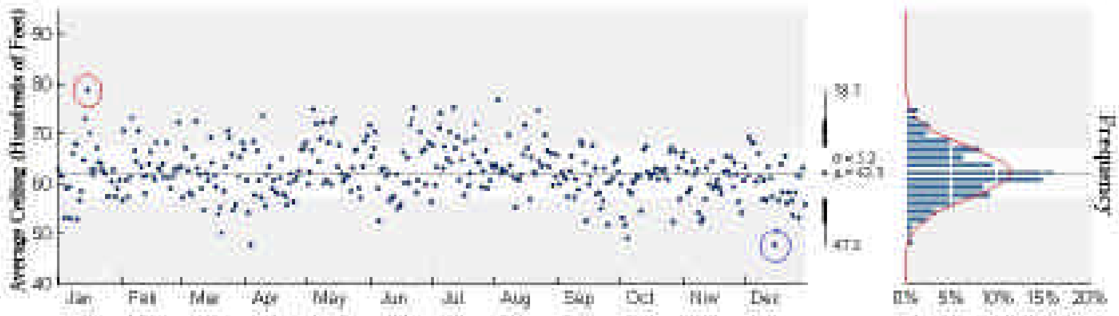

Figure 38. Average ceiling conditions for 2000.

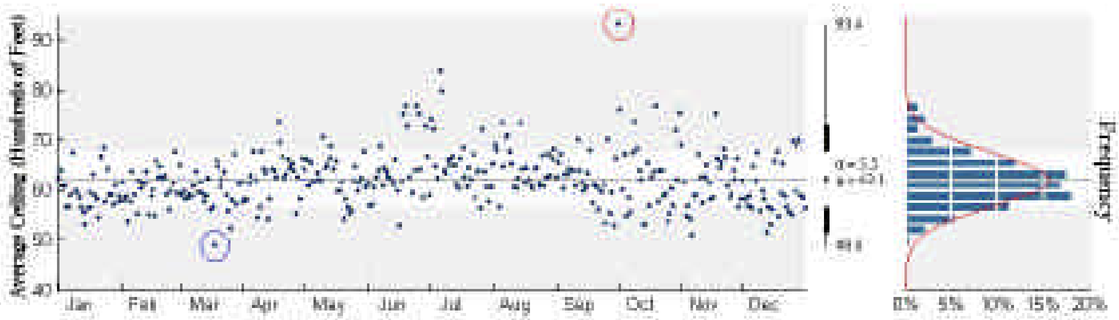

Figure 39. Average ceiling conditions for 2001

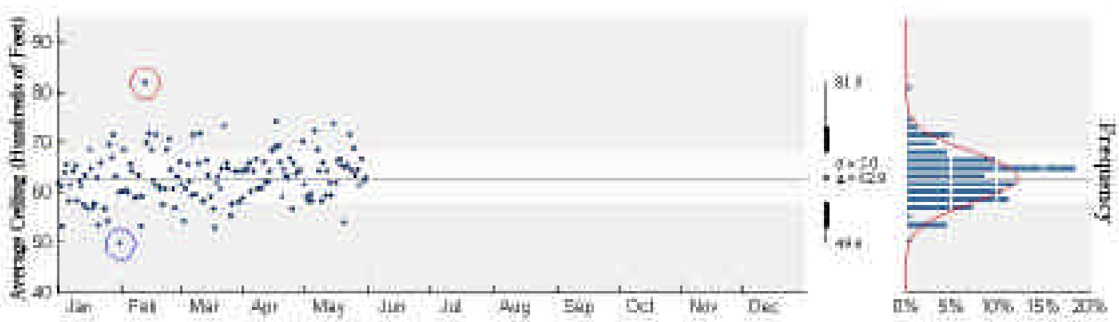

Figure 40. Average ceiling conditions for 2002.

4.1.9 Airport Visibility Conditions

ASPM provides data to describe the visibility conditions at airports. **Figure 41** through **Figure 43** indicate the visibility conditions for the ASPM-50 airports for 2000 to date. Note: There was a severe outlier in the July 11, 2000, so the visibility data record for that date was removed from the 2000 data set.

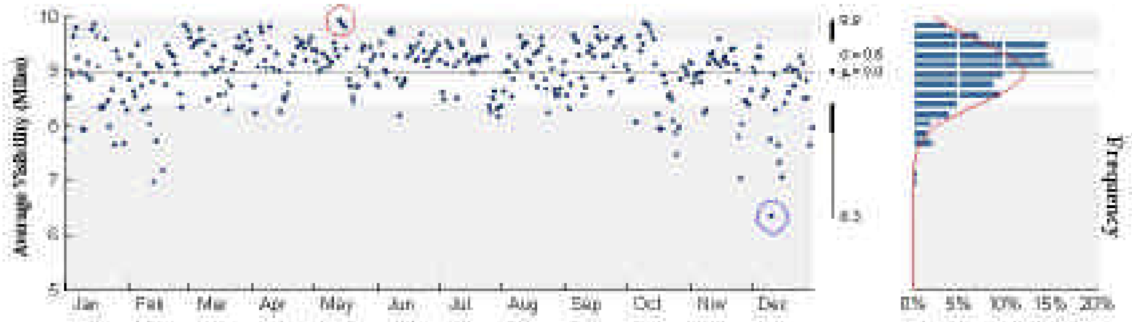

Figure 41. Average visibility conditions for 2000.

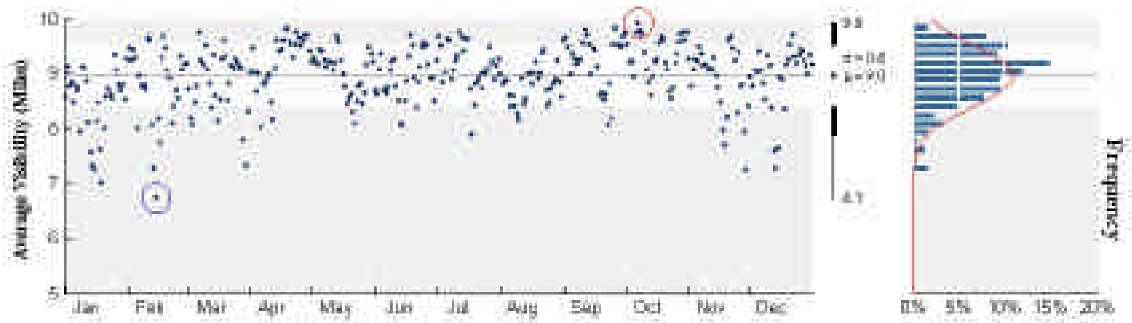

Figure 42. Average visibility conditions for 2001.

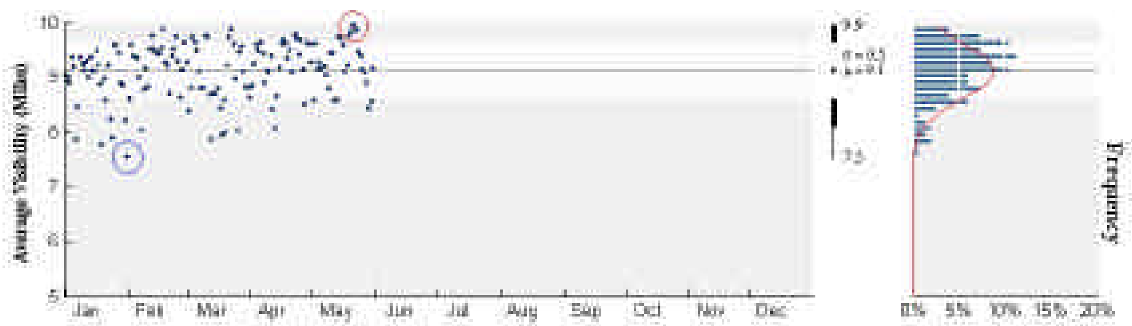

Figure 43. Average visibility conditions for 2002.

4.1.10 Changes in Runway Configuration Conditions

As shown in **Figure 44** through **Figure 47**, ASPM provides data describing the total number of runway configuration changes for ASPM-50 airports. While not a direct measurement of weather conditions, this statistic is highly linked to wind speed and direction. Additionally, it is linked to traffic demand. Runway configuration changes were fairly constant from the beginning of 2000 through the summer of 2001. However, in August 2001 there is a significant increase. The events of September 11, 2001 seem to have mitigated this increase in configuration changes (probably due to lower volume) for a short time. Note that the mean number of runway configuration changes for 2002 is over twice that of the mean for 2000.

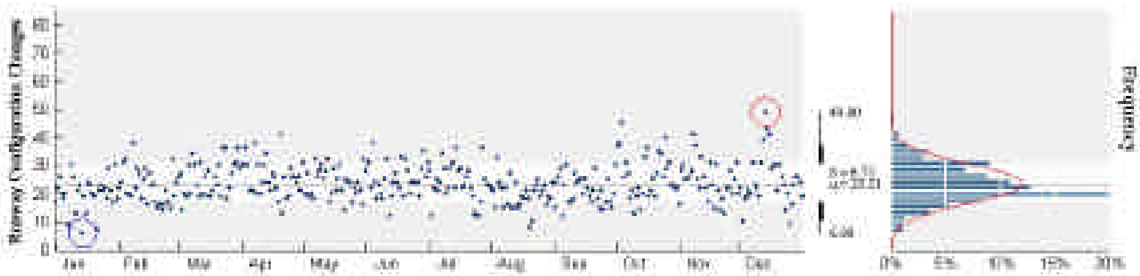

Figure 44. Average number of runway configuration changes for 2000.

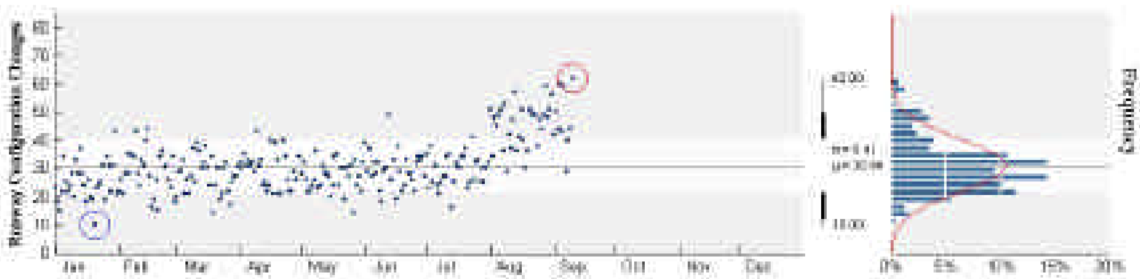

Figure 45. Average number of runway configuration changes for 2001.

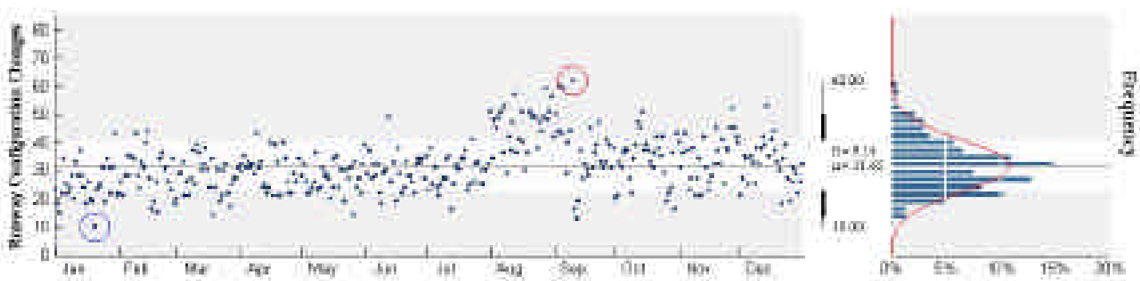

Figure 46. Average number of runway configuration changes for 2000.

Figure 47. Average number of runway configuration changes for 2002.

4.2 Statistics of NAS Controls

Next, aggregate properties of the following NAS controls are reviewed:

- Ground Delay Programs
- Cancellations
- Ground Stops
- MIT Restrictions
- Airborne Holding

4.2.1 Ground Delay Programs

During a GDP, the arrival flow rate into an airport with a capacity shortfall (or abundance of demand) is reduced to avoid inundating the airport with more arrivals than it can safely accommodate. Arrival demand is brought in line with arrival capacity by creating a virtual queue of arrival slots. The number of slots in any given hour corresponds to the estimated number of aircraft that the airport can land. Each aircraft estimated (or scheduled) to arrive during the GDP time horizon is assigned to an arrival slot (time interval). From this Controlled Time of Arrival (CTA), a Controlled Time of Departure (CTD) is computed for each aircraft by subtracting the estimated flying time. The difference between the CTD of a flight and the estimated arrival time just prior to implementing the GDP is the amount of FAA-issued ground delay that flight must absorb. The net effect of a GDP is to transfer anticipated airborne holding back onto the ground.

There are many reasons why a GDP might be run. Sometimes, traffic flow into an airport is reduced to slow down the flow through an unrelated piece of airspace. Other times, there is an unusually large arrival demand that exceeds normal airport acceptance rates. But most of the time, a GDP simply reflects deterioration in airport conditions. Weather is the most frequent culprit. Thus, frequency and magnitude of GDPs make a strong statement about NAS conditions.

A simplistic metric of GDPs is how many are being run on a given day. This gives a rough indication of how many airports are in a state of demand-capacity imbalance. A typical day in the NAS might have 1 or 2 GDPs in place, while 5 or more GDPs indicates NAS-wide problems.

A more refined metric should take into account the scope and magnitude of the GDP. The period of duration varies with the GDP, with the most common duration being in the 4-6 hour range. In addition to its temporal scope, there is a geographical scope. Flights bound for a GDP airport are often exempt from FAA-issued ground delay based on the proximity of their origin airport from the GDP airport. It is quite common to restrict the application of ground delays to those flights originating within traffic control centers that are immediately adjacent to the traffic control center that houses the GDP airport (called a *first tier* program). A smaller geographic scope of a program is a way of mitigating the potential damages of weather forecast uncertainty: larger scope programs (e.g., the entire NAS) tend to capture flights departing earlier in time and, therefore, run the risk of assigning them ground delays that cannot be recovered in the event of a GDP cancellation.

Another GDP statistic to consider is average ground delay per flight. This can be computed by dividing total ground delay by the number of flights to which ground delay was applied ("affected flights"). We have already seen that the numerator is essentially constant. The denominator will vary with the geographic scope of the program. Thus, this statistic would say very little about the overall scope and magnitude of the program.

For this reason, it is often beneficial to consider the total number of minutes of ground delay issued in the GDP rather than average ground delay per flight. This figure will increase with the number of flights in the program and with the duration of the program, but is independent of the geographic scope chosen for the program. Of course, one could compute average delay by dividing total ground delay by number of flights involved in the program (whether they were exempt from ground delay or not).

For this study, GDP data were collected from the ATCSCC (Herndon, VA). **Figure 48** and **Table 5** show the aggregate GDP data across the NAS for the years 1998 through 2001. These data indicate an increase in the use of GDPs each year. When all the GDPs are analyzed in terms of the airport where the GDP is issued, as shown in **Figure 49** through **Figure 58**, it is clear that there are certain airports for which GDPs are issued more than others. These airports are: ATL, BOS, EWR, LAX, LGA, ORD, PHL, and SFO. The three leading airports where GDPs are issued (based on 2000 and 2001 data) are: SFO, LGA, and ORD.

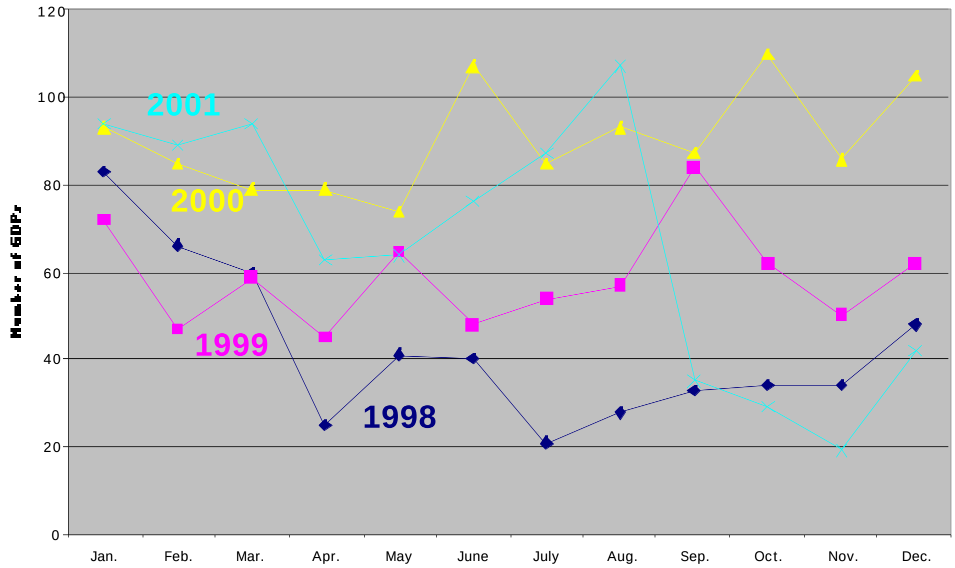

Figure 48. The number of GDPs issued per month in 1998 through 2001.

Table 5. Statistics for GDPs issued during 1998 through 2002.

Year	Number of GDPs	Average per Day
1998	513	1.4
1999	705	1.9
2000	1083	3.0
2001	799	2.8*
2002	398	1.6**

Notes: * The 2001 average is determined using Jan.-Aug. (243 days) data only.
** The 2002 number of GDPs and average is based on using Jan.-Aug. (243 days) data only.

In **Figure 49** through **Figure 52**, airports that ran GDPs in the period from January 1996 to April 2002 are compared.

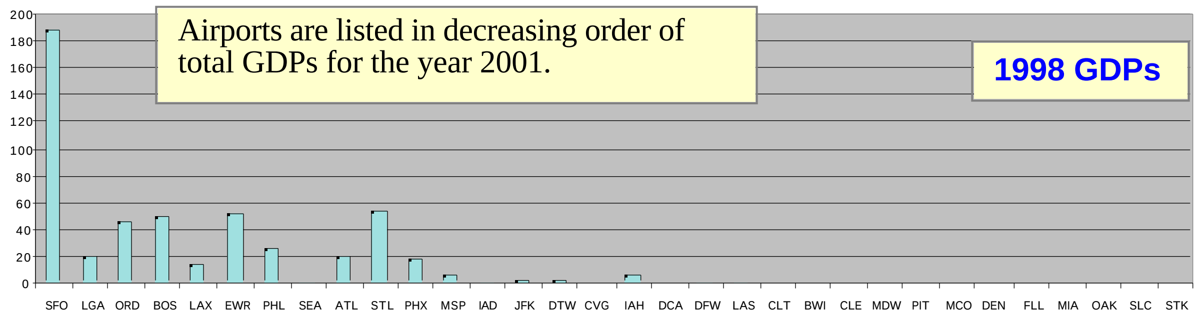

Figure 49. Airports where GDPs occurred in 1998.

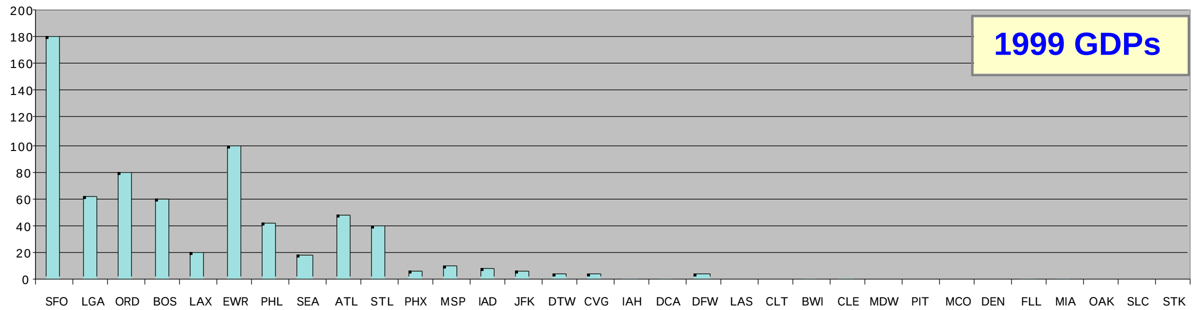

Figure 50. Airports where GDPs occurred in 1999.

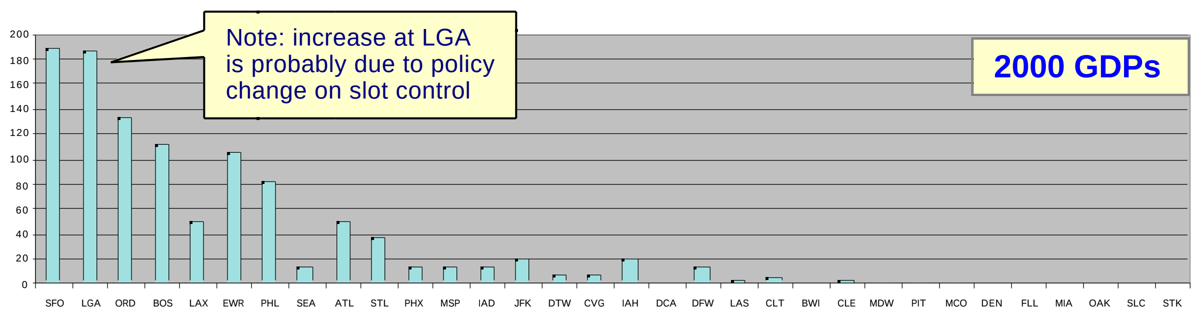

Figure 51. Airports where GDPs occurred in 2000.

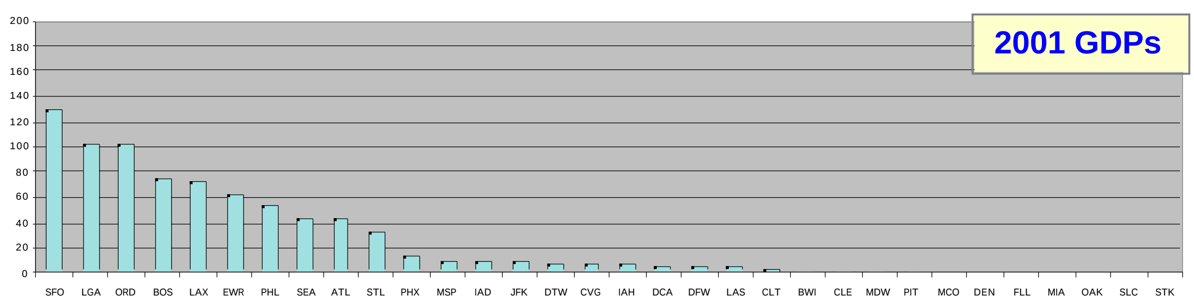

Figure 52. Airports where GDPs occurred in 2001.

Figure 53 through **Figure 58** further describe the statistics of GDPs from 1998 to 2002.

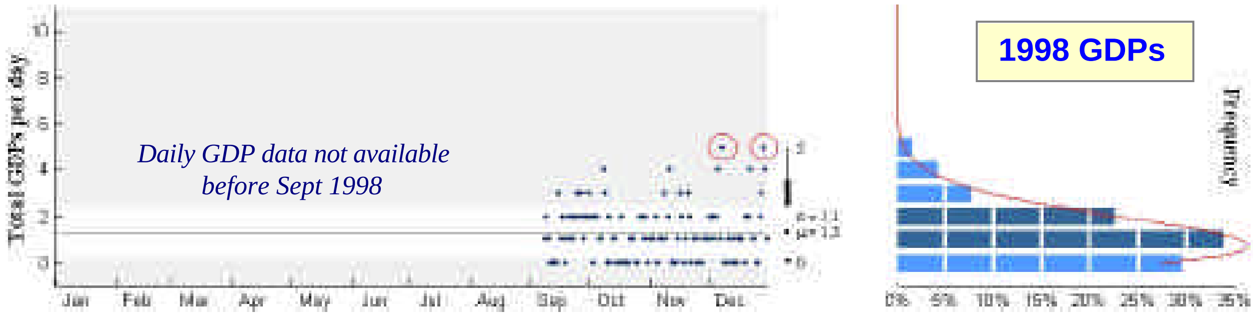

Figure 53. Number of daily GDPs that occurred in 1998[4].

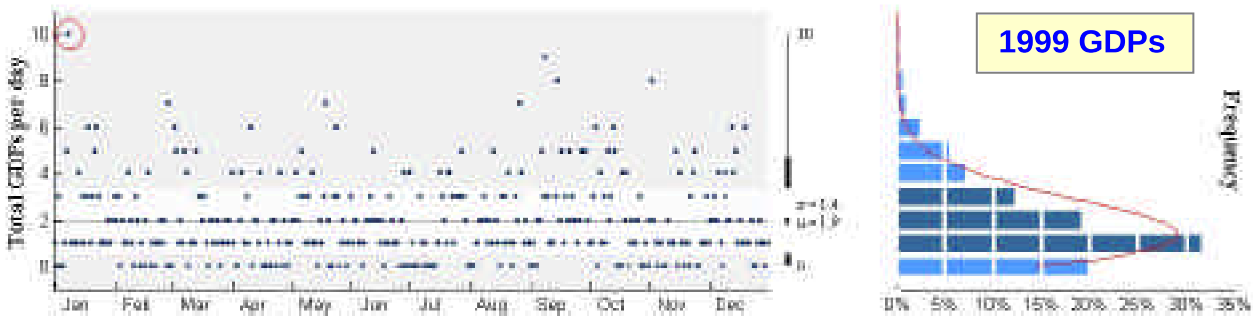

Figure 54. Number of daily GDPs that occurred in 1999.

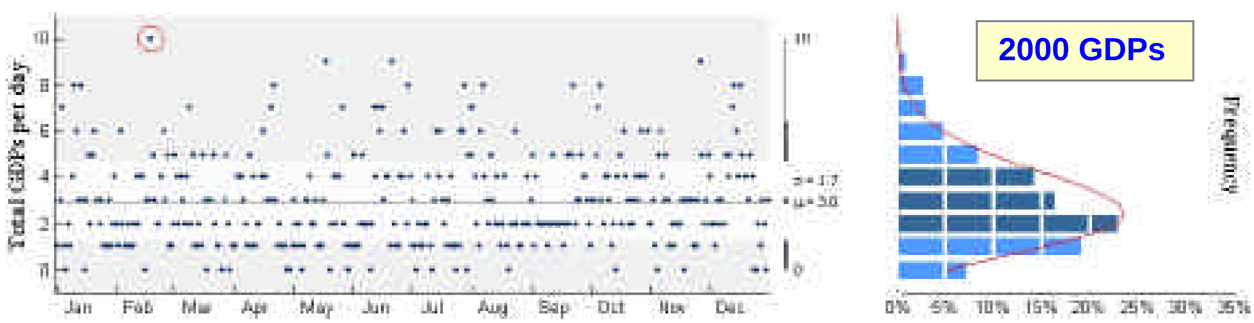

Figure 55. Number of daily GDPs that occurred in 2000.

[4] A Poisson distribution was assumed for determining the mean and standard deviation for GDP figures.

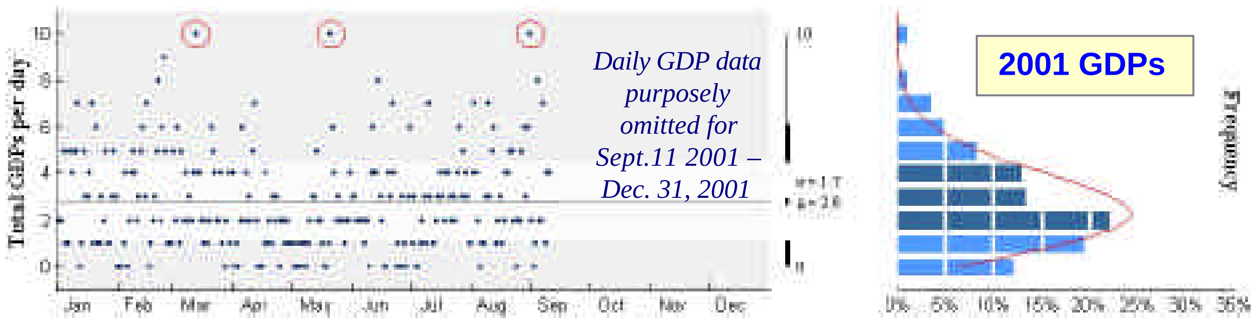

Figure 56. Number of daily GDPs that occurred in 2001 up to Sept. 10, 2001.

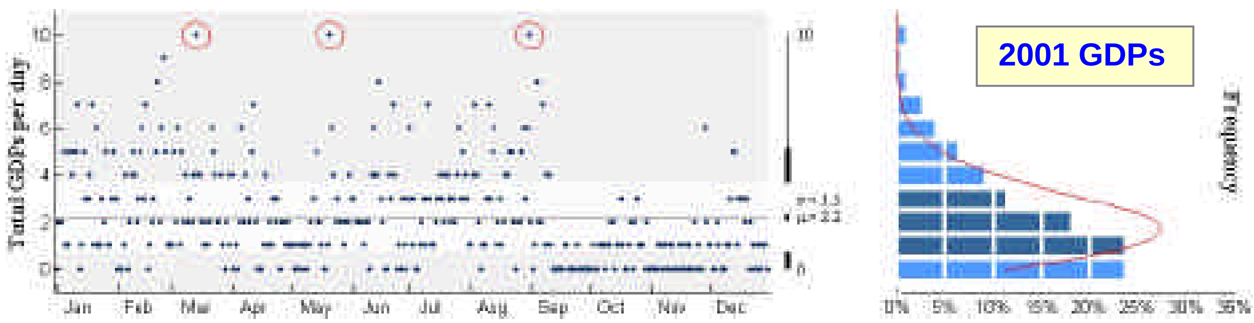

Figure 57. Number of daily GDPs that occurred in 2001.

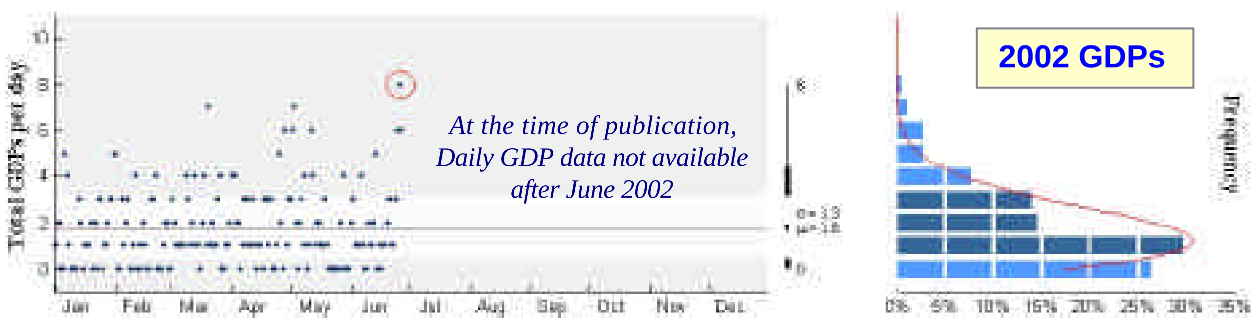

Figure 58. Number of daily GDPs that occurred in 2002.

Next, the geographic distribution of GDPs is illustrated. In **Figure 59** through **Figure 63**, the areas of the circles at each airport are proportional to the number of GDPs, which occurred at the respective airport. Those colored red have a total number of GDPs that is higher than one standard deviation above the mean. Those colored yellow are within one standard deviation from the mean. Blue is used for airports with a total number of GDPs less than one standard deviation below the mean (thus, these circles are very small). These figures show that GDPs are concentrated at the major hub airports[5], SFO, and in the northeast corridor. This spatial distribution does not change much year to year.

[5] Appendix G presents the major hub airports as defined by the BTS.

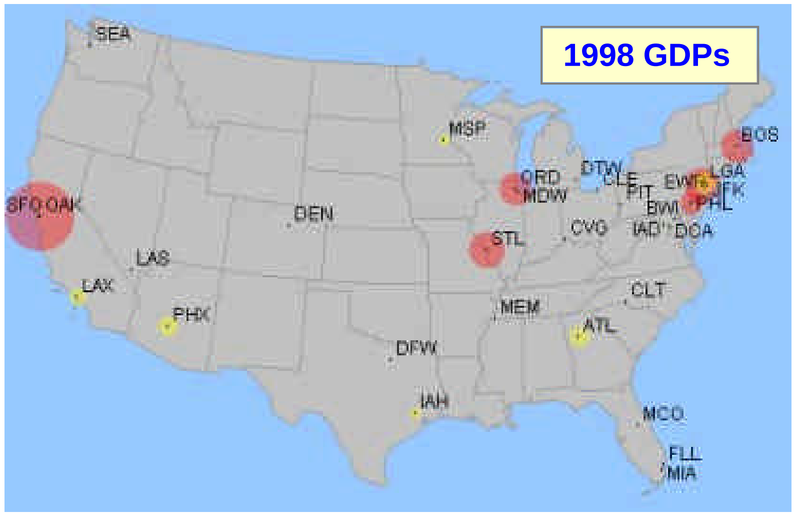

Figure 59. Geographical location of airports (and magnitude) where GDPs occurred in 1998.

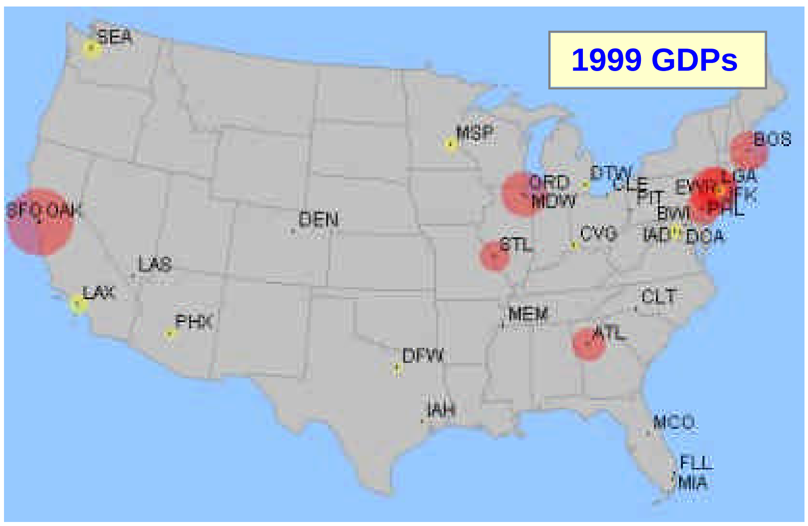

Figure 60. Geographical location of airports (and magnitude) where GDPs occurred in 1999.

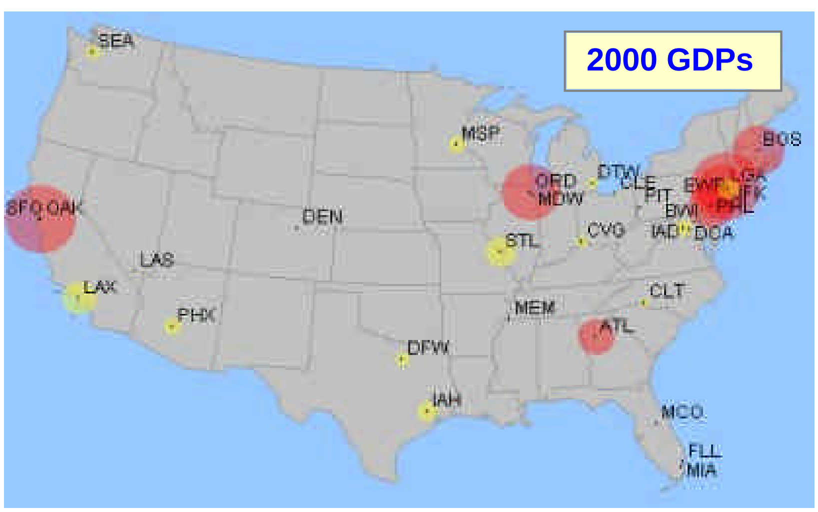

Figure 61. Geographical location of airports (and magnitude) where GDPs occurred in 2000.

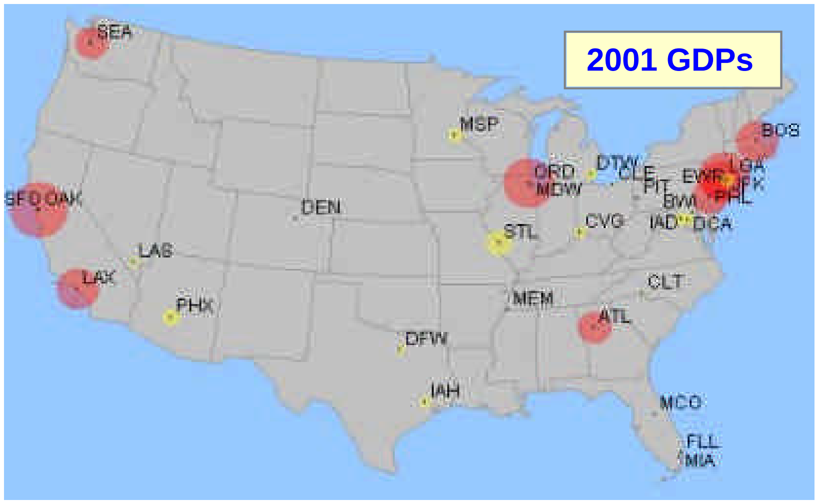

Figure 62. Geographical location of airports (and magnitude) where GDPs occurred in 2001.

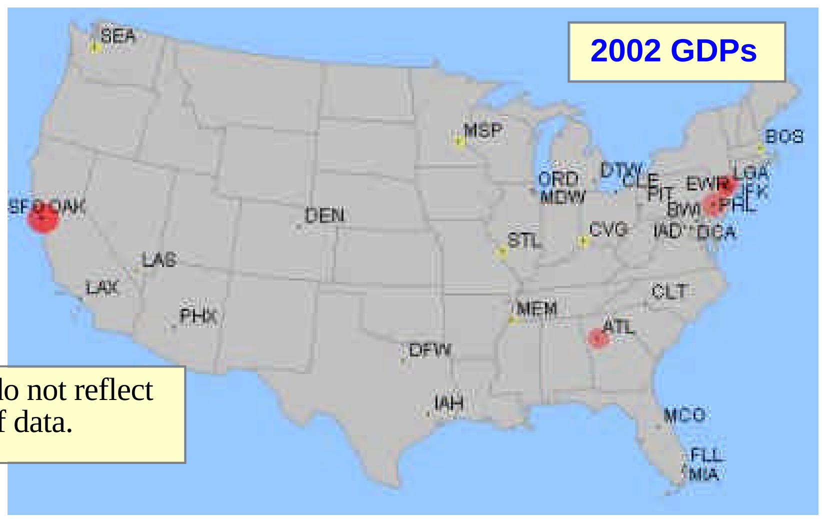

These data do not reflect a full year of data.

Figure 63. Geographical location of airports (and magnitude) where GDPs occurred in 2002.

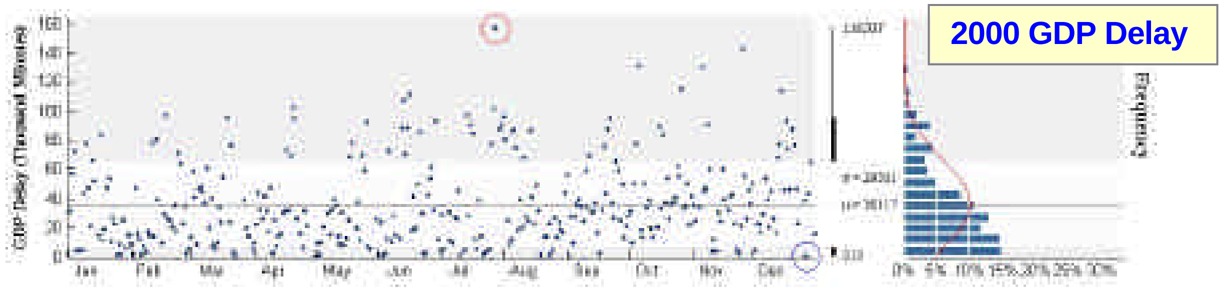

Figure 64. Number of delay minutes due to GDPs that occurred in 2000.

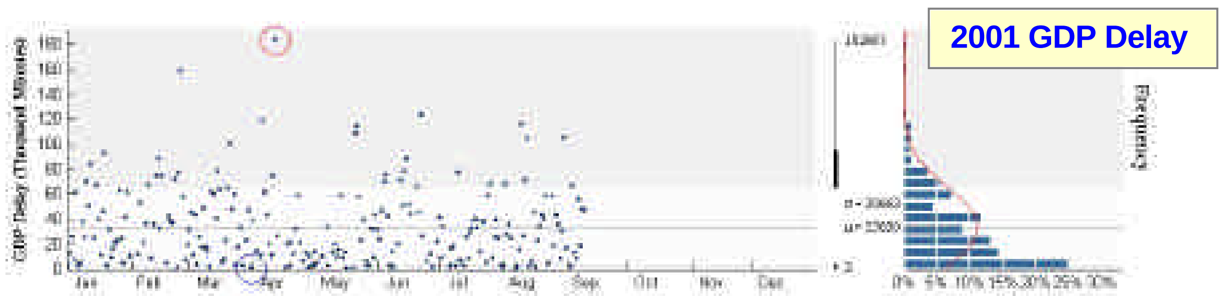

Figure 65. Number of delay minutes due to GDPs that occurred in 2001 prior to September 11.

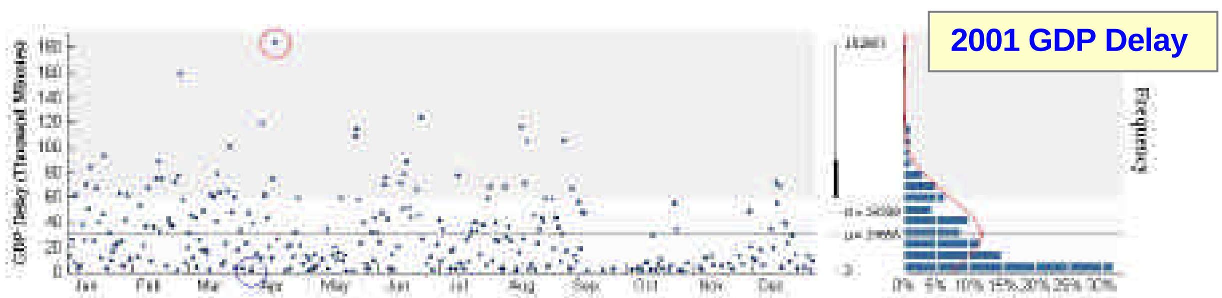

Figure 66. Number of delay minutes due to GDPs that occurred in 2001.

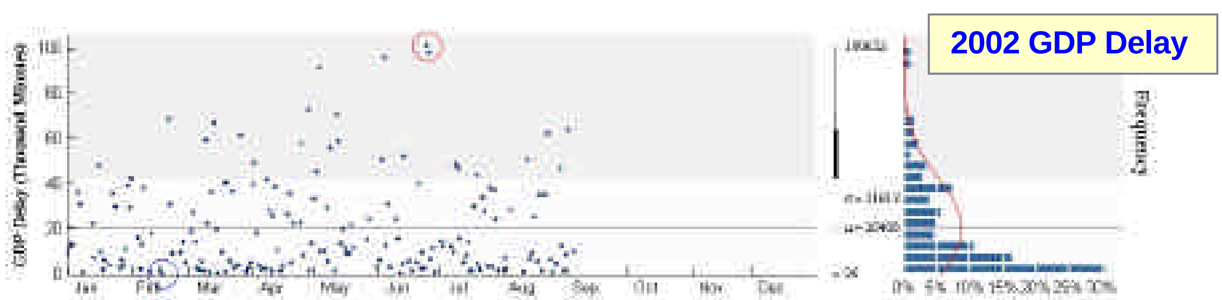

Figure 67. Number of delay minutes due to GDPs that occurred in 2002.

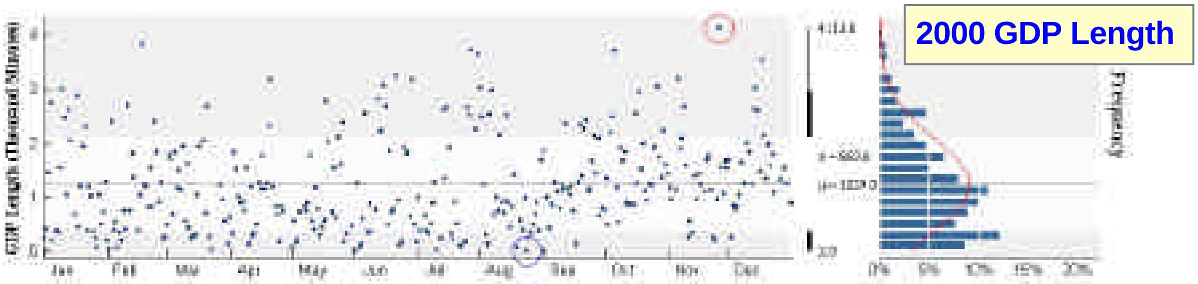

Figure 68. Duration of GDPs that occurred in 2000.

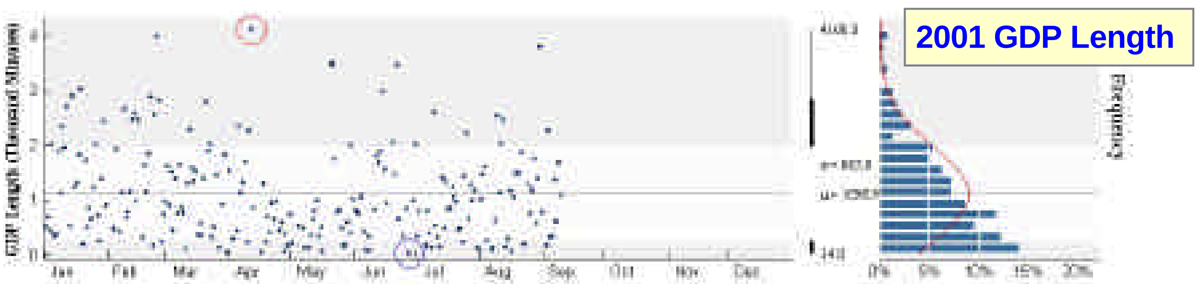

Figure 69. Duration of GDPs that occurred in 2001 prior to September 11.

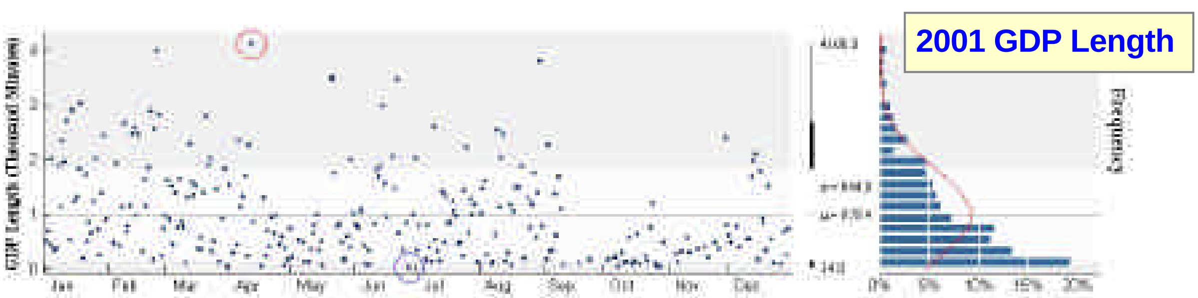

Figure 70. Duration of GDPs that occurred in 2001.

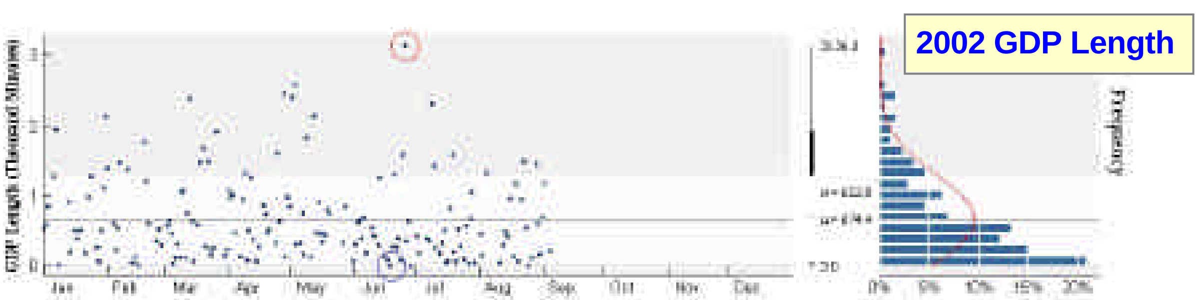

Figure 71. Duration of GDPs that occurred in 2002.

4.2.2 Cancellations

ASPM cancellations are recorded for arrival as well as departure traffic; these statistics are illustrated next in **Figure 72** and **Figure 73.** Note that the peak number of cancellations for this period is highly related to the location of the weather activity. **Appendix C** provides details of how these cancellations vary with respect to weather conditions.

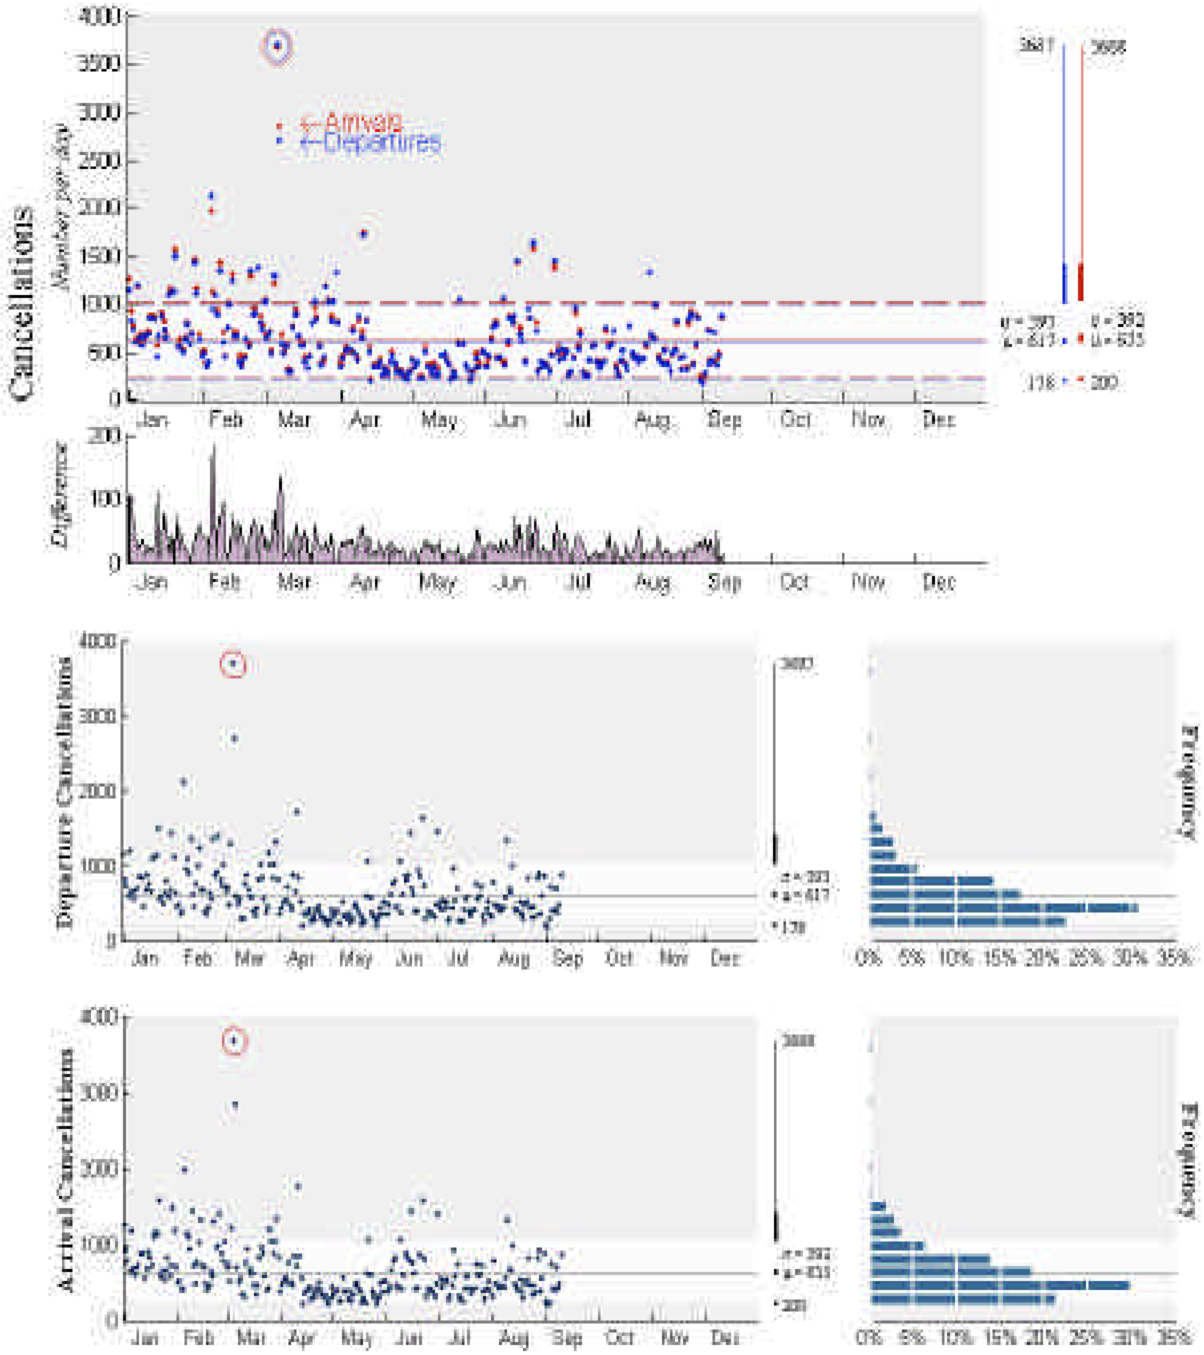

Figure 72. Cancellations at 50 major airports up to Sept. 10 in the year 2001.

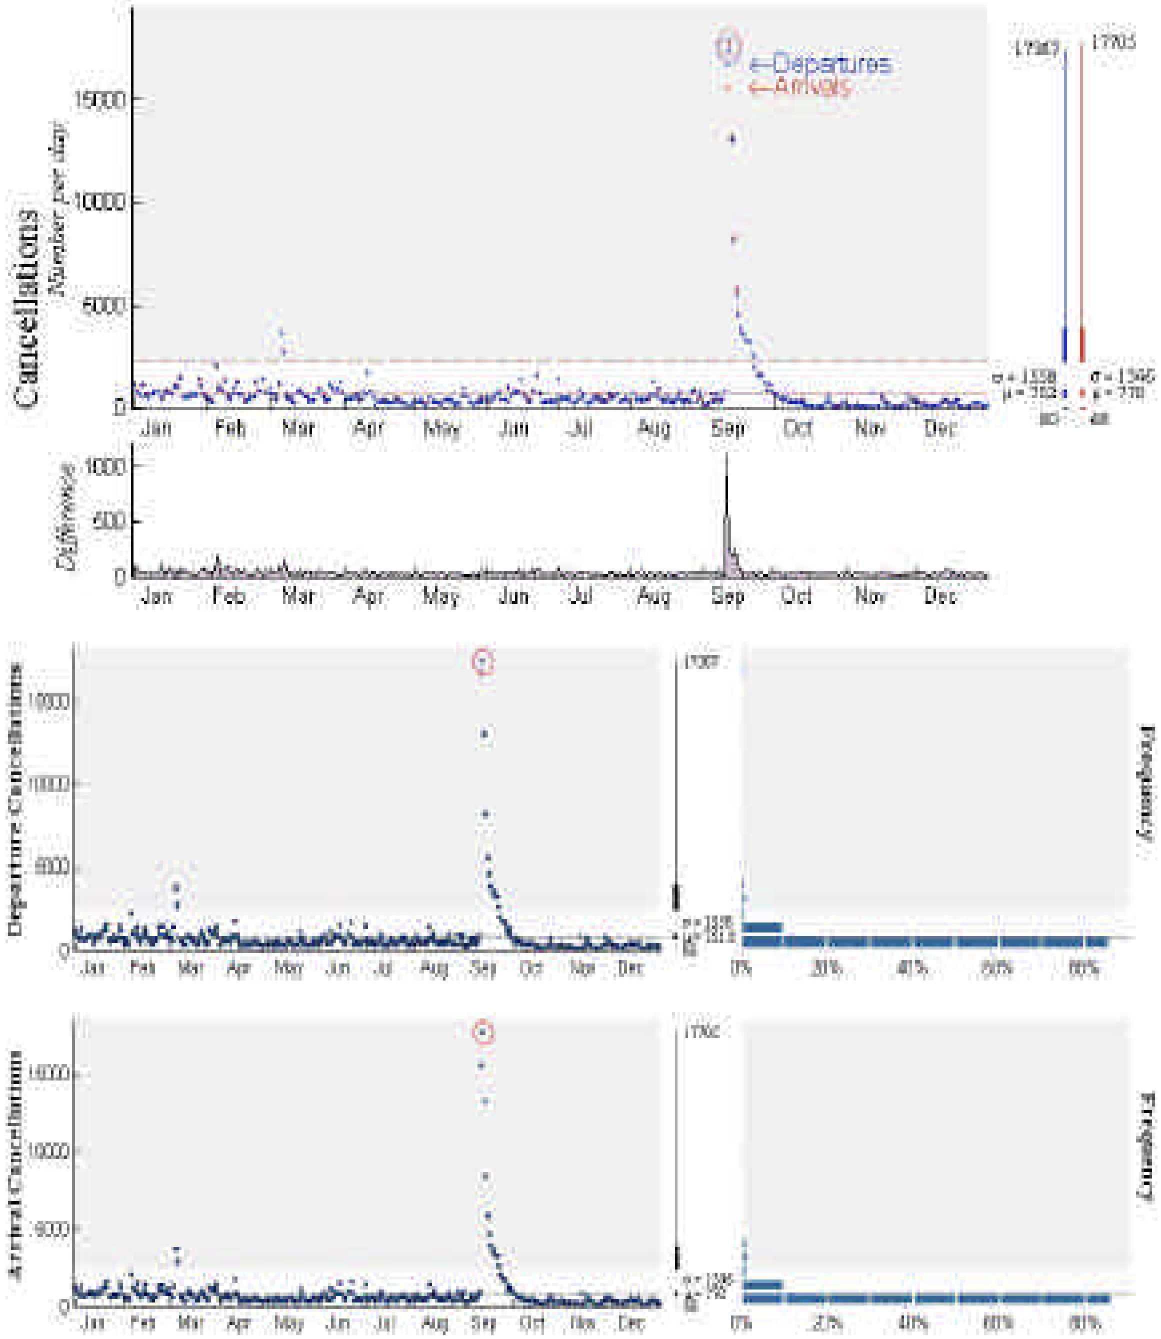

Figure 73. Cancellations at 50 major airports up to Sept. 10 in the year 2001.

4.2.3 Ground Stops

A Ground Stop (GS) is a traffic flow initiative in which all traffic is banned from departure for a period of specification. This is usually implemented to reduce the flow of traffic in an unexpected situation, such as highly unpredictable convective weather activity. The criterion for inclusion of a GDP is based on arrival time. Unlike a GDP, the criterion for inclusion in a GS is departure time. For instance, a 11:00Z - 11:30Z GS at the Cleveland ARTCC (ZOB) means that no flights within ZOB are allowed to depart between 11:00Z and 11:30Z. A GS is usually short-lived, highly tactical maneuver. For this reason, the GS frequency is an indication of unpredictable, disruptive events in the NAS.

For this study, GS data were collected from the ATCSCC (Herndon, VA). **Table 6** and **Figure 74** through **Figure 80** show the aggregate GS data across the NAS for the years 2000 through 2002. **Figure 81** through **Figure 83** show the geographic distribution of GS data across the NAS for 2000 through 2002. An increase during the summer months is primarily resulting from the convective weather season. In September 2001, a drop in the traffic volume across the NAS occurred after the September 11th tragedy, and thus, the number of GSs during the remainder of the year was low.

Table 6. Statistics for GSs issued during 2000 through 2002.

Year	Number of GSs	Average per Day
2000	491*	1.8*
2001	955	2.6
2002	786**	3.2**

Notes: * Due to a lack of data, the 2000 average is determined using April – Dec. (275 days) only.
** The 2002 number of GSs and average is based on using Jan. – Aug. (243 days) data only.

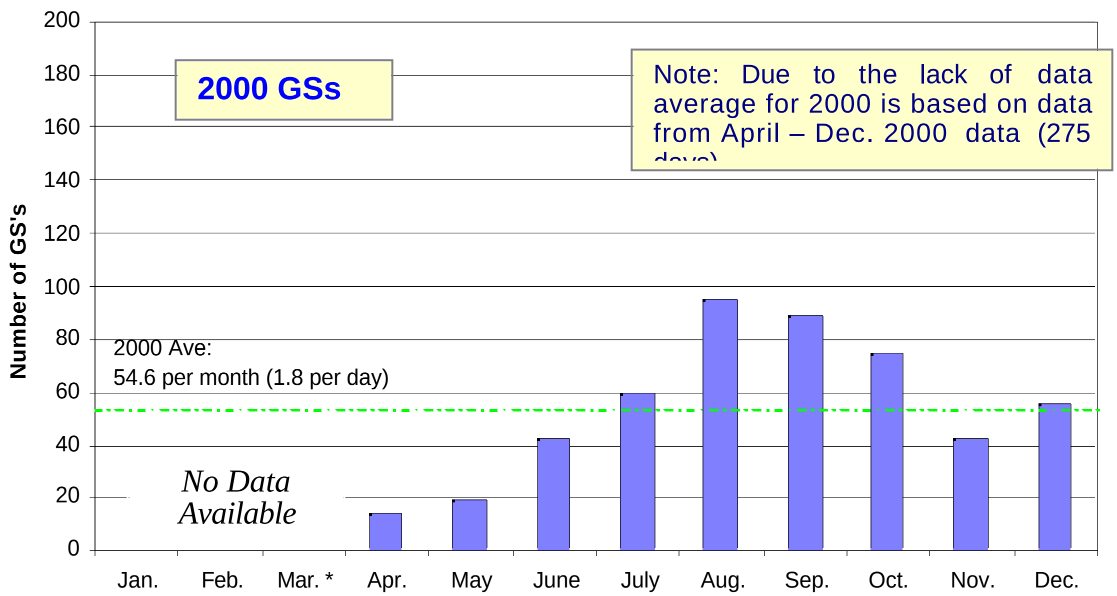

Figure 74. The number of ground stops issued per month in 2000.

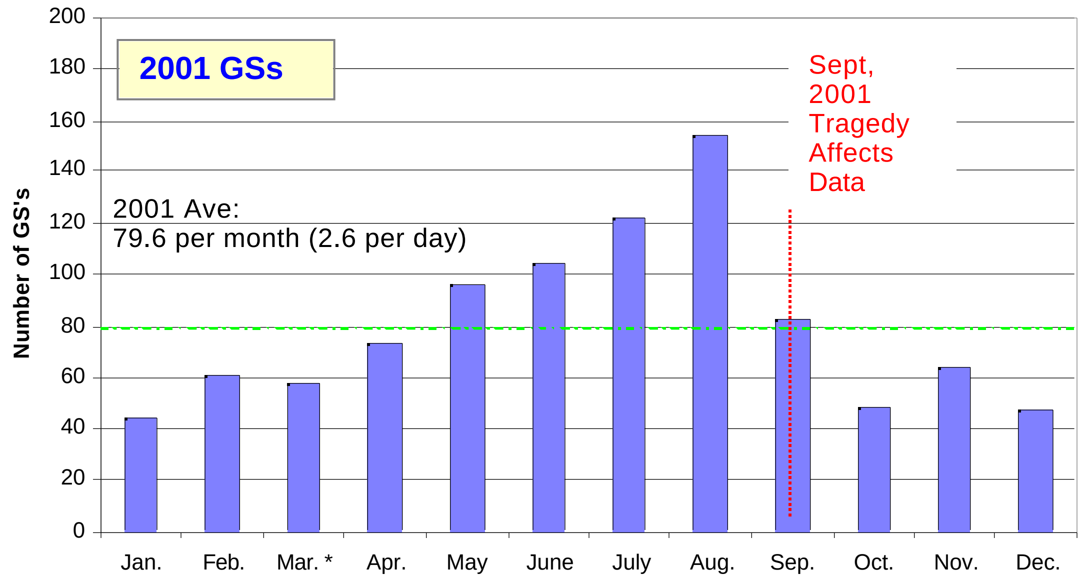

Figure 75. The number of ground stops issued per month in 2001.

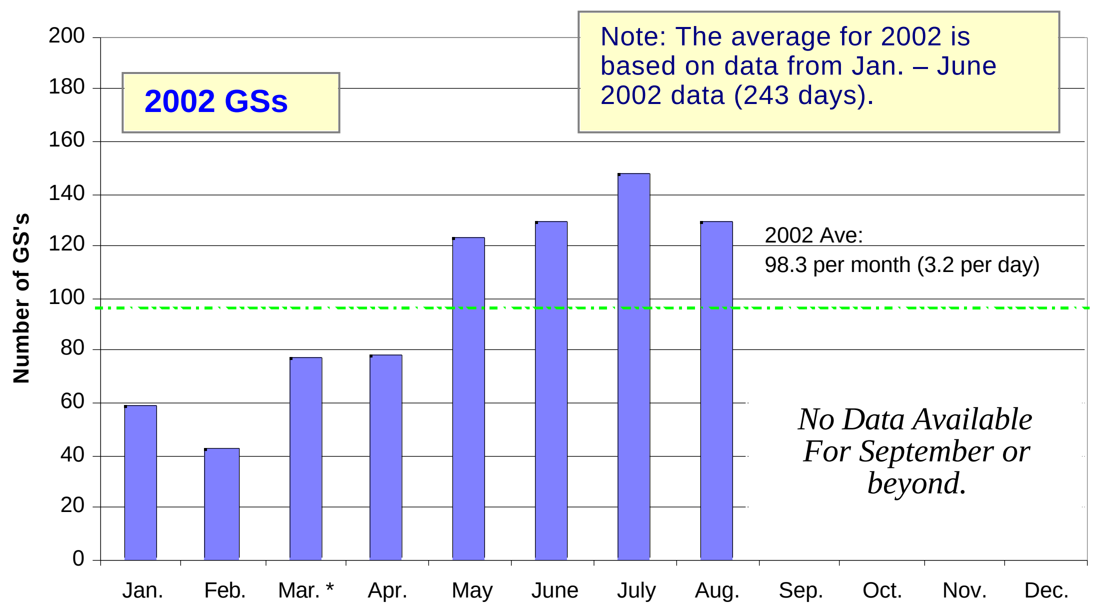

Figure 76. The number of ground stops issued per month in 2002.

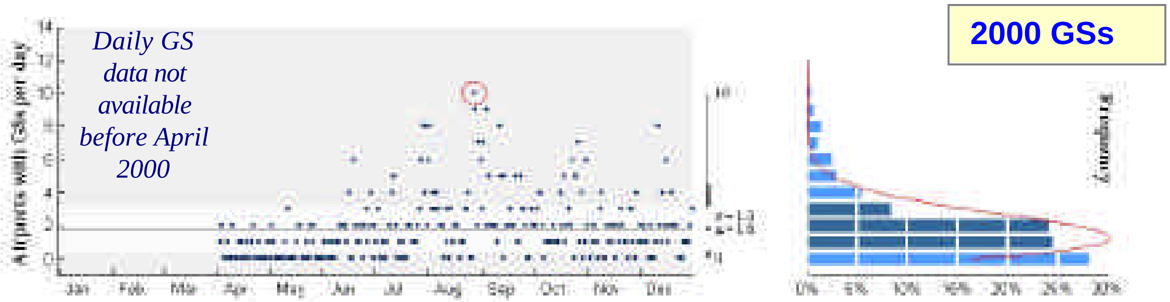

Figure 77. Number of daily GSs that occurred in 2000[6].

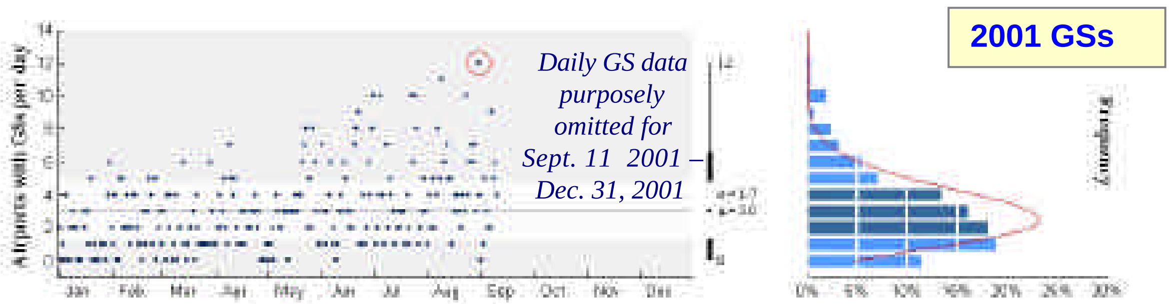

Figure 78. Number of daily GSs that occurred in 2001 up to Sept. 10, 2001.

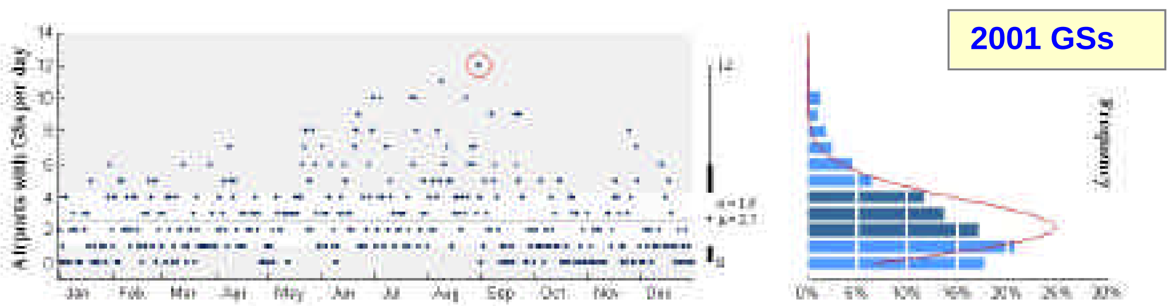

Figure 79. Number of daily GSs that occurred in 2001.

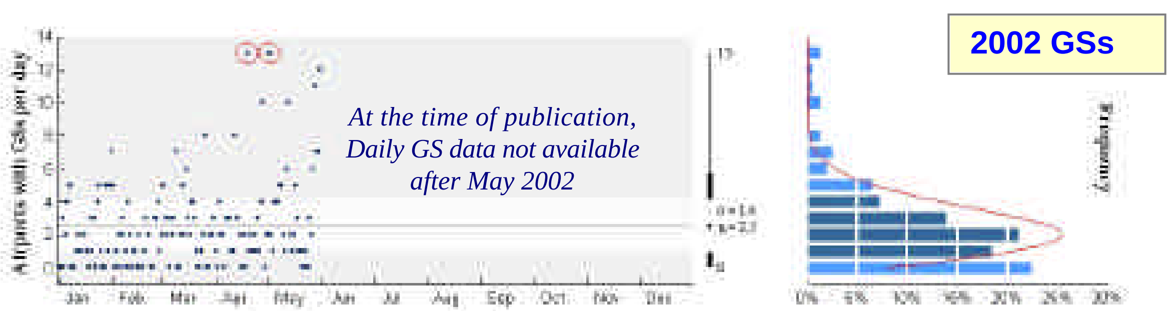

Figure 80. Number of daily GSs that occurred in 2002.

[6] These are GSs that were run independent of GDPs. A Poisson distribution was assumed for determining the mean and standard deviation for GS figures.

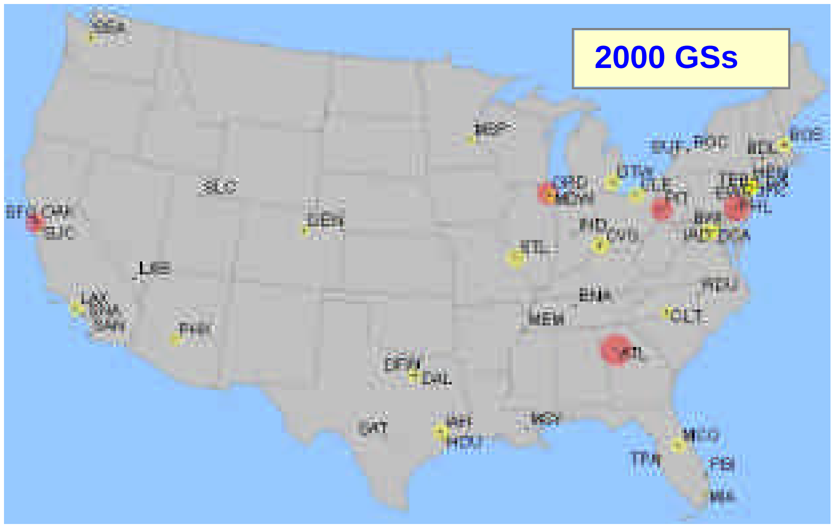

Figure 81. Geographical location of airports (and magnitude) where GSs occurred in 2000.

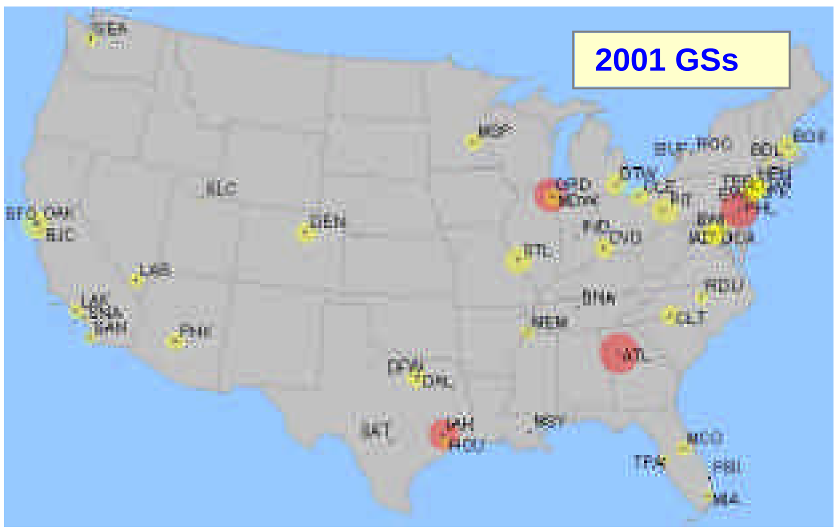

Figure 82. Geographical location of airports (and magnitude) where GSs occurred in 2001.

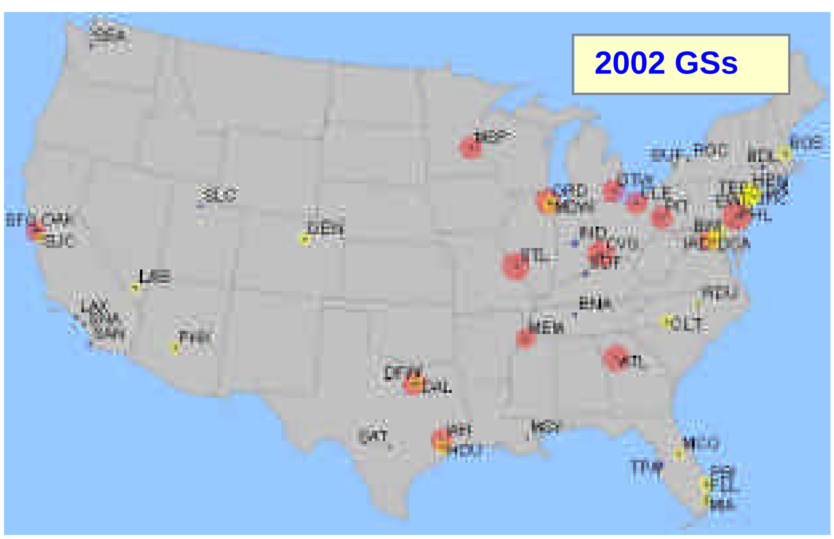

Figure 83. Geographical location of airports (and magnitude) where GSs occurred in 2002.

4.2.4 MIT Restrictions

TFM places MIT restrictions on lines of aircraft traveling in the same direction to reduce the rate of flow. A MIT of 30 miles means that each aircraft within the flow on a jet route must maintain at least 30 miles of separation between itself and the aircraft in front of it. MIT restrictions are usually placed at the boundary of adjacent sectors or centers to provide the spacing necessary for safe travel with the downstream sector or center. Some MIT restrictions are standard operating procedure. But most of the time, they are indicative of a combination of excessive demand and degraded capacity.

Since MIT restrictions often impact more than one center, they are coordinated through and electronically recorded by the ATCSCC. For each restriction, there is a record of the stream of traffic to which it applies, the start and end times of the restriction, and the required spacing. The primary user of this database is the Quality Assurance (QA) department of the ATCSCC. We obtained NAS-wide MIT restrictions data for years 2000 and 2001 from ATCSCC QA. (In the course of our research, we were not able to obtain MIT data for year 2002.) We used this data to compute year-long summary statistics for each of the 20 centers in the continental US and for the entire NAS. The summary statistics in **Table 7** were computed as follows:

- The average duration of a restriction equals the end time minus the start time.
- The average number of restrictions per day equals the total number of restrictions divided by 366 (for the 2000 leap year) or 365 days (for the 2001 year).
- The average restriction size is the total of all restriction sizes divided by the total number of restrictions.
- Restriction-hours, which we used to quantify the magnitude of a MIT restriction, equals the product of the duration and the spacing requirement. Units are in mile-hours.
- Standard deviations (SDs) are computed in an analogous fashion.

The centers are sorted by decreasing average number of restrictions.

Based on these data, information about the MIT restrictions can be derived as illustrated in **Figure 84.** Note that the number of restrictions (count) is a reasonable surrogate for magnitude (note: magnitude is restriction-hours divided by 50, to scale it down for presentation purposes). That is, these curves follow each other fairly closely, except for ZTL. Size tends to decrease with count, that is, if a center needs more restrictions, then it probably needs more spacing too. The duration is fairly constant across the center (with notable exceptions: ZBW and ZJX have unusually lengthy restrictions). Note: duration values are in 15-minute increments, to make it scale with the rest of the curves. Finally, ZOB has unusually large restrictions (in miles).

Table 7. Aggregate MIT restrictions statistics center by center for 2000.

Center	Ave. Number Per Day	SD Number Per Day	Ave. Duration (Hours)	SD Duration (Hours)	Ave. Size (Miles)	SD Size (Miles)	Ave. Restriction-Hours Per Day (Mile-Hours)	SD Restriction-Hours Per Day (Mile-Hours)
ZNY	27.99	15.71	2.41	1.34	17.33	5.67	1179.67	29.04
ZDC	27.28	12.02	2.1	1.47	20.53	7.49	1264.2	40.61
ZAU	24.87	10.77	1.89	0.91	18.72	7.71	867.73	21.72
ZTL	22.61	12.15	1.36	1.33	16.4	3.51	541.04	35.29
ZOB	21.66	12.89	1.92	1.26	24.57	6.12	1031.13	33.03
ZID	12.19	9.17	1.35	1.23	20.61	5.34	346.27	26.73
ZLA	10.36	6.87	2.18	1.57	15.51	6.39	342.65	24.23
ZBW	5.5	4.18	2.67	1.04	16.78	3.79	246.39	21.34
ZME	5.41	4.77	1.43	0.95	18.17	5.15	146.08	22.93
ZHU	3.03	2.56	0.92	0.53	13.2	4.08	39.31	11.46
ZOA	2.92	2.55	1.69	0.66	14.16	4.67	70.92	13.34
ZJX	2.78	3.37	2.86	2.28	17.76	5.95	144.53	50.35
ZAB	2.64	2.54	1.43	1.2	13.81	5.2	54.54	21.11
ZFW	2.49	2.65	1.54	1.72	11.79	4.01	48.52	22.38
ZKC	1.96	2.35	1.58	0.79	19.2	4.96	61.09	20.21
ZMP	1.26	1.24	1.55	0.97	15.3	5.42	31.47	27.56
ZMA	0.57	1.11	2.44	1.47	12.12	4.1	17.48	23.47
ZDV	0.42	1.42	1.72	0.67	15.31	5.55	11.53	16.65
ZLC	0.2	0.86	1.83	1.06	14.27	6.96	5.75	22.78
ZSE	0.02	0.18	2.72	0.6	11.67	3.73	0.53	13.76

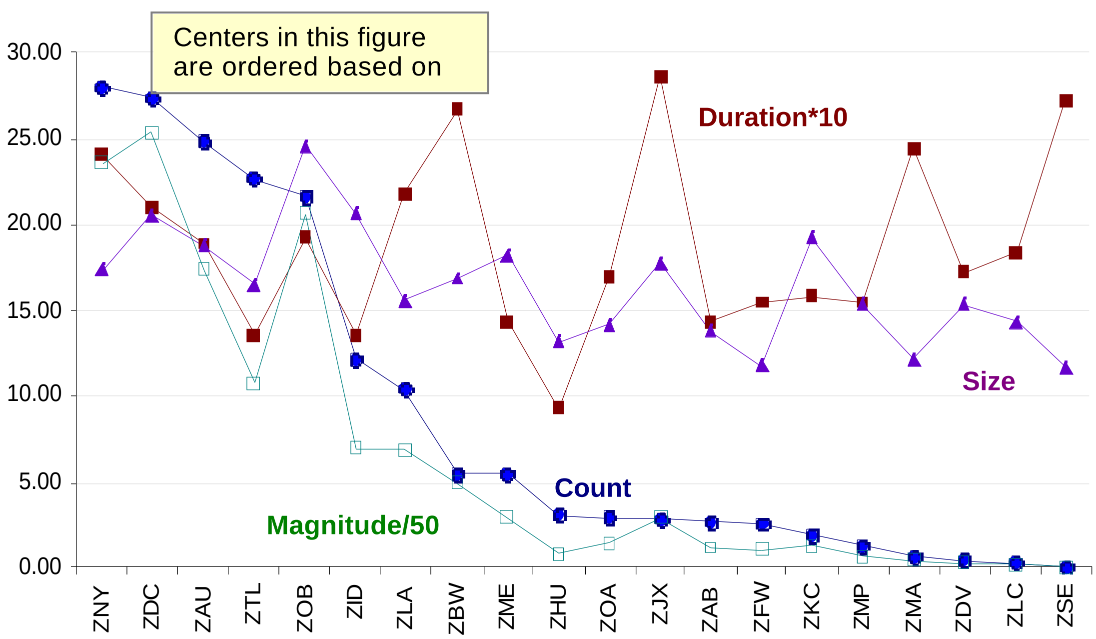

Figure 84. MIT data comparison for the year 2000.

Table 8. Aggregate MIT restrictions statistics center by center for 2001.

Center	Ave. Number Per Day	SD Number Per Day	Ave. Duration (Hours)	SD Duration (Hours)	Ave. Size (Miles)	SD Size (Miles)	Ave. Restriction-Hours Per Day (Mile-Hours)	SD Restriction-Hours Per Day (Mile-Hours)
ZOB	15.54	9.84	1.34	0.79	22.47	5.30	475.44	21.24
ZDC	14.65	11.47	2.00	1.21	21.22	6.79	657.35	35.24
ZTL	13.70	14.66	2.07	1.02	18.66	5.09	523.81	21.23
ZNY	10.84	8.31	1.27	0.80	16.70	3.49	239.73	17.25
ZAU	7.21	9.47	1.78	0.70	23.52	6.61	304.63	21.85
ZLA	6.41	5.53	2.03	0.81	15.80	6.49	207.13	19.94
ZID	5.57	5.32	1.24	0.78	20.92	5.93	147.74	21.26
ZOA	3.61	2.89	1.67	0.63	13.41	4.94	82.80	14.28
ZME	3.35	3.59	1.40	0.88	16.96	5.72	84.48	21.75
ZHU	3.06	2.59	0.94	0.60	14.12	4.22	43.77	12.88
ZJX	2.17	3.00	2.33	1.04	16.48	4.86	84.71	21.81
ZFW	1.92	2.29	1.20	0.73	11.04	3.68	27.28	11.98
ZBW	1.71	2.55	1.41	0.69	18.92	5.48	47.58	18.69
ZKC	1.49	2.37	2.25	0.88	18.16	4.49	60.36	19.20
ZAB	0.92	1.27	1.94	1.12	13.27	4.54	24.21	20.20
ZMA	0.91	1.71	1.50	0.83	15.44	5.50	21.69	16.70
ZMP	0.39	0.85	1.38	0.60	15.17	5.00	8.63	14.82
ZDV	0.16	0.63	1.87	0.63	18.73	7.73	5.56	17.10
ZSU	0.08	0.33	1.73	0.73	17.50	6.55	2.69	23.04
ZLC	0.01	0.10	1.50	0.00	20.00	0.00	0.16	0.00

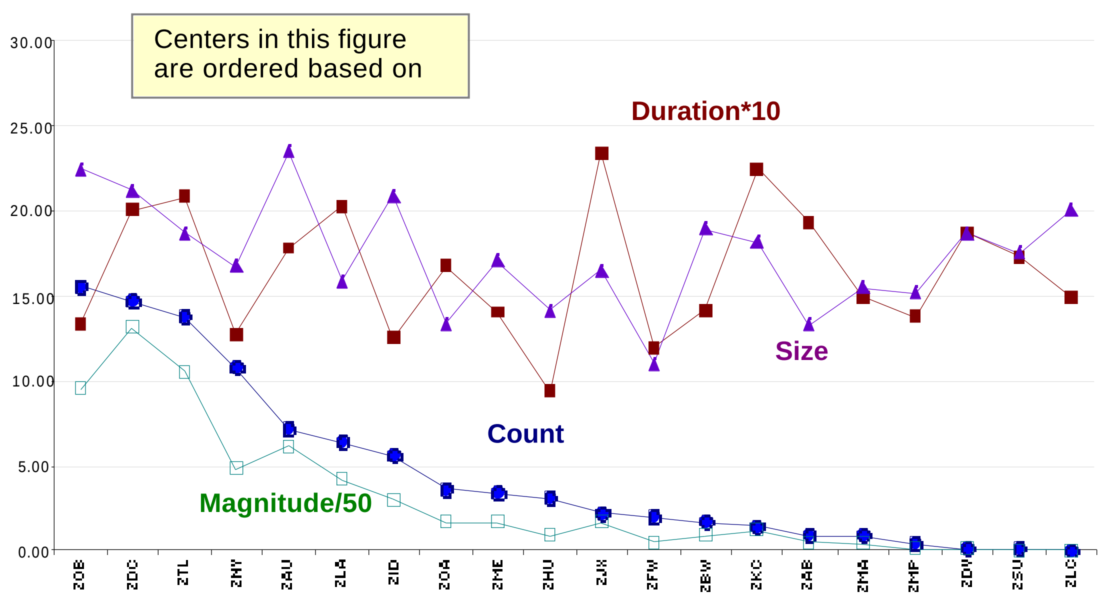

Figure 85. MIT data comparison for the year 2001.

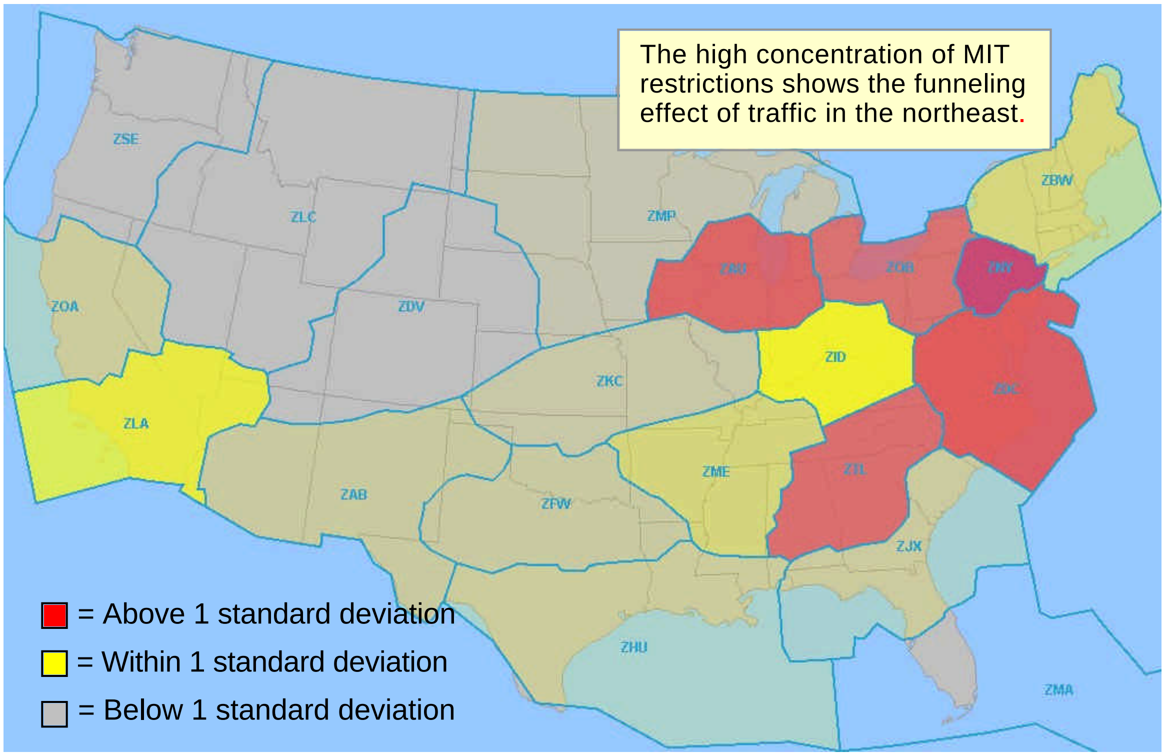

Figure 86. Geographical distribution for the average daily number of MIT restrictions per center for 2000.

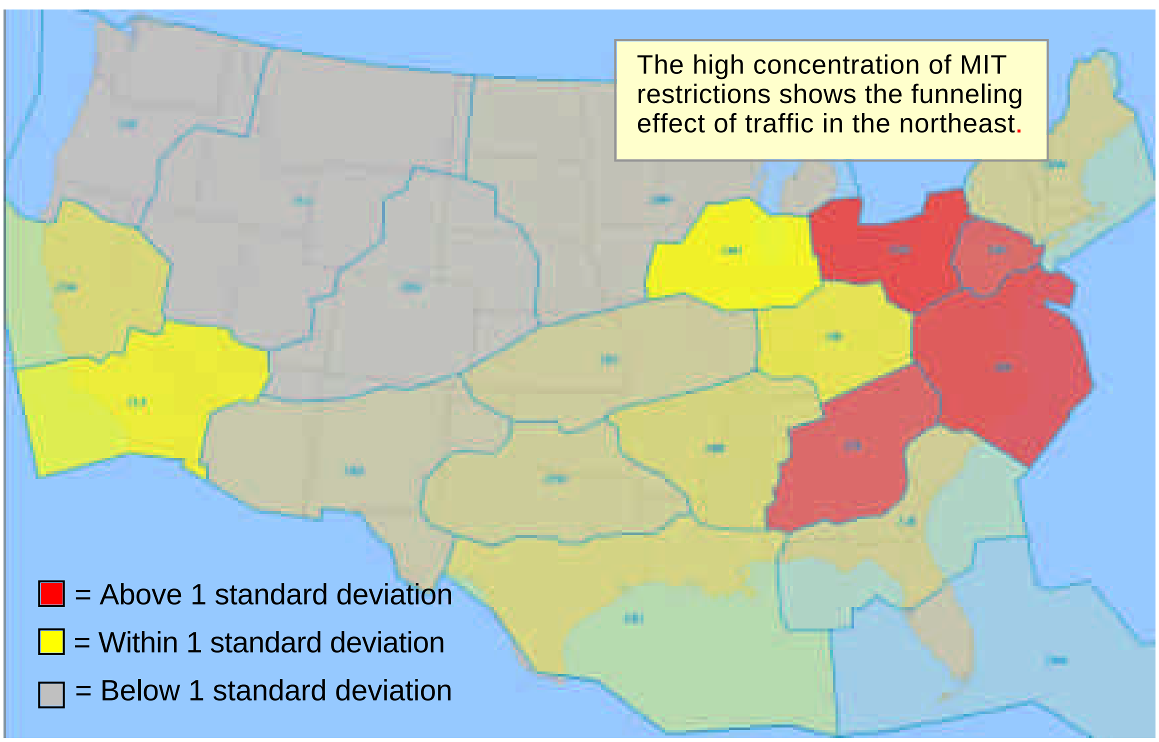

Figure 87. Geographical distribution for the average daily number of MIT restrictions per center for 2001.

4.2.5 Airborne Holding

Holding is used for en route aircraft when sudden changes in the traffic flow inhibit aircraft from landing in or from entering into neighboring centers (e.g., exiting ZOB into ZNY) or into airports in the TRACON area. Often, airborne holding is preceded by a GS or some other event, like a sudden change in MIT restrictions at a sector boundary. Thus, a GS or a MIT restriction change may be correlated with holding data.

Two ways to delay an aircraft en route are path stretching and speed reduction (or some combination thereof). Speed reduction is hard to easily detect, so we will limit our attention to path stretching, which takes multiple forms (e.g., **Figure 88**):

- Circular holding patterns – a race track type of loop usually at a designated fix location, and
- Vectoring, S-turns, and Path deviations – one or more repeated pattern of sharp, direction heading changes designed to “stretch the path” and thus delay traffic.

Detection of airborne holding from historical data is not a simple matter. In some parts of the country, databases are maintained of holding instances. Even these tend to be incomplete, since they are based on hand-written records of controllers; in periods of heavy demand, they may be too busy to record holding events. Thus, records can be incomplete when we would most like to have them.

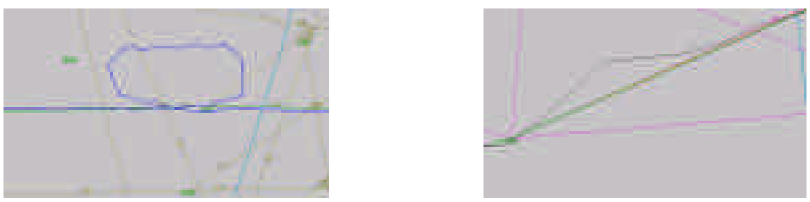

(a) Circular Holding *(b) Path Stretching*

Figure 88. Two example forms of holding.

Using flight track data, the easiest type of holding to detect (in an automated fashion) is circular holding patterns. We used the circular holding algorithm in POET to compile circular holding pattern statistics based on ETMS data. The algorithm seeks flight instances in which a flight track crosses over itself. The area of the holding loop must be greater than some minimum area (set to zero). Also, in order to avoid potential false positives in the terminal area (where a 1 minute update in the ETMS data may cause problems), the circular holding pattern loop must be at least 30 miles from the origin airport and at least 10 miles away from the destination airport. It is important to note that this algorithm has binary output: it simply says whether or not a flight endured holding under these criteria, not the number of circles nor the duration of holding.

Figure 89 illustrates the number of aircraft that experienced circular holding each day for year 2001 data. An insufficient amount of data was available during the course of this study to include a full set of data for the year 2000. Furthermore, there are 86 days worth of data missing from the 2001 data set. For this reason, circular holding pattern data was not included in the cluster analysis. Note that the number of planes that experience circular holding is variable throughout 2001 that appears to increase during the summer convective weather season months.

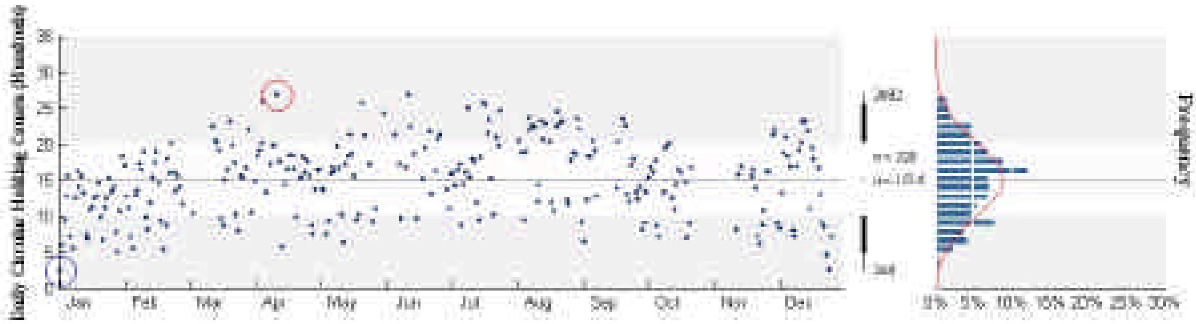

Figure 89. Distribution of circular holding for the year 2001.

A regression-based algorithm was used to infer terminal area holding. It works as follows. Once a flight takes off, ETMS makes periodic refinements of the estimated time of arrival (ETA). ETMS has knowledge of the current flight status (e.g., filed flight plan, current trajectory and speed), but has no knowledge of future delay events. Suppose that a flight were held by TFM for 15 minutes. In the absence of other forms of predictive interference, all of the ETA predictions made by ETMS prior to the holding event would be 15 minutes less than the actual arrival time, while all the ETA predictions made after the holding event would be equal to the actual arrival time. There would be a noticeable 'jump' in the value of the ETA predictions made during the time of the holding event. The amount of the jump would be commensurate with the amount of holding. An algorithm developed by Metron Aviation uses this basic principle to infer from ETA history how much holding a flight incurred, and the approximate start and stop times of the holding event. The algorithm is necessarily statistical, since there is inherent noise in the ETMS predictive model. Terminal area holding is inferred by considering only those holding events that occurred within a reasonable proximity to the arrival airport (e.g., 30 minutes). Such holding may include circular holding patterns, but is not limited to circular holding.

Figure 90 illustrates the number of aircraft that experienced terminal area holding each day for year 2001 data. The figure shows the total number of minutes (over all flights) of unanticipated delay that occurred within 30 minutes of arrival. The ASPM 50 airports were used in this analysis. One can clearly see the dip in September, associated with the September 11 tragedy. An insufficient amount of data was available during the course of this study to include a full set of data for the year 2000. For this reason, terminal area holding data was not included in the cluster analysis.

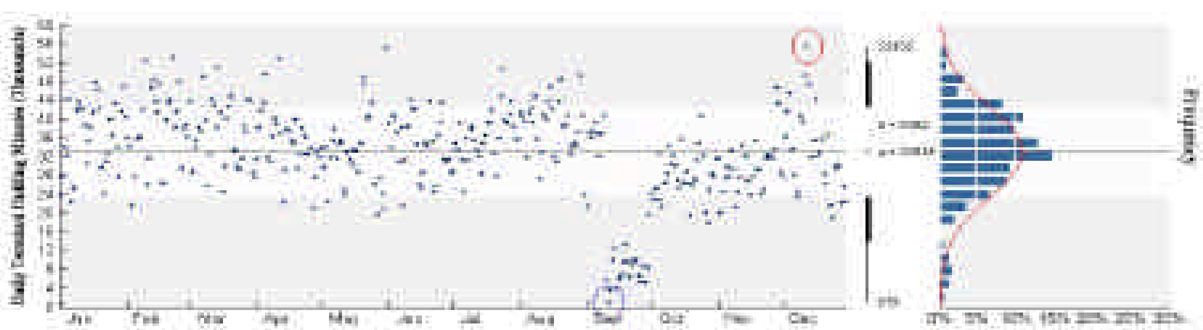

Figure 90. Distribution of terminal area holding in 2001.

4.3 Statistics of NAS Performance

The following performance statistics are reviewed next:

- OAG-Based Gate Departure Delays
- Ave. Airport Departure Delay
- Ave Gate Delay
- Ave. Taxi-Out Delay
- Ave. Airborne Delay
- Ave. Taxi In Delay
- Ave Arrival Delay
- Ave. Block Delay
- Total Delays
- Delays Caused by Weather
- ASPM Airport Performance Metric

Performance statistics describe delay in the NAS. The following discusses how average daily delay minutes were calculated using the *ASPM Quarter Hour Delays and Performance Scores* data. The total minutes of gate delay, taxi out delay, and airport departure delay were summed for each day and divided by the total count of departures as measured by ASPM. The total minutes of taxi in, airborne delay, block delay and OAG-based arrival delay were summed for each day and divided by the total count of arrivals for that day as measured by ASPM.

Taxi delay was computed in ASPM based on an estimated parameter that represents taxi time under optimal operating conditions. This parameter was based on aircraft queue lengths by carrier and airport, and was called unimpeded taxi-out time. Unimpeded taxi-out time is the taxi-out time with neither congestion, nor weather, nor any other factors, which could cause delay in an aircraft's movement from leaving the gate to the taking off. Aspects such as the combination of gates and runways, carriers at each airport, location of the carriers' gates relative to the runways that were used for takeoff, and seasonal considerations were all taken into account.

4.3.1 OAG-based Gate Departure Delay

OAG-based gate departure delays, as shown in **Figure 91** through **Figure 94**, are based on ASPM data. Gate delay is defined as the actual gate departure time minus the OAG-scheduled departure time.

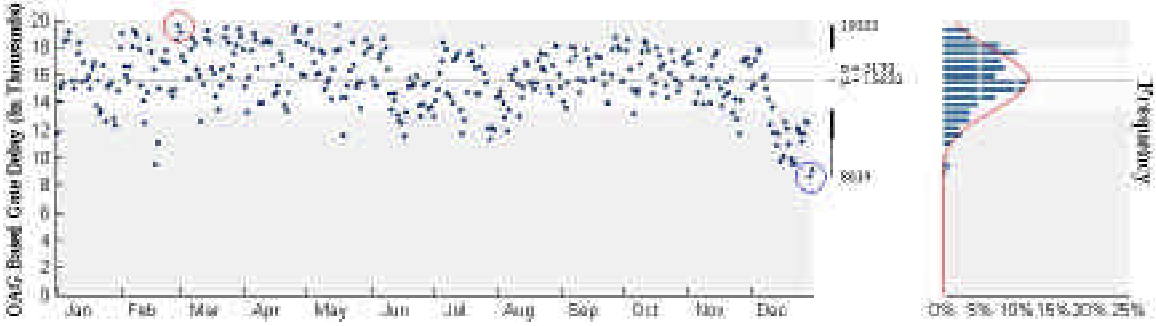

Figure 91. Total minutes of gate departure delay for domestic flights in 2000.

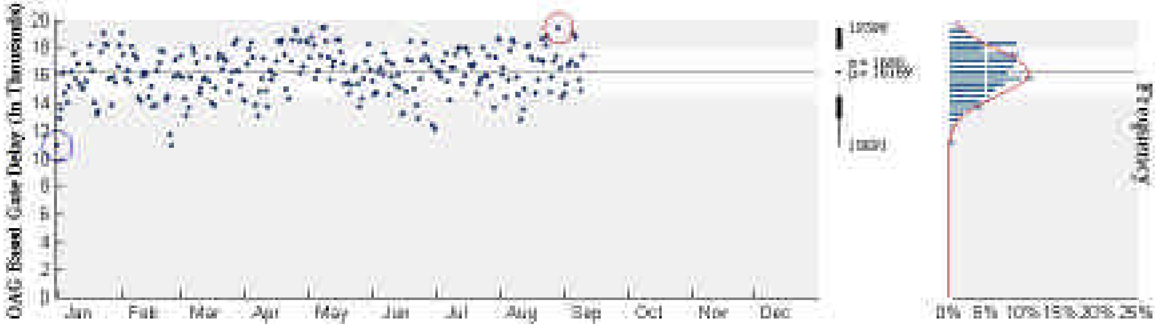

Figure 92. Total minutes of gate departure delay for domestic flights up to Sept. 2001.

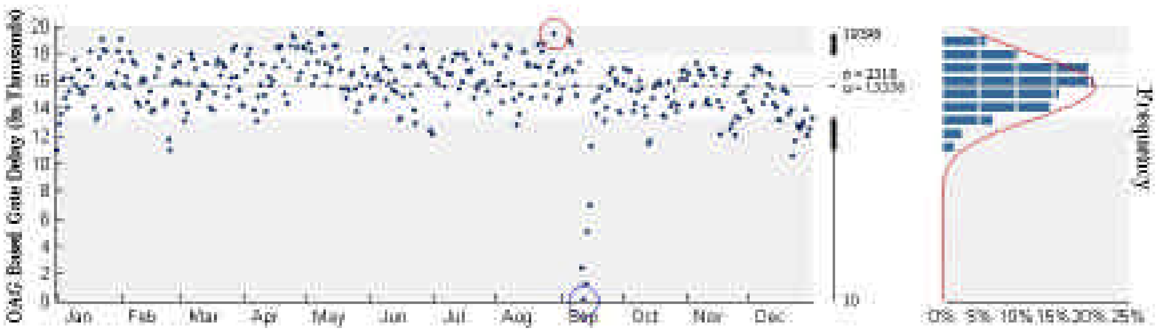

Figure 93. Total minutes of gate departure delay for domestic flights in 2001.

Figure 94. Total minutes of gate departure delay for domestic flights in 2002.

4.3.2 Average Airport Departure Delay

Airport Departure Delay is defined as the actual time off, minus the scheduled gate departure time and the unimpeded taxi-out time. **Figure 95** through **Figure 98** presents the airport departure delay statistics.

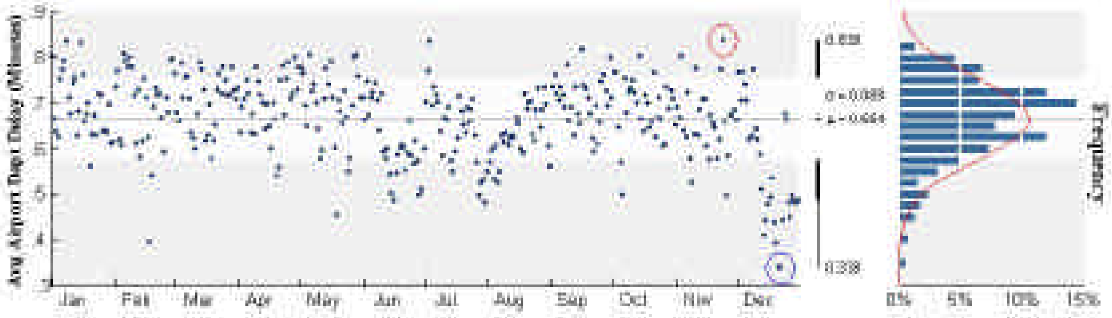

Figure 95. Distributions for airport departure delay for domestic flights in 2000.

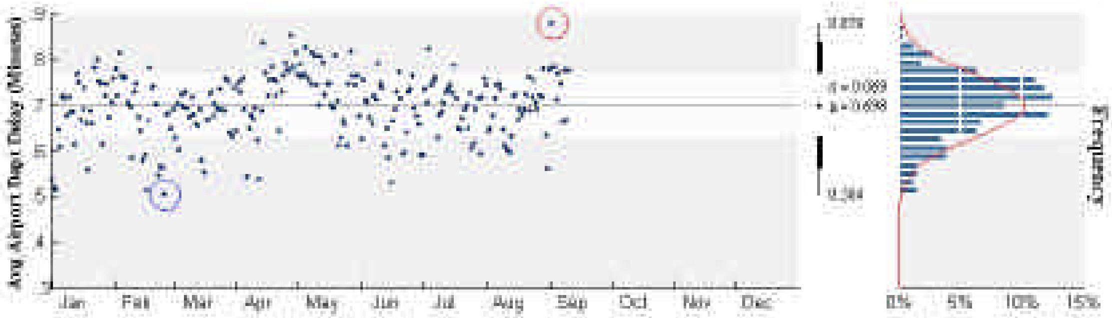

Figure 96. Distributions for airport departure delay for domestic flights up to Sept. 2001.

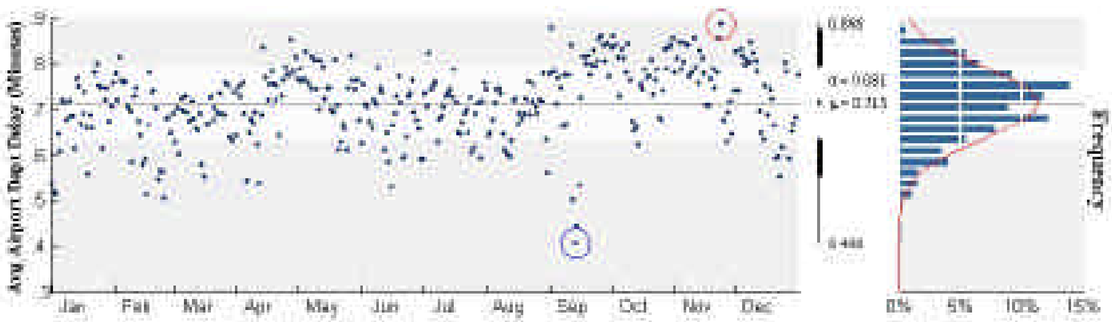

Figure 97. Distributions for airport departure delay for domestic flights in 2001.

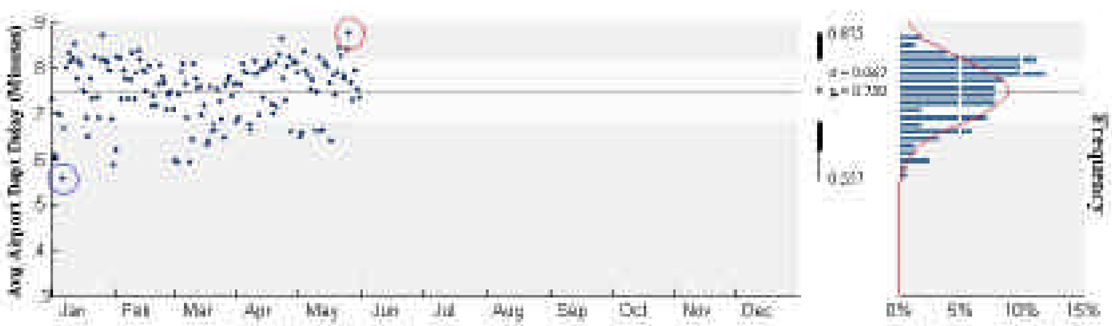

Figure 98. Distributions for airport departure delay for domestic flights in 2002.

4.3.3 Average Gate Delay

Average Gate Delay statistics, as shown in **Figure 99** through **Figure 102** are defined as the actual gate departure times, minus the scheduled gate departure times.

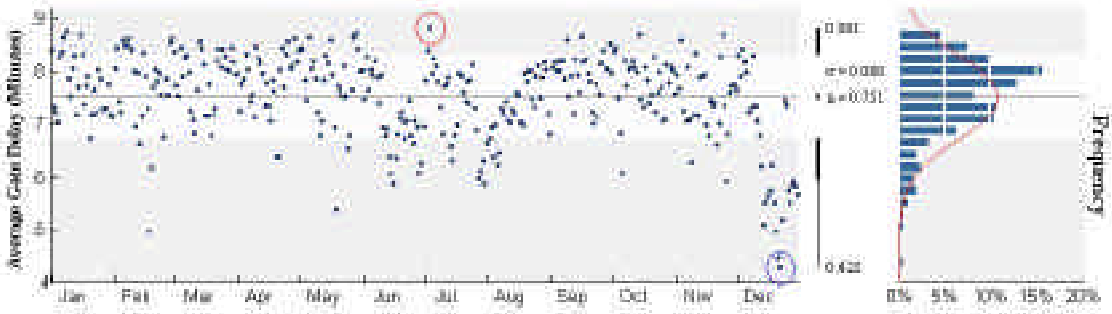

Figure 99. Distributions for airport gate delay for domestic flights in 2000.

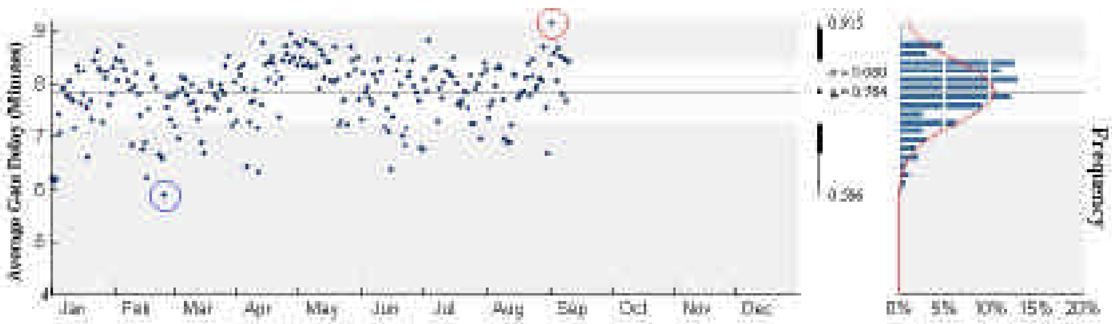

Figure 100. Distributions for airport gate delay for domestic flights up to Sept. 2001.

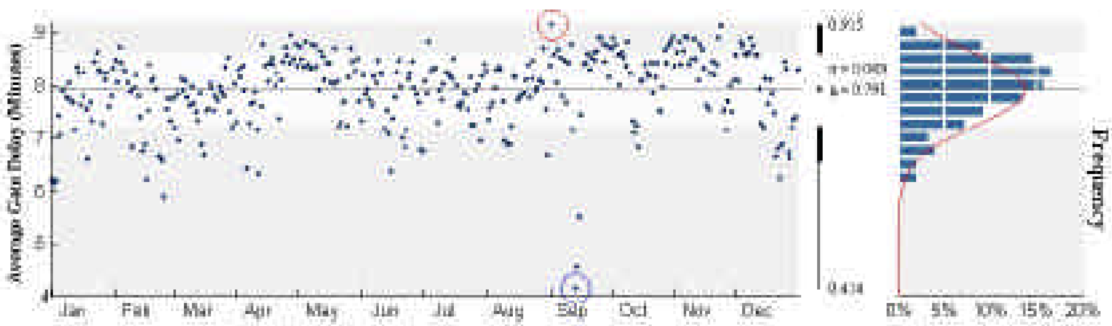

Figure 101. Distributions for airport gate delay for domestic flights in 2001.

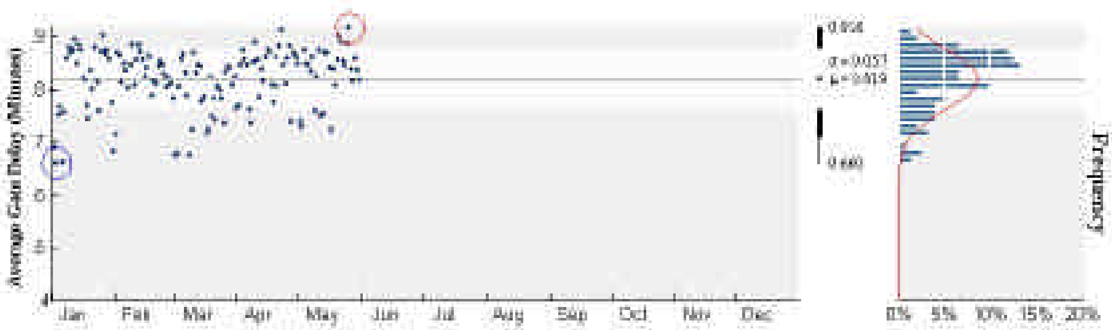

Figure 102. Distributions for airport gate delay for domestic flights in 2002.

4.3.4 Average Taxi-Out Delay

ASPM data are used to investigate average taxi-out delay, as illustrated in **Figure 103** through **Figure 106**.

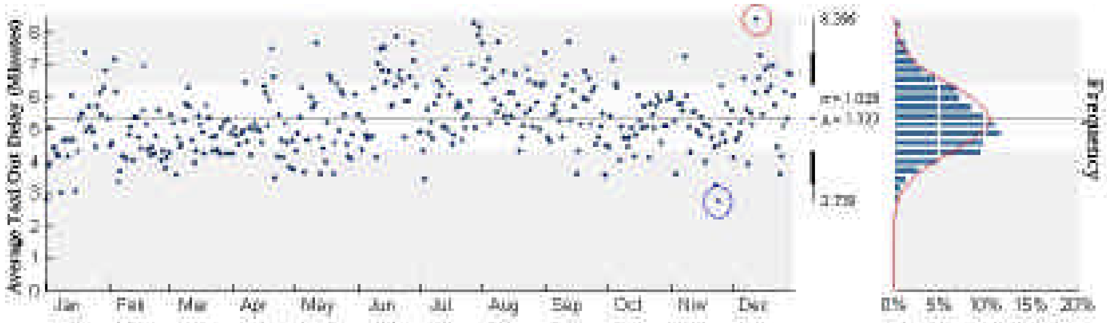

Figure 103. Distributions for average taxi out delay for domestic flights in 2000.

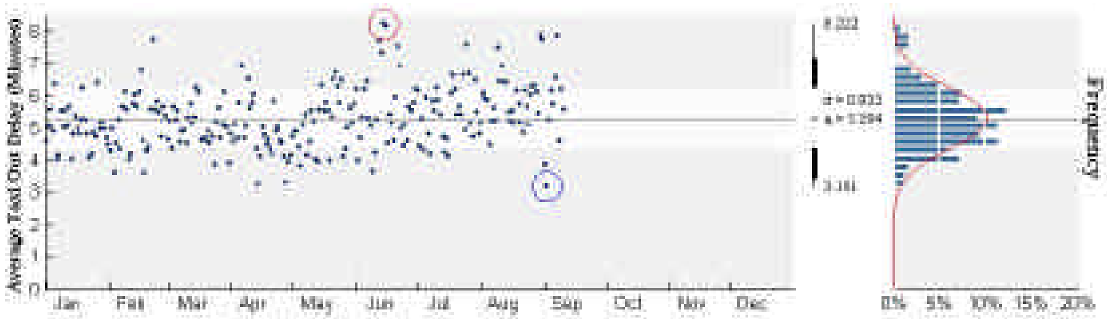

Figure 104. Distributions for average taxi out delay for domestic flights up to Sept. 2001.

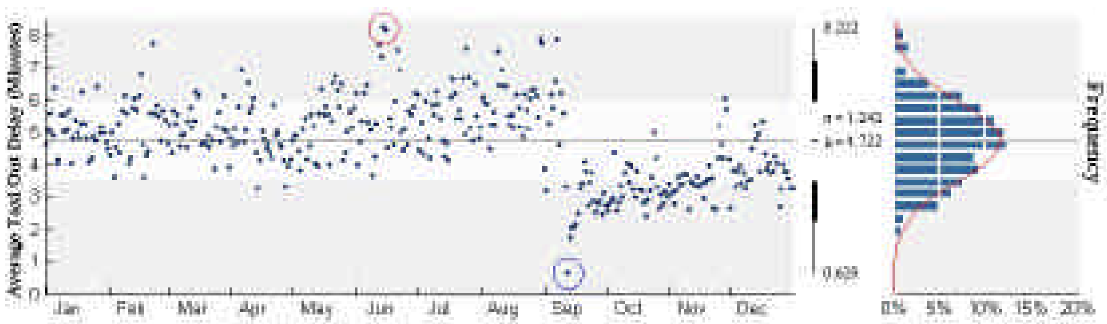

Figure 105. Distributions for average taxi out delay for domestic flights in 2001.

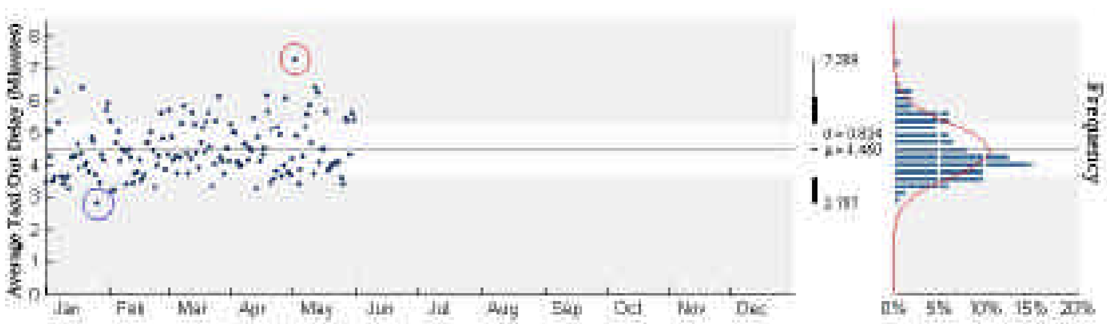

Figure 106. Distributions for average taxi out delay for domestic flights in 2002.

4.3.5 Average Airborne Delay

ASPM is used to identify airborne delay, as illustrated in **Figure 107** through **Figure 110**. Airborne delay is defined as the actual airborne time minus the carrier submitted time en route.

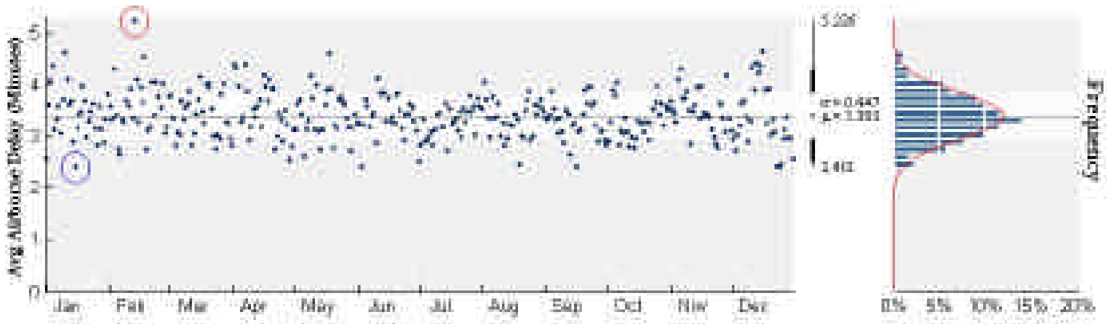

Figure 107. Distributions for airborne delay for domestic flights in 2000.

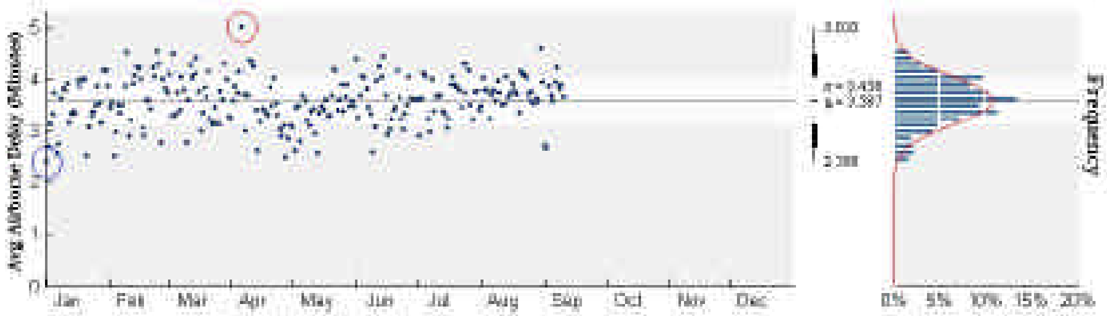

Figure 108. Distributions for airborne delay for domestic flights up to Sept. 2001.

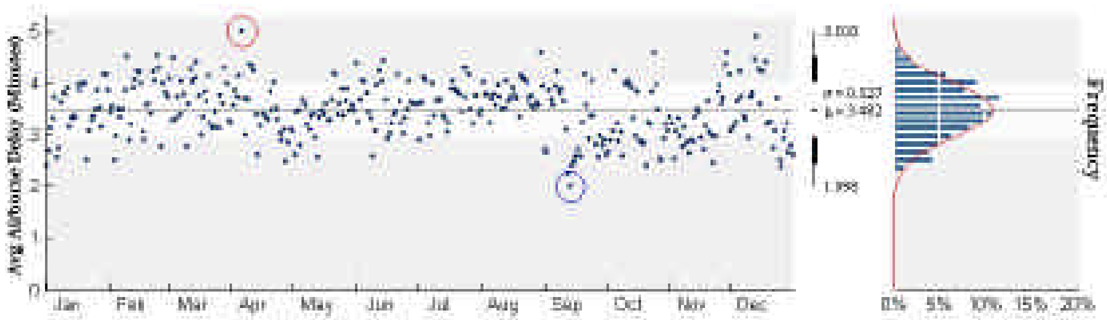

Figure 109. Distributions for airborne delay for domestic flights in 2001.

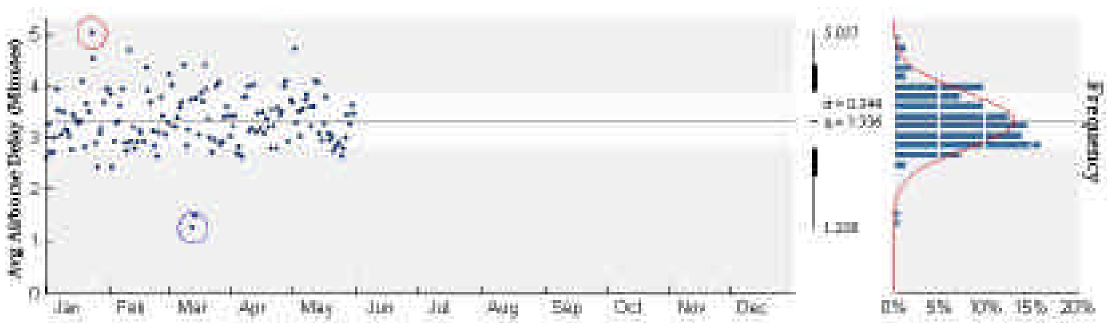

Figure 110. Distributions for airborne delay for domestic flights in 2002.

4.3.6 Average Taxi-In Delay

ASPM data are used to investigate taxi-in delay, as illustrated in **Figure 111** through **Figure 114**. Taxi-In Delay is defined as actual Taxi-In time minus the ASPM defined Unimpeded Taxi-In time.

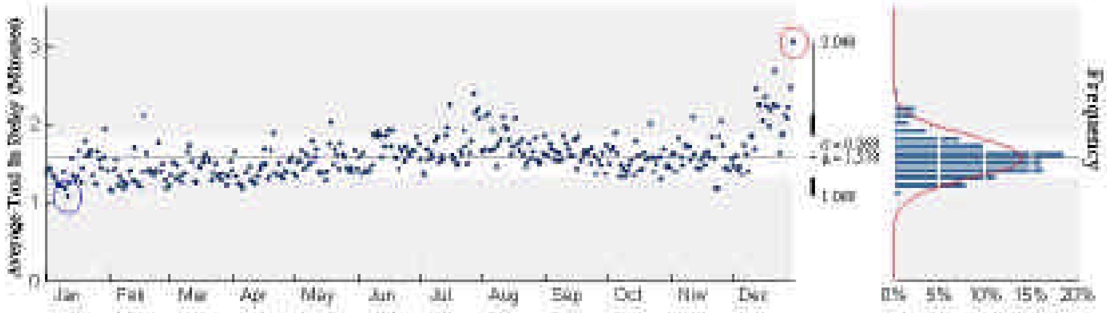

Figure 111. Distributions for taxi in delay for domestic flights in 2000.

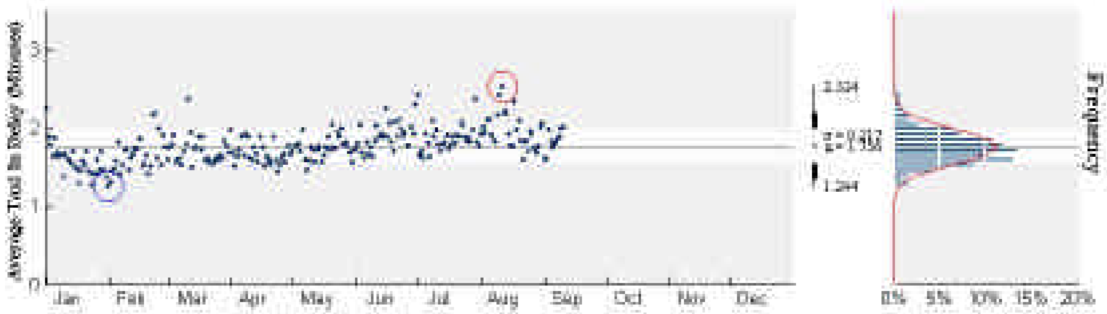

Figure 112. Distributions for taxi in delay for domestic flights up to Sept. 2001.

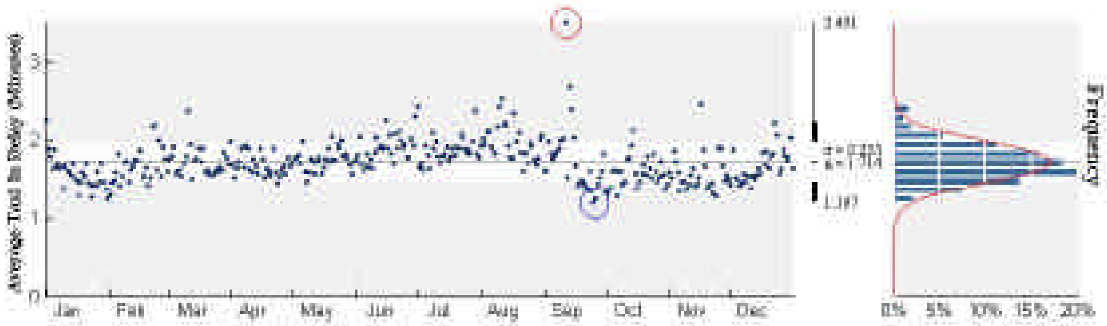

Figure 113. Distributions for taxi in delay for domestic flights in 2001.

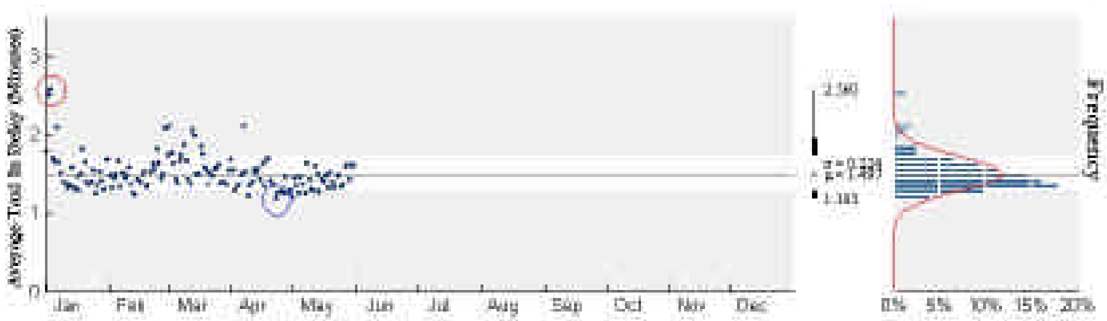

Figure 114. Distributions for taxi in delay for domestic flights in 2002.

4.3.7 Average Arrival Delay

The average airport arrival delay is the average delay for gate-in, defined as the actual gate-in minus the scheduled gate-in. Average arrival delay statistics are illustrated in **Figure 115** through **Figure 118**.

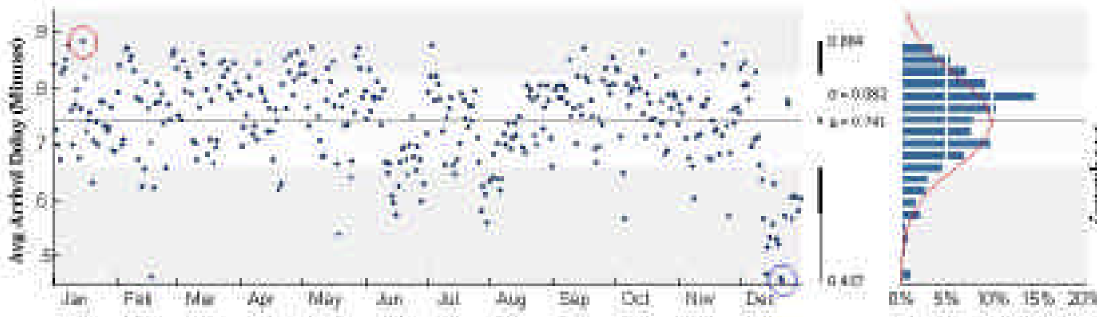

Figure 115. Distributions for average arrival delay for domestic flights in 2000.

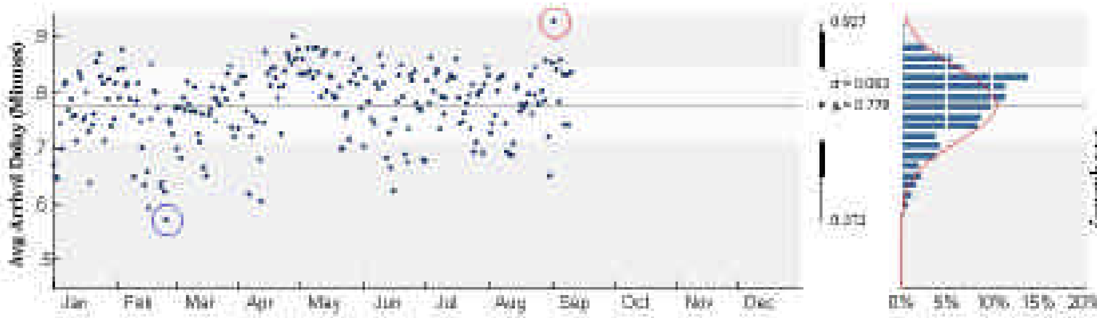

Figure 116. Distributions for average arrival delay for domestic flights up to Sept. 2001.

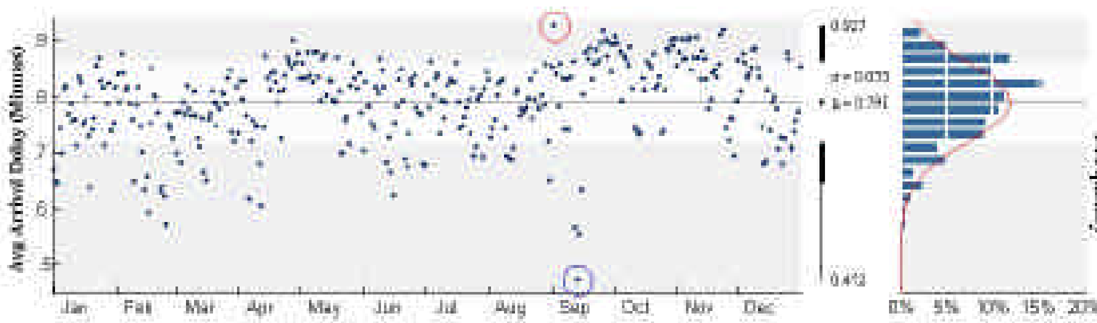

Figure 117. Distributions for average arrival delay for domestic flights in 2001.

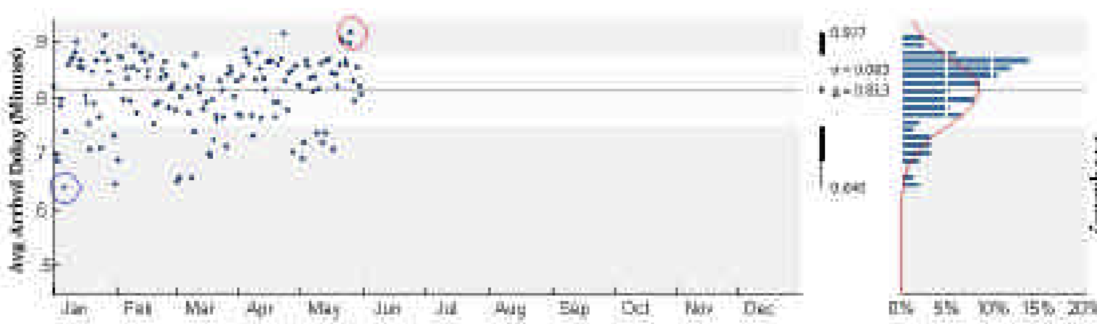

Figure 118. Distributions for average arrival delay for domestic flights in 2002.

4.3.8 Comparison of Taxi-In and Taxi-Out Delay

ASPM data are used to compare total taxi-in and taxi-out delay minutes per day. **Figure 119** through **Figure 122** illustrate taxi-in and taxi-out delay comparisons. The vast difference in the average and standard deviation of the two suggests that the root causes are vastly different. The vast difference in the average and standard deviation of the two suggests that the root causes are vastly different. Fundamentally, the differences (in mean and standard deviation) arise from the fact that there tends to be substantial queuing for departures (to keep pressure on the runways) as they wait for gaps in the arrivals stream which leads to the increases in Taxi Out delays. Departures are much more subject to NAS constraints (e.g., MIT restrictions) as well as interactions with arrival flows in terms of sharing runways. By comparison, unless an arrival has to cross the runway or wait for a gate to be available, it usually proceeds nearly directly to its gate at or near an "unimpeded-like" Taxi In time. Thus, the mean is lower and the variance is quite small.

As far as why some days the Taxi Out delay is on the same order as the Taxi In delay, there can be several reasons. First, depending on the day of the week, the traffic loading might be quite less (e.g., Sunday vs. Thursday). Also, depending on the "dynamics" of the rest of the NAS, the phasing of banking operations at hub airports might be more or less optimal. When there are late arrivals, departures that might nominally be favored end up having to wait as the arrivals trickle in. For future research, it might be interesting to look separately at the behavior of hub airports (Chicago, Denver) vs. a constant pressure airport like LGA.

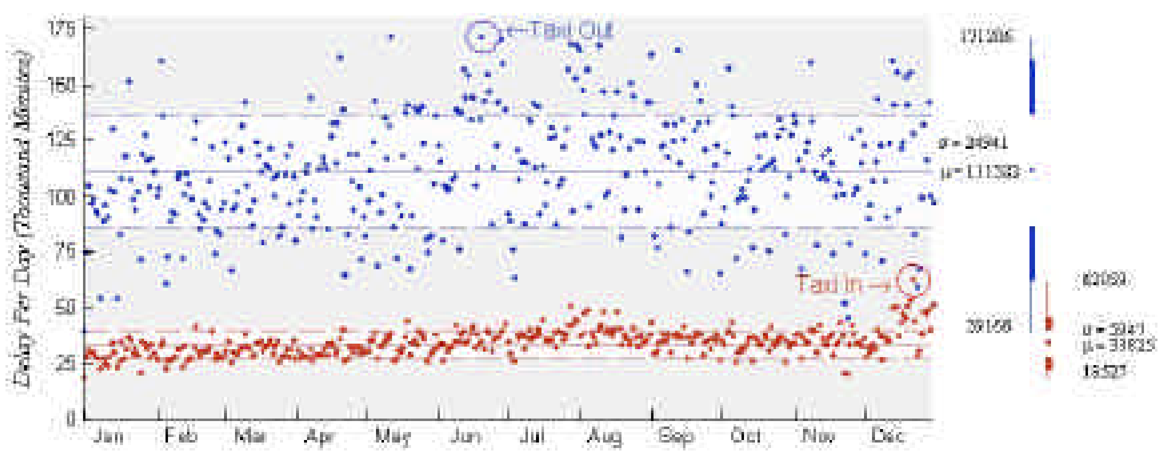

Figure 119. Distributions for taxi-in vs. taxi-out delay for domestic flights in 2000.

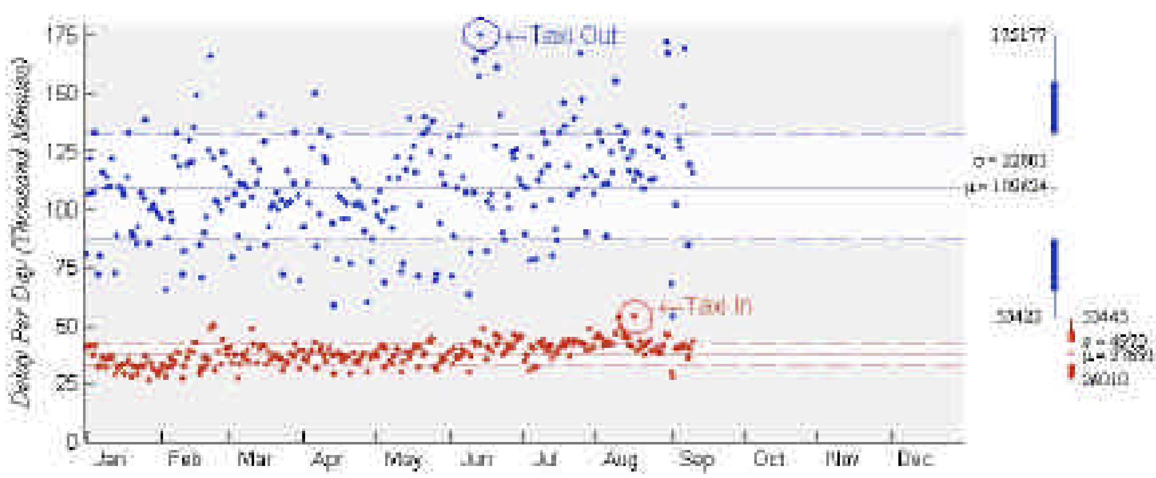

Figure 120. Comparison of taxi-in vs. taxi-out delay for domestic flights up to Sept. 2001.

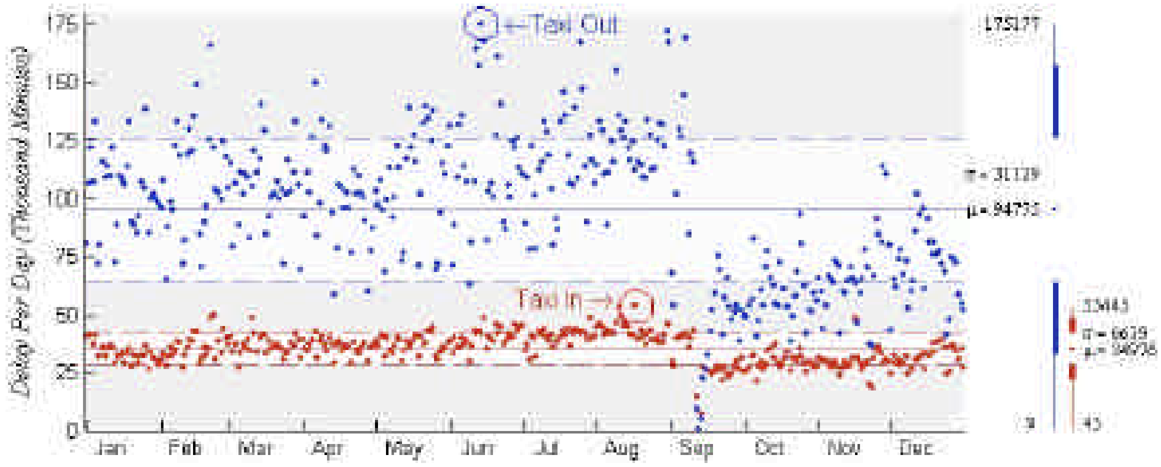

Figure 121. Comparison of taxi-in vs. taxi-out delay for domestic flights in 2001.

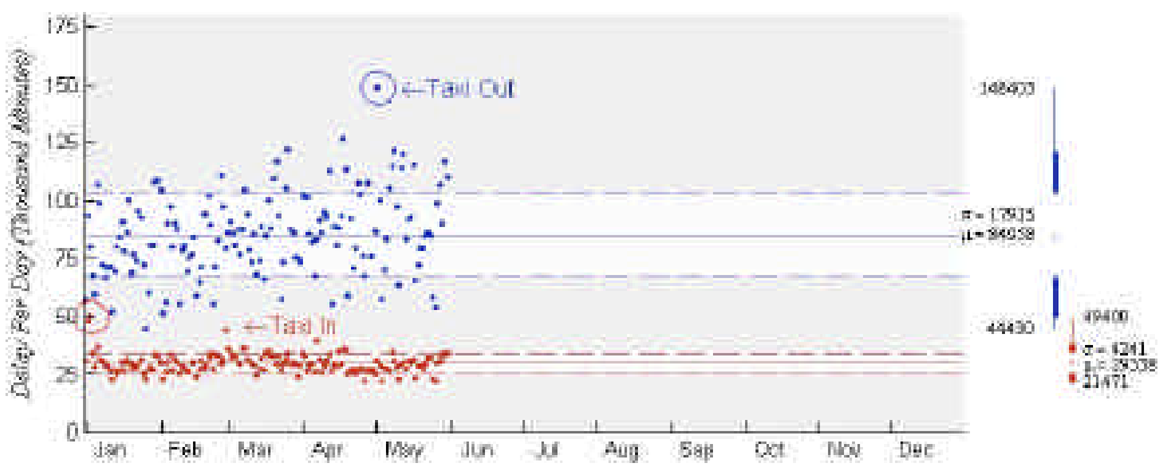

Figure 122. Comparison of taxi-in vs. taxi-out delay for domestic flights in 2002.

4.3.9 Average Block Delay

Block times span gate out to gate in. Block delay is defined as the actual Gate-to-Gate minus the scheduled Gate-to-Gate. Block delay statistics are presented in **Figure 123** through **Figure 126**.

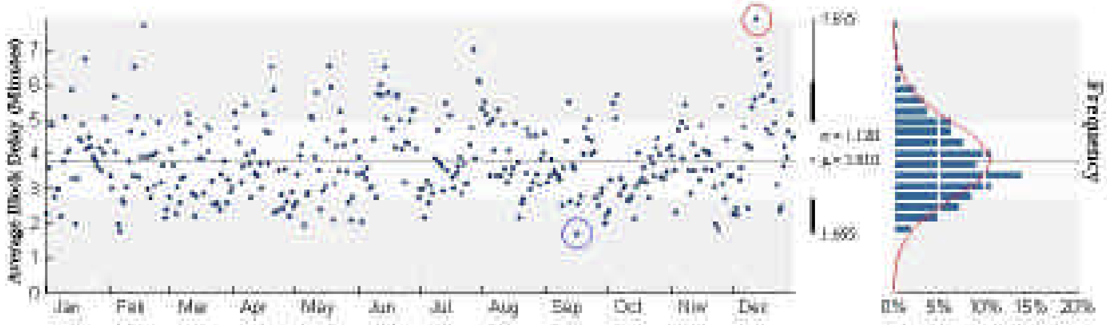

Figure 123. Distributions for block delay for domestic flights in 2000.

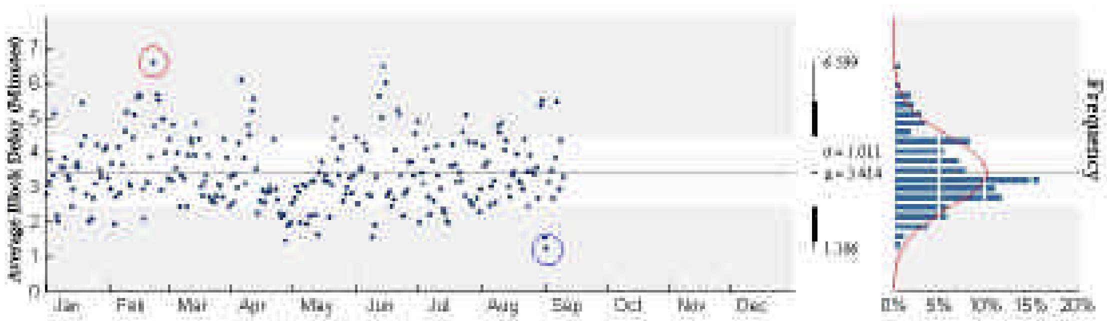

Figure 124. Distributions for block delay for domestic flights up to Sept. 2001.

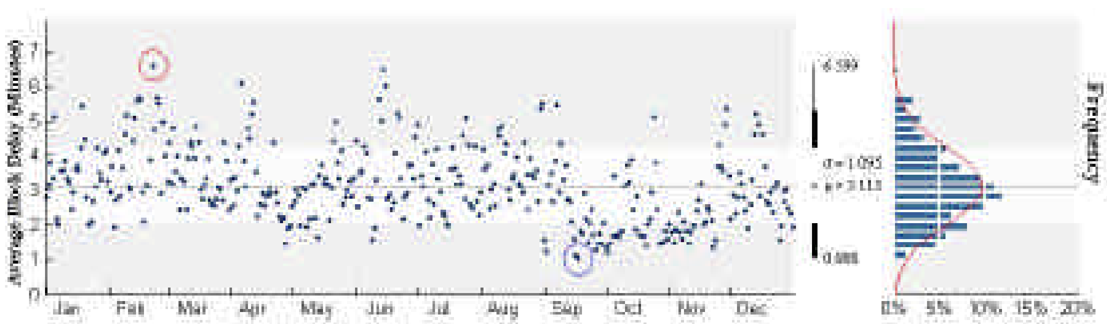

Figure 125. Distributions for block delay for domestic flights in 2001.

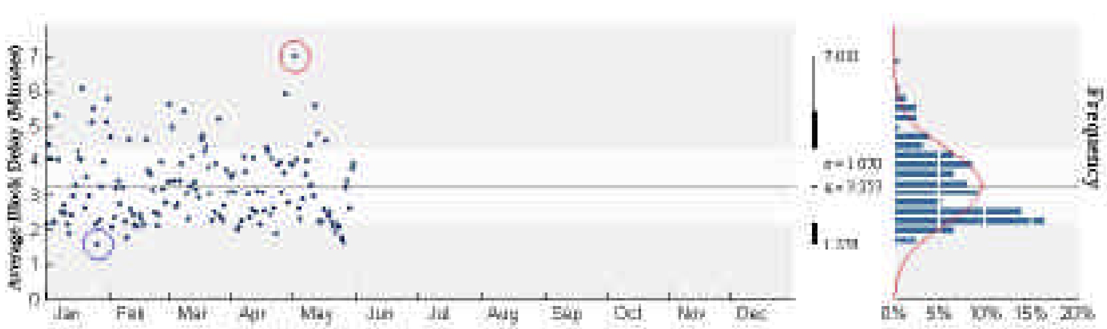

Figure 126. Distributions for block delay for domestic flights in 2002.

4.3.10 Total Delays

Total delays based on OPSNET data are presented in **Figure 127** through **Figure 130.** This is a count of operations which have reportable delay based on OPSNET standards.

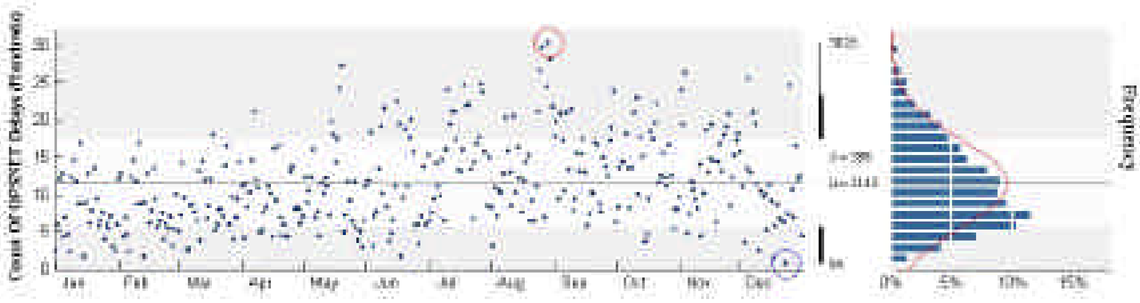

Figure 127. Total delayed operations in 2000.

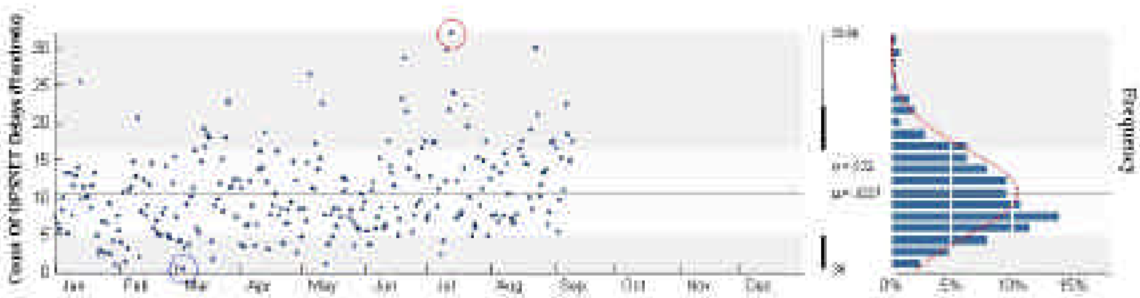

Figure 128. Total delayed operations in 2001 up to September, 2001.

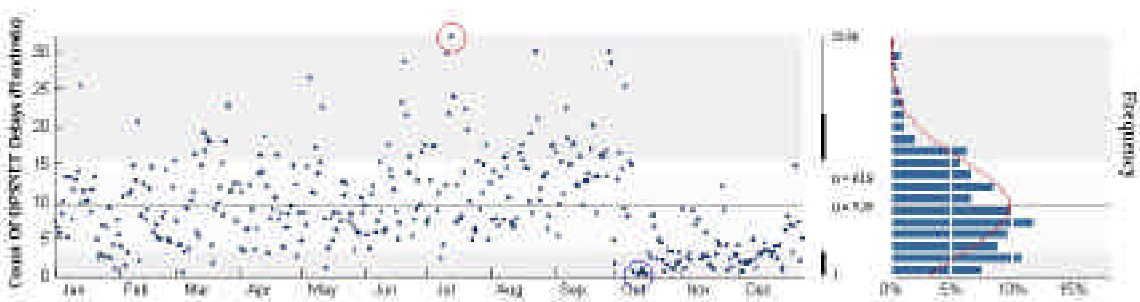

Figure 129. Total delayed operations in 2001.

Figure 130. Total delayed operations in 2002.

4.3.11 Delays Caused by Weather

One finding of our interview with ATCSCC subject matter experts was that the notion of a "typical" day in the NAS depends on the season. They specified that a typical day in the convective weather season (May through September) is quite different from a typical day in the non-convective weather season. Traffic flow initiatives are required more often in the convective weather season, and often under more stressful conditions. This is due to the increased air travel in the summer months and the unpredictable nature of convective weather activity.

Figure 131 confirms the dramatic increase in weather-related delays during the convective weather season. For this reason, we split our data sets into two groups: days within the convective weather season and days outside the convective weather season. Otherwise, natural variance within these two groups will undoubtedly complicate the results of the cluster analysis. A separate analysis could be done to refine the boundaries of the convective season, but we chose May 15 to September 15.

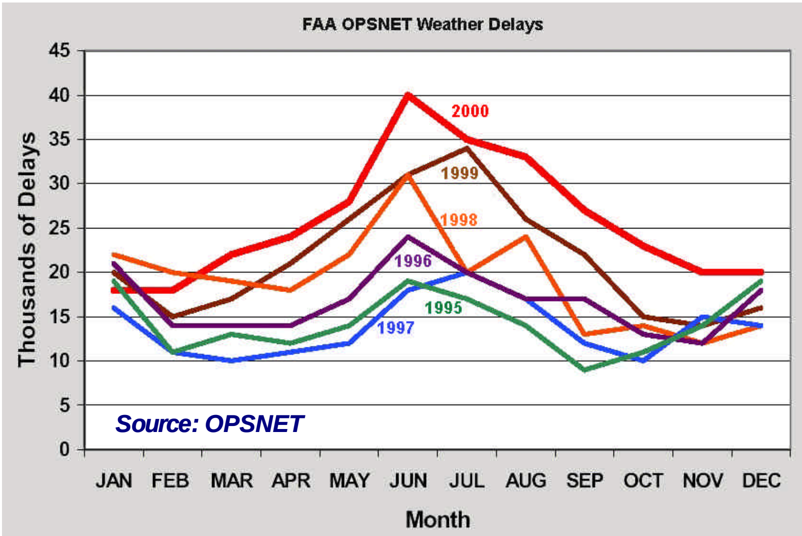

Figure 131. Comparison of yearly data for weather related delays.

When an airport is not operating in its optimum configuration, an impact condition is identified and recorded in OPSNET. OPSNET accepts a set of weather conditions that the user is responsible for determining. Data entry specifics are contained in the OPSNET User's Manual:

(1) Wind. Winds that cause less than optimum runway configuration, wind shear, or other adverse conditions.
(2) Rain. The presence of rain affecting the operating condition of runway(s) and/or taxiway(s).
(3) Snow and/or Ice. The presence of snow or ice affecting the operating condition of runway(s) and/or taxiway(s), including any combination of: (a) Poor or nil braking action; or (b) Snow/ice removal operations.
(4) Low Ceilings. Cloud conditions adversely affecting the operation, at or below takeoff, landing, or VFR requirements.
(5) Low Visibility. Reduced visibility adversely affecting the operation, at or below takeoff, landing, or VFR requirements.
(6) Tornado or Hurricane. The presence of a tornado or a hurricane.
(7) Thunderstorm. The presence of a thunderstorm.

Figure 132 through **Figure 135** present the total number of OPSNET weather related delays, which occurred during the period from January 1, 2000 to August 21, 2002. A 21-day (3 week) moving average is overlaid on each figure to give a feel for the trend of the data.

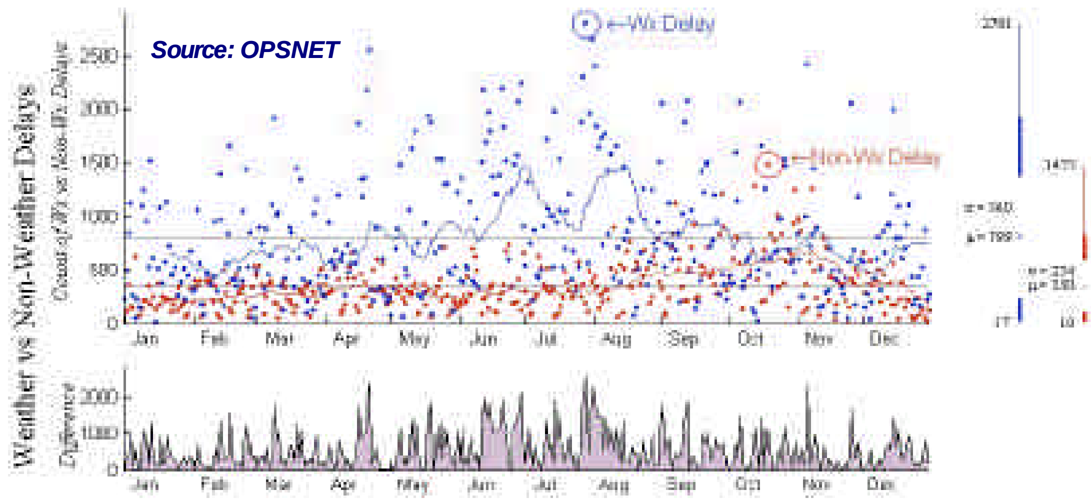

Figure 132. Weather Related Delays vs. Non- Weather Related Delays (and 21 day moving average) in 2000.

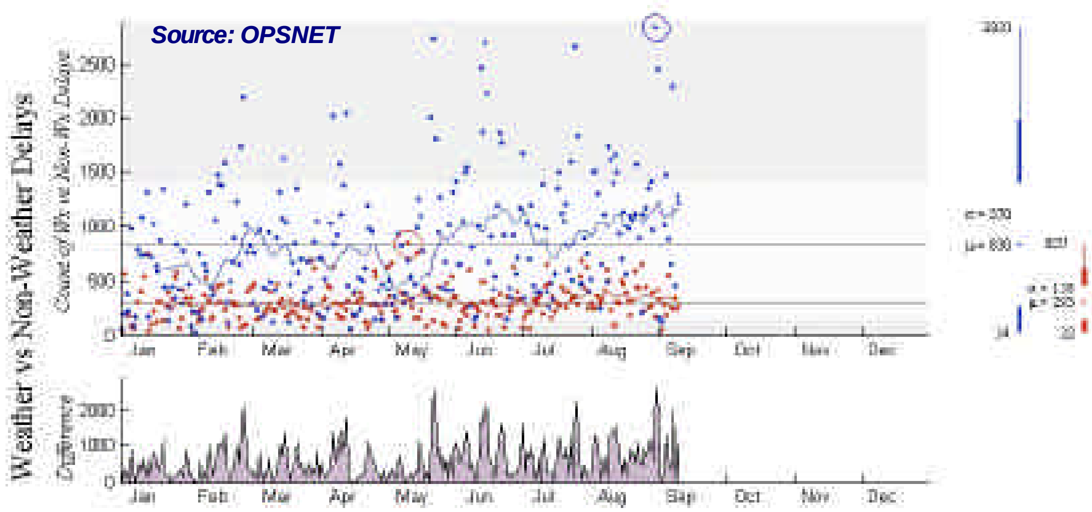

Figure 133. Weather Related Delays vs. Non-Weather Related Delays (and 21 day moving average) in 2001 prior to September 11.

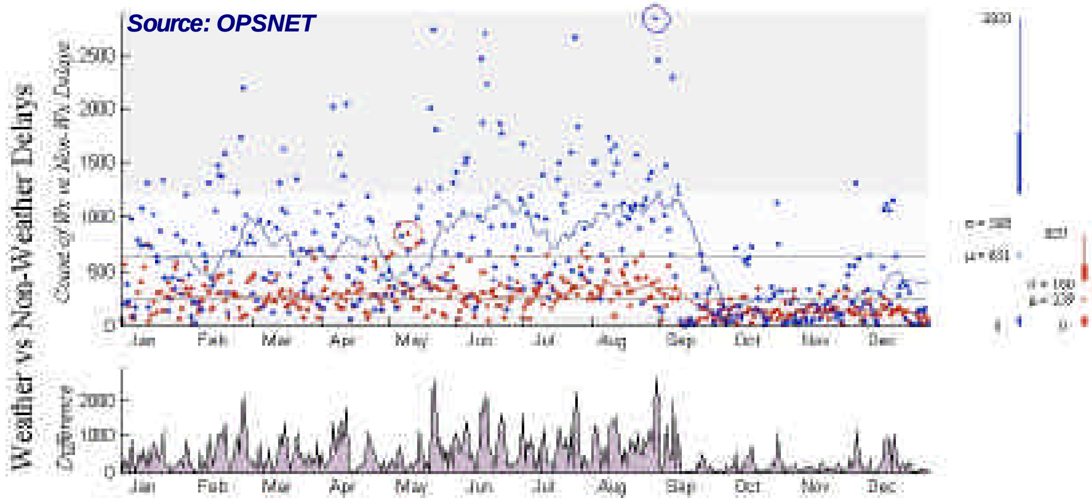

Figure 134. Weather Related Delays vs. Non-Weather Related Delays (and 21 day moving average) in 2001.

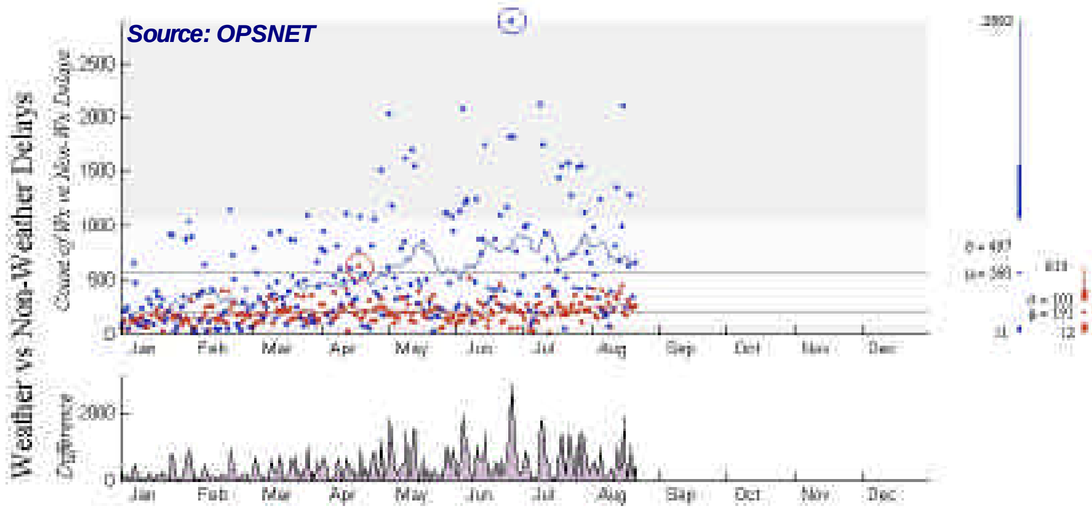

Figure 135. Weather Related Delays vs. Non-Weather Related Delays (and 21 day moving average) in 2002.

4.3.12 ASPM Airport Performance Metric

Definitions provided in the *Documentation for Airport Utilization Metrics* [FAA99, FAA02] were used in order to calculate daily airport performance scores for the ASPM 21 airports. These metrics were developed by MITRE's Center for Advanced Aviation System Development (CAASD) for measuring the performance of the NAS by measuring how well each airport's arrival and departure capacities are used when there is demand to be met. See [C01] for a complete explanation of the airport performance metric. This metric is intended to locate where further analysis could be done to identify contributing factors when performance is either exceptionally good or significantly less than optimal. It also takes into account the relative importance of meeting arrival and departure demand in each time period. **Figure 136** through **Figure 139** illustrates ASPM airport performance metric statistics.

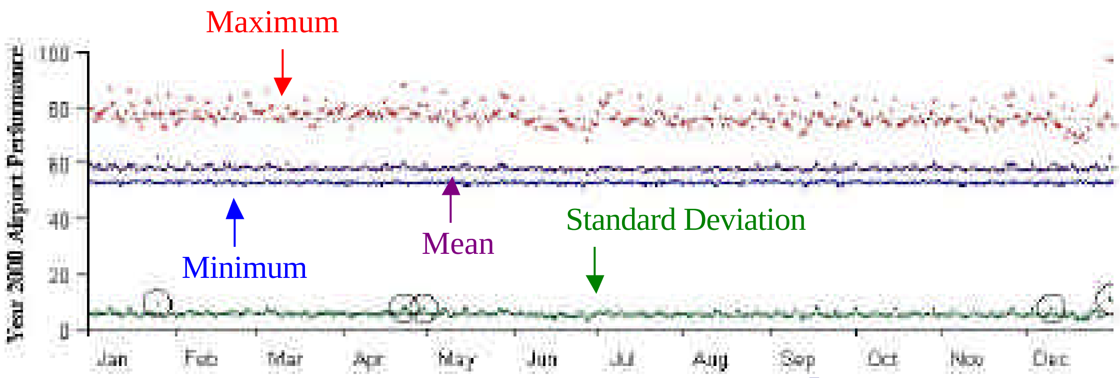

Figure 136. Airport Performance Scores[7] for 2000.

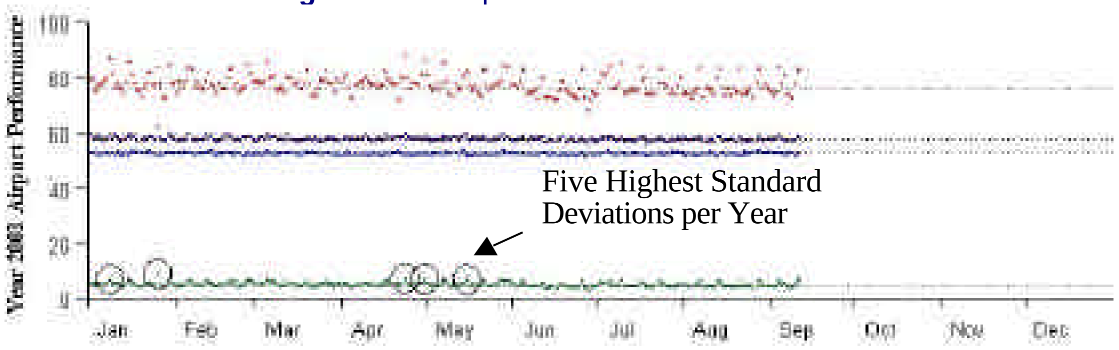

Figure 137. Airport Performance Scores for 2001 through September 10, 2001.

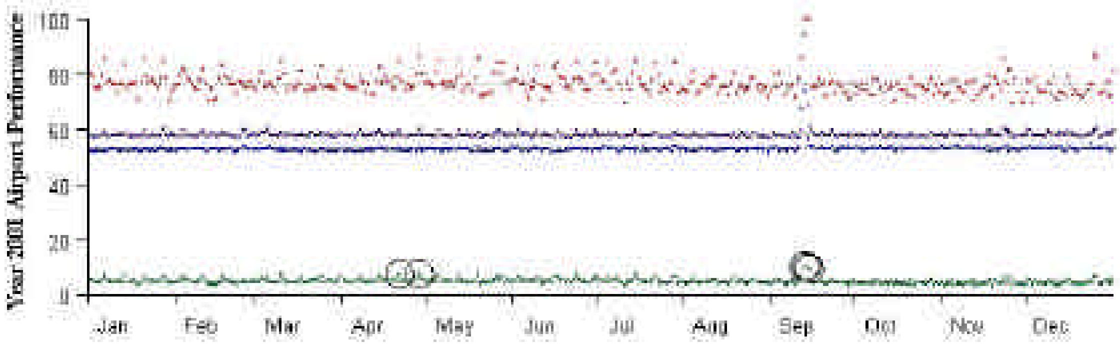

Figure 138. Airport Performance Scores for 2001.

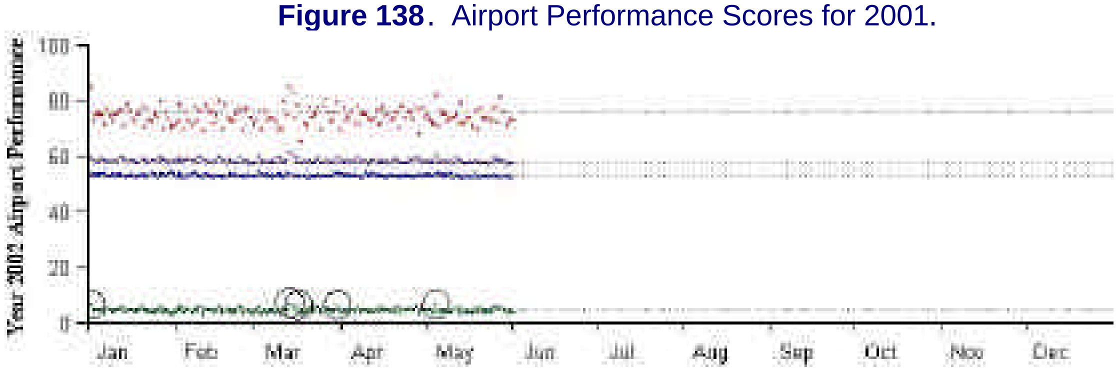

Figure 139. Airport Performance Scores for 2002.

[7] Statistics are out of 100 point scale and are determined over the ASPM 21 airports for each day.

5 Cluster Analysis

In this Chapter, we pursue two interests: the design of an optimal NAS feature vector, and the identification of several "types" of days in the NAS that capture major variations in overall operational behaviors of the NAS.

We design an "optimal" NAS feature vector with components that fully characterize the macroscopic behavior of the NAS on any given day. Statistical criteria are invoked to classify each component in the NAS feature vector as either typical or atypical of NAS performance. Whereas the previous chapter surveyed many potential variables to include into the NAS feature vector, we now wish to classify the NAS feature vector with as few variables as necessary. Too large a number of components may make the NAS feature vector difficult to interpret. Furthermore – as we show in this Chapter – we have reason to believe that many of the potential components of a NAS feature vector may be strongly correlated. For instance, if taxi-out times were unusually high, one would expect block delays to be high as well.

To understand variable dependencies, we use *cluster analysis* to partition the variables into groups so that within each group, the variables display similar behavior. From any of these cluster groups, one may select a single variable as representative. Or, variables within groups may be summed (e.g., number of GDPs + number of GSs), thus reducing the number of components to be considered for a NAS feature vector.

The second part of this Chapter investigates the classification of several "types" of days in the NAS. Once again we use a cluster analysis approach to pursue this objective. Given the "optimal" NAS feature vector as a basis, we investigate the collection of NAS feature vectors from Jan. 1, 2000 through Sept. 10, 2001, to identify the natural clusters of "types" of days. In the analysis, we did not force any partitions. Rather, the data naturally reveals certain "types" of days in the NAS. In particular, the analysis shows that weather and GDPs play an important role in determining the "types" of days in the NAS.

The term "cluster analysis" is necessarily broad and encompasses a wide variety of clustering algorithms. Even within a particular statistical endeavor, there can be many ways to cluster the data into meaningful groups. For instance, distance-based cluster algorithms map the variables into *n*-dimensional space, and then check for geometric proximity, using any of a number of metrics. As clusters develop, the trick is how to define distance between multiple objects. Some concept of cluster "center" must be applied. Nevertheless, cluster analysis is a mature science.

For the most part, clustering algorithms fall into one of the three categories:

- *Tree-based clustering*. In tree-based clustering, data are broken into groups, by successive branching on variables with distributions that demand partitioning. The end result is a tree graph such that all branches at a node point toward a sub tree with significantly different behavior from the node variable.
- *K-means clustering*. In K-means clustering, an algorithm moves objects in and out of groups (clusters) with the goal of maximizing the Analysis of Variance (ANOVA) significance. That is, there should be similar variance within each cluster but dissimilar variance between clusters. This method relies on the analyst's ability to specify the number of clusters in advance.
- *Two-way joining*. In two-way joining, clusters are formed for "cases" and variables at the same time. This assumes that there are multiple cases (e.g. patients in a clinic) for each variable.

In our analysis, we opted for a combination of tree-based clustering and K-means clustering. Our strategy thereby afforded us both robust control over and review of the clustering process. See [R71], [H75], or [G99] for a reference on cluster analysis.

The best way to understand cluster analysis is through analogy. Suppose someone asks you what your "typical" day is like. Specifically, what time do you get out of bed in the morning? Your immediate reply might be "weekday or weekend?" knowing that your habits differ on these two types of days; it would be meaningless to consider averages across these days. For if you wake at 7:00 am on weekdays, but sleep in until 10:00 am on weekends, then your average time of rising is 8:30 am. Clearly this statistic is misleading, because there may be virtually no day on which you get up at 8:30 am. Although this fulfills the notion of "typical" in some average sense it lacks a sense of frequency or modality. By contrast, the average rising time of a weekday has meaning in the "typical" sense, since rising times are fairly tightly grouped around 7:00 am on weekdays. Moreover, we have a strong chance of finding a typical weekday (historically speaking) in which you rose very close to 7:00 am. Overall, a major purpose of cluster analysis is to insure that within each cluster, data are actually present close to the mean, and with reasonable frequency.

The purpose of the cluster analysis across NAS feature vectors is to determine whether certain variables that should be split into multiple groups with similar behavior. In our rise-and-shine analogy, the cluster analysis would first detect the weekday-vs.-weekend difference and split out two groups before even asking the question what is typical. Moreover, the analysis might even identify a third category, called "vacation days" with highly erratic rising times. It may seem reasonable to add the vacation days to the weekend category, but the high variance might compromise the integrity of the weekend data set. The purpose of the cluster analysis is to determine from an objective, scientific standpoint whether or not this is a reasonable thing to do.

5.1 Data Associated with September 11, 2001

Data in this study are split distinctly before and after September 11, 2001. . It is well known that air traffic volumes precipitously dropped immediately after Sept. 11 (recall **Figure 1**). Clearly, Sept. 11 – and a few days thereafter – should be treated as a special event. What about the response and explanatory behavior of the other variables after September 11? Are the days after September 11 just low-volume instances of days prior to September 11? The answers to these questions help us determine whether we should restrict our attention to pre-September 11 data or to span over it.

Principal Component Analysis (PCA) is a statistical procedure that transforms a number of (possibly) correlated variables into a smaller number of uncorrelated variables called *principal components*. We adopted a form of PCA known as *oblique* PCA and *centroid-based clustering* [A73, H76] (available within SAS, S, and other statistical software). With these methods, we found that most of the candidate NAS feature vector variables contribute a similar variation in the pre- <u>and</u> post-September 11 datasets. Nevertheless, we also found some unique differences in the two datasets that require more detailed investigation outside the scope of this effort. So to be cautious, we proceeded only with the pre-September 11 dataset to obtain optimal clusters.

Even though most of the variables in the original dataset have similar effect toward explaining the observed variation in the pre- and post-September 11 data, some of the selected model variables had markedly different behavior in the pre and post era. For example, the scheduling groups of variables were quite different in the two periods. These differences were clearly visible in a tour of the data obtained from the dynamic graphics package X-gobi. Thus, in view of these strong differences we found it best to divide the data into pre and post September 11. A complete understanding of the differences between these two datasets remains an open research question outside the scope of the current effort.

5.2 Cluster Analysis Process and Results

A total of 65 variables were considered in the cluster analysis process. Although some of these variables were excluded from the statistical analysis survey presented in **Chapter 4,** they all are included in the cluster analysis to identify the most dominant variables for describing the optimal

NAS feature vector. **Table 9** presents the variables included in the analysis, and **Figure 141** shows a relative comparison of the variable ranges.

Table 9. Descriptive information about the variables considered in cluster analysis.

Number	Variable	Variable Description
1	Date	Dates from January 1, 2000 to May 31, 2002
2	OAG_DEP_CT	Daily OAG Scheduled Departure Count
3	OAG_ARR_CT	Daily OAG Scheduled Arrival Count
4	DEP_CNT	Daily Departure Count
5	ARR_CNT	Daily Arrival Count
6	DEP_CANCEL	Daily Departure Cancellations Count
7	ARR_CANCEL	Daily Arrival Cancellations Count
8	OGATE_DELC	Daily Count of OAG-Based Gate Delays
9	GDEP_ONTIM	Average Percent OAG-Based On Time Gate Departures (per quarter hour)
10	OARPT_DEPC	Daily Count of OAG-Based Airport Departure Delays
11	ADEP_ONTIM	Average Percent OAG-Based On time Airport Departures (per quarter hour)
12	ODELARR_C	Daily Count of OAG-Based Arrival Delays
13	GARR_ONTIM	Average Percent OAG-Based On Time Arrivals (per quarter hour)
14	OGATE_DEL	Daily Total OAG-Based Gate Delay (minutes)
15	DELAY_TO	Daily Total Taxi Out Delay (minutes)
16	O_ARPT_DEP	Daily Total OAG-Based Airport Departure Delay (minutes)
17	DEL_AIR	Daily Total Airborne Delay (minutes)
18	DEL_TI	Daily Total Taxi In Delay (minutes)
19	BLOCK_CNT	Daily Count Of Flights With Block Delay
20	Block_DEL	Daily Total Block Delay (minutes)
21	O_DEL_ARR	Daily Total OAG-Based Arrival Delay (minutes)
22	WND_SPEED	Maximum Reported Wind Speed (Knots)
23	AVGGATEDEL	Average Daily Gate Delay (minutes)
24	AVGDEL_TO	Average Daily Taxi Out Delay (minutes)
25	AVG_DEL_TI	Average Daily Taxi In Delay (minutes)
26	AVGDEL_AIR	Average Daily Airborne Delay (minutes)
27	AVG_BLOCK	Average Daily Block Delay (minutes)
28	AVGARPTDEP	Average Daily Airport Departure Delay (minutes)
29	AVGODELARR	Average Daily OAG-Based Airport Arrival Delay (Gate In) (minutes)
30	CEILING	Average Daily Ceiling
31	VISIBLE	Average Daily Visibility
32	IFR	Daily Percentage of quarter hours in IFR conditions
33	VFR	Daily Percentage of quarter hours in VFR conditions
34	AVG_ARR_RATE	Average Hourly AAR over all airports throughout the day
35	Avg_ARPT_PER	Average over 21 ASPM Airports of Daily Airport Performance Score
36	Min_ARPT_PER	Minimum over 21 ASPM Airports of Daily Airport Performance Score
37	Max_ARPT_PER	Maximum over 21 ASPM Airports of Daily Airport Performance Score
38	SD_ARPT_PER	Std. Deviation over 21 ASPM Airports of Daily Airport Performance Score

39	Var_ARPT_PER	Variance over 21 ASPM Airports of Daily Airport Performance Score
40	GDPs	Daily GDP Count
41	GDP_Length	Length of GDP (minutes)
42	GDP_Delay	Total Delay attributed to the GDP (minutes)
43	WxRel_MIT	Daily Count of Weather-Related MIT Restrictions
44	Volume_MIT	Daily Count of Volume-Related MIT Restrictions
45	Total_MIT	Daily Count of Total MIT Restrictions
46	Total_Ops	Total Operation Count From OPSNET
47	Total_Delays	Total Delay Count From OPSNET
48	Dep_Del	Total Departure Delay Count From OPSNET
49	Arr_Del	Total Arrival Delay Count From OPSNET
50	Enroute_Del	Total En route Delay Count From OPSNET
51	TMS_Del	Total Traffic Management Initiative Delay Count From OPSNET
52	AirCarrier_Del	Total Air Carrier Delay Count From OPSNET
53	AirTaxi_Del	Total Air Taxi Delay Count From OPSNET
54	GA_Del	Total General Aviation Delay Count From OPSNET
55	Military_Del	Total Military Delay Count From OPSNET
56	Wx_Del	Total Weather Related Delay Count From OPSNET
57	TermVol_Del	Total Terminal Volume Related Delay Count From OPSNET
58	CenterVol_Del	Total Center Volume Related Delay Count From OPSNET
59	Equip_Del	Total Equipment Related Delay Count From OPSNET
60	Runway_Del	Total Runway Related Delay Count From OPSNET
61	Other_Del	Total Other Delay Count From OPSNET (includes Noise Abatement, Flight Check, Fire, Bomb Threat, VIP Movement, Air Show, Aircraft Emergency, Stuck Mike, External Radio Frequency Interference, etc.)
62	DelsPer1000Ops	Delays per Thousand Operations from OPSNET
63	AvgDel_Min	Avg minutes of Delay from OPSNET
64	TotDel_Min	Total Minutes of Delay from OPSNET
65	PercOpsDelayed	Percentage of Operations Delayed from OPSNET

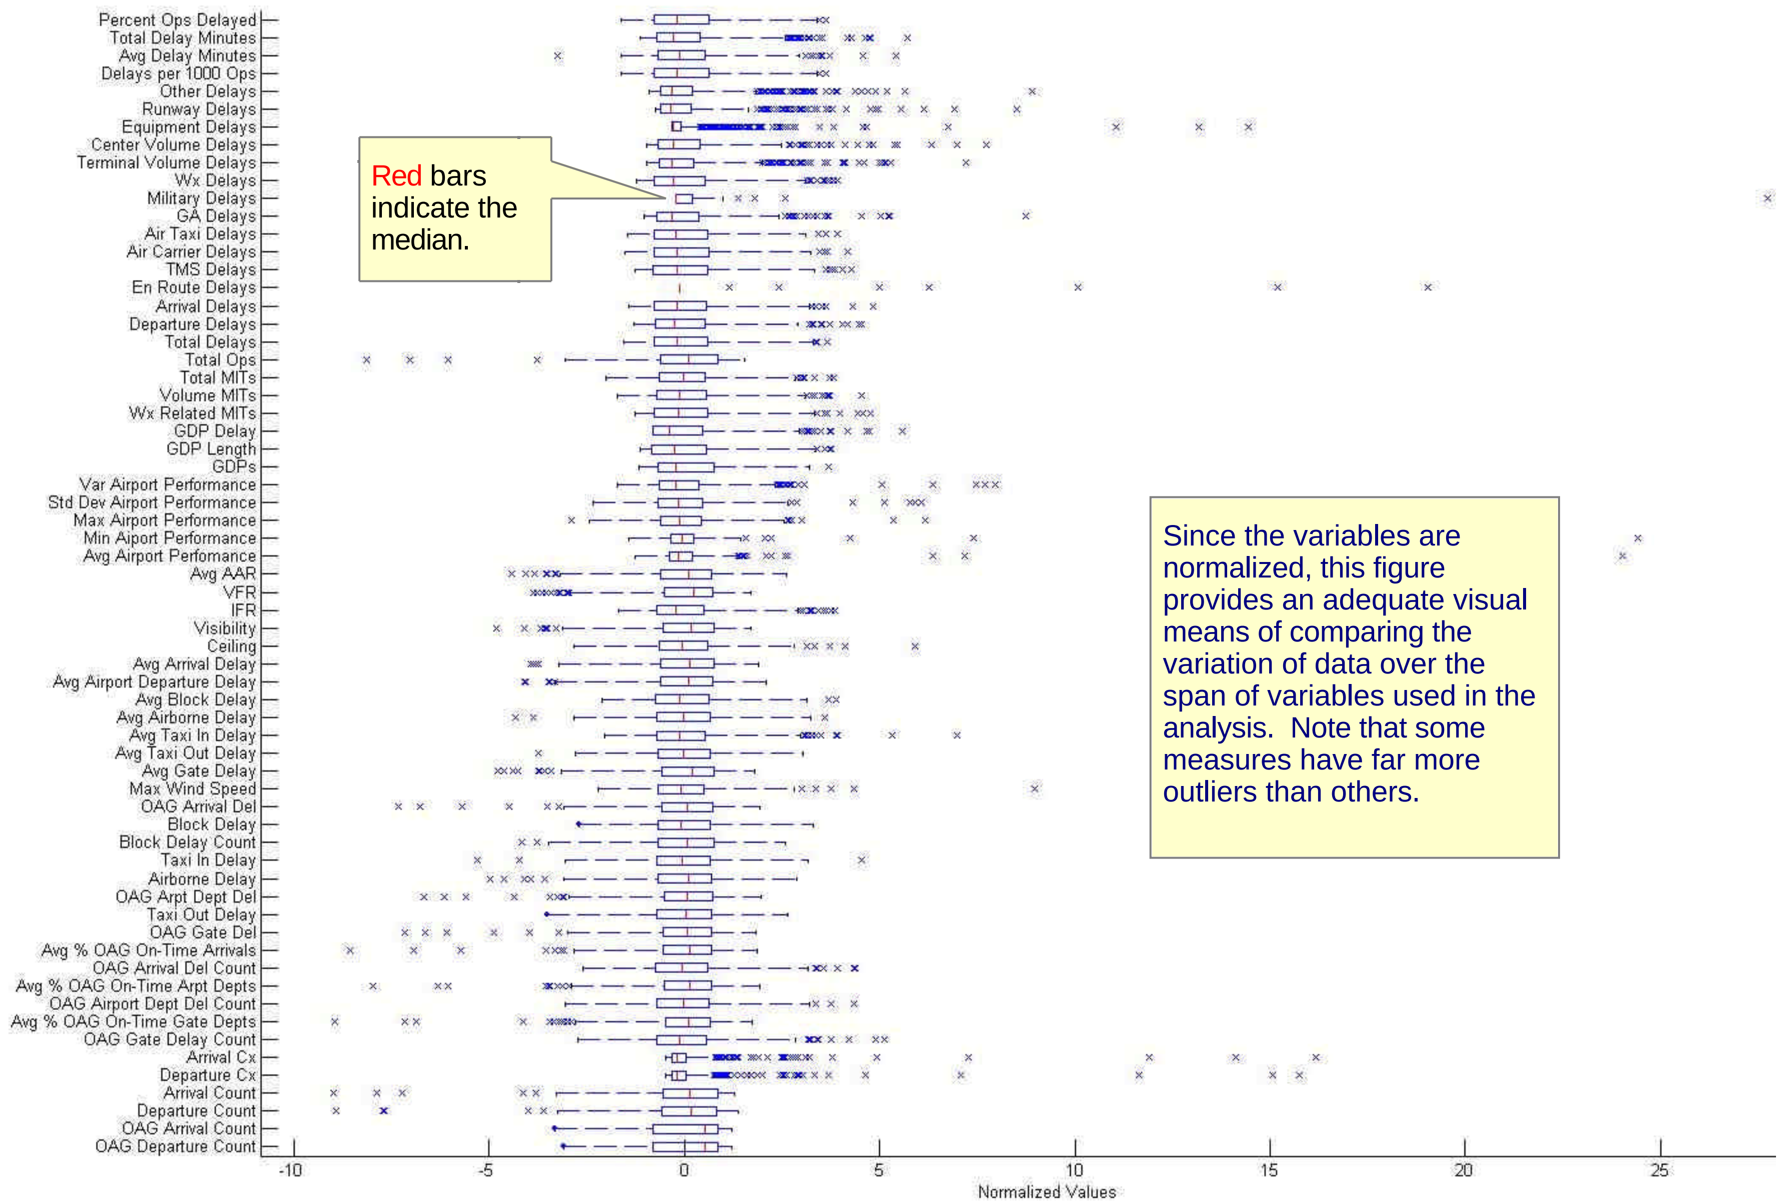

Figure 140. Analysis of states, control actions, and performance measures of the NAS.

In both the optimal feature vector analysis and the types-of-day analysis, a two-phase cluster analysis was pursued.

The process for the identifying the optimal NAS feature vector is described next. The candidate 65 variables were clustered (bundled) by similar statistical behavior to reduce the variables (**Table 9**) to a more intuitive, manageable set. This set of key variables defines what we call the "optimal" NAS feature vector. The clustering procedure used an interactive combination of a clustering algorithm and subjective review for integrity of interim results. **Figure 141** depicts the process.

The first step in **Figure 141** is an oblique PCA clustering algorithm (from SAS), which forms an initial clustering. The algorithm was run with an initial setting of the maximum number of clusters. Each cluster was reviewed for uniformity and membership count. In general, a clustering algorithm will return the maximum number of clusters, because its internal objective function can always be reduced by creating one or more singleton clusters. So, the algorithm was run again with the maximum number of clusters decremented by 1. This cycle, Loop 1 in **Figure 141**, was repeated until cluster memberships were deemed reasonable. (A good rule of thumb is that each cluster should have at least 2% membership.)

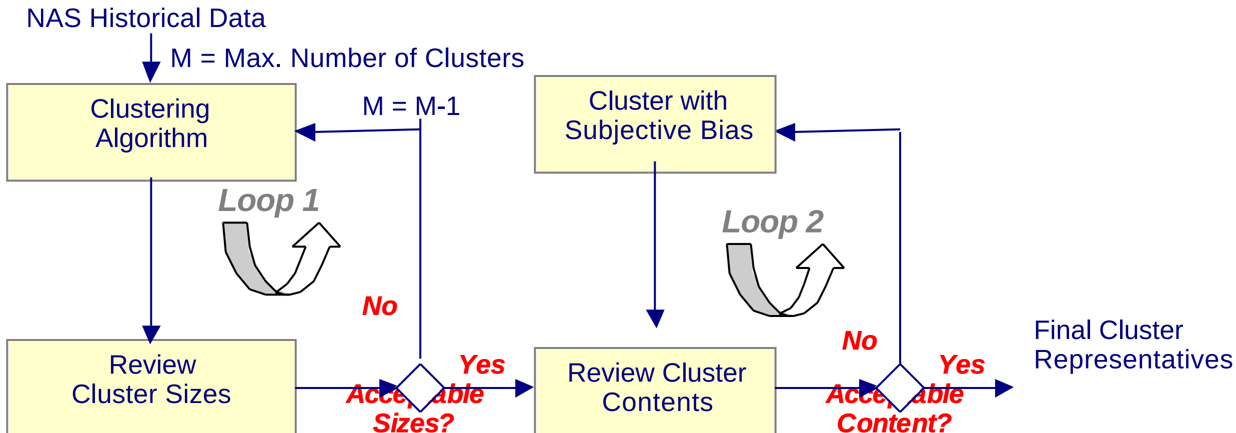

Figure 141. Cluster Analysis flow chart (note: Loop 2 was planned for in our analysis but it was not required, that is, the content was determined to be acceptable).

Once the cluster membership counts were acceptable, the clusters were reviewed for content. The algorithm clustered variables based on similar statistical behavior (variance). Intuitive knowledge of NAS behavior and variable nuances was applied to ensure that the algorithmic groupings are consistent with known, or suspected, relationships. For instance, one would expect to see daily departure counts and arrival counts to be in the same cluster.

If content of one or more clusters is unacceptable, then the process would enter Loop 2. There are two options: to override the clustering algorithm with subjective bias – that is, force certain groupings – or to appeal to another type of clustering algorithm. Fortunately, our process never had to enter Loop 2.

5.2.1 Phase I: Variable Clustering to Determine an Optimal NAS Feature Vector

We next step through the cluster analysis process that determines the optimal NAS feature vector. While results are summarized in this section, detailed notes taken during the clustering process are given in **Appendix D**.

In order to determine the optimal NAS feature vector, Loop 1 was executed six times. With each iteration, we decided whether certain variables should be eliminated.. The criteria for potential elimination of a variable *v* were:

- *v* is redundant, i.e., there exists another variable *w* with an unusually strong correlation with *v* (hence there is no need for both *v* and *w*);
- *v* is clearly "homeless", meaning that it has an extremely weak association with all of the clusters;
- *v* is essentially constant over time.

We begin with easily identified, simple correlations and eliminate variables that illustrate clear redundancies. Variables *v32* and *v33* (average conditions of IFR and VFR) have a perfect correlation (correlation coefficient = 1). By definition, these variables sum to 1 on each day. Only one is needed, so we retain only the IFR variable. Next, **Figure 142** shows that the variables "Volume-related MIT" and "Total MIT Restrictions" are highly correlated. One of these could be eliminated, so we arbitrarily chose the former.

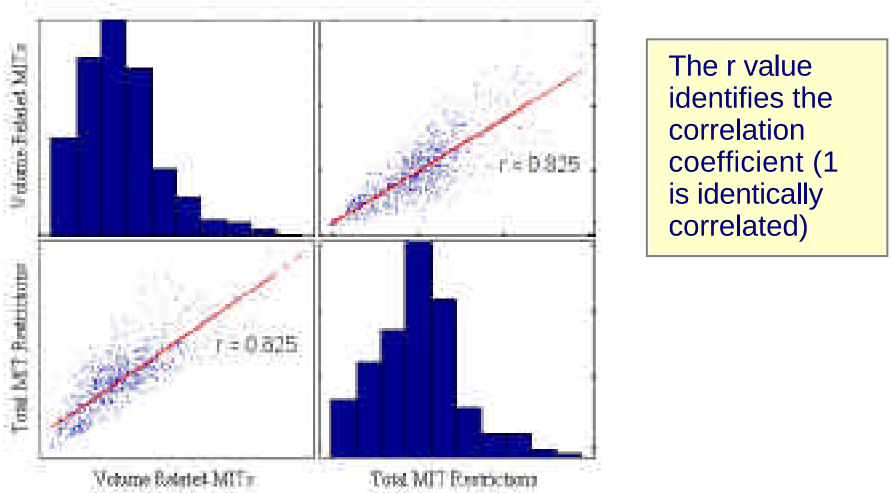

Figure 142. A scatter plot[8] shows that Volume-related MIT and Total MIT restrictions are highly correlated.

Unfortunately, none of the weather variables by itself proved to be a prominent variable in any of the variable clusters. Two of the weather variables, "Average Daily Visibility" and "Maximum Reported Wind Speed", fell into a cluster with the GDP variables. Also within this cluster were average hourly airport acceptance rates, and daily count of weather-related MIT restrictions. This assembly makes sense, since the formal issuance of delay by the FAA is usually a reaction to adverse weather conditions. But the variable "Cloud Ceilings" was completely eliminated due to weak association with any of the bundles. Its lack of correlation with Maximum Reported Wind Speed and Average Daily Visibility, for instance, is demonstrated in, **Figure 143**.

[8] Note: Scatter plots use **blue** and **red** data points to indicate when the data is statistically significant. The **blue entries** represent two variables that had a p-value in the correlation matrix less than 0.05, while the **red entries** have p-value greater or equal to 0.05. The p-value is the probability of observing this correlation if the actual correlation was really 0. Essentially, that is saying that a p-value below 0.05 means that there is a very low probability that the observed correlation occurred by chance.

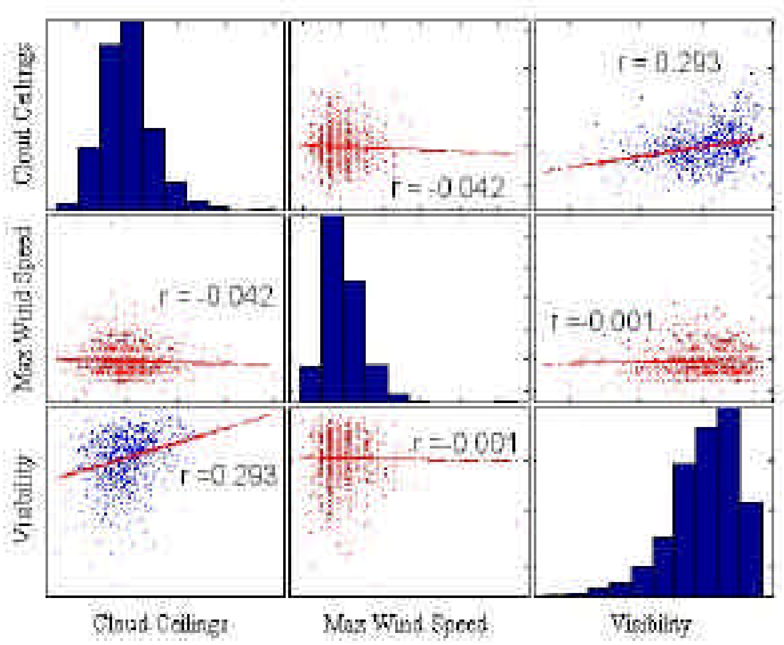

Figure 143. Low r values indicate little or no correlation with Maximum Reported Wind Speed and Average Daily Visibility.

Another cluster consisted of cancellation statistics and two of the airport performance scores. Direct examination of these variables revealed a strong association. But the latter two were eliminated because they are virtually constant over time, as seen in **Figure 144**, and therefore add no descriptive value to a description of the NAS. This left Cluster 6 with purely cancellation statistics.

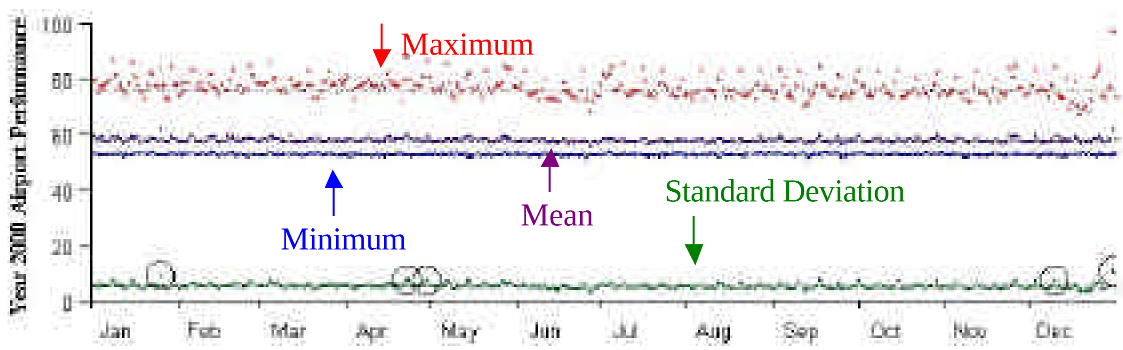

Figure 144. The mean and minimum airport performance scores are almost constant and thus do not reveal any meaningful information for the cluster analysis.

We next eliminated the IFR variable from the study on the grounds that is does not really belong in any of the bundles (weak association) and that it would add no value to NAS description as its own bundle. It also lacked correlation with the cloud ceilings variable, which was also eliminated.

Cluster 7 was mostly center and volume delays. Equipment and runway delays (*v62* and *v63*) showed a very weak association, but were eliminated since a proper cluster membership could not be found for them.

In all, we eliminated 8 of the 65 variables. **Table 10** documents the final cluster analysis results.

Table 10. All Cluster Memberships of the variables considered in cluster analysis.

No.	Variable	Cluster	Cluster	Units	$1\text{-}R^2$	Time Format

			Ranking		Ratio	
1	DATE	n/a	n/a	mm-dd-yyyy	n/a	n/a
2	OAG_DEP_CT	4	3	Count	0.34	Local Time (by Airport)
3	OAG_ARR_CT	4	4	Count	0.34	Local Time (by Airport)
4	DEP_CNT	4	2	Count	0.27	Local Time (by Airport)
5	ARR_CNT	4	1	Count	0.25	Local Time (by Airport)
6	DEP_CANCEL	6	2	Count	0.2	Local Time (by Airport)
7	ARR_CANCEL	6	1	Count	0.19	Local Time (by Airport)
8	OGATE_DELC	1	1	Count	0.07	Local Time (by Airport)
9	GDEP_ONTIM	3	2	%	0.05	Local Time (by Airport)
10	OARPT_DEPC	1	5	Count	0.17	Local Time (by Airport)
11	ADEP_ONTIM	3	4	%	0.12	Local Time (by Airport)
12	ODELARR_C	1	4	Count	0.16	Local Time (by Airport)
13	GARR_ONTIM	3	6	%	0.14	Local Time (by Airport)
14	OGATE_DEL	3	5	Minutes	0.12	Local Time (by Airport)
15	DELAY_TO	4	5	Minutes	0.52	Local Time (by Airport)
16	O_ARPT_DEP	3	1	Minutes	0.03	Local Time (by Airport)
17	DEL_AIR	4	6	Minutes	0.55	Local Time (by Airport)
18	DEL_TI	4	9	Minutes	0.93	Local Time (by Airport)
19	BLOCK_CNT	4	7	Count	0.57	Local Time (by Airport)
20	Block_DEL	8	2	Minutes	0.39	Local Time (by Airport)
21	O_DEL_ARR	3	3	Minutes	0.1	Local Time (by Airport)
22	WND_SPEED	8	11	Knots	0.97	Local Time (by Airport)
23	AVGGATEDEL	1	6	Minutes	0.18	Local Time (by Airport)
24	AVGDEL_TO	8	7	Minutes	0.71	Local Time (by Airport)
25	AVG_DEL_TI	3	7	Minutes	0.77	Local Time (by Airport)
26	AVGDEL_AIR	8	6	Minutes	0.68	Local Time (by Airport)
27	AVG_BLOCK	8	1	Minutes	0.38	Local Time (by Airport)
28	AVGARPTDEP	1	2	Minutes	0.07	Local Time (by Airport)
29	AVGODELARR	1	3	Minutes	0.13	Local Time (by Airport)
30	CEILING	eliminated	n/a	Hundreds of Feet	n/a	Local Time (by Airport)
31	VISIBLE	8	8	Miles	0.79	Local Time (by Airport)
32	IFR	eliminated	n/a	%	n/a	Local Time (by Airport)
33	VFR	eliminated	n/a	%	n/a	Local Time (by Airport)
34	AVG_ARR_RATE	8	9	Flights per Hour	0.82	Local Time (by Airport)
35	Avg_ARPT_PER	eliminated	n/a	Score (out of 100)	n/a	Local Time (by Airport)
36	Min_ARPT_PER	eliminated	n/a	Score (out of 100)	n/a	Local Time (by Airport)
37	Max_ARPT_PER	5	3	Score (out of 100)	0.08	Local Time (by Airport)
38	SD_ARPT_PER	5	1	Score (out of 100)	0.02	Local Time (by Airport)
39	Var_ARPT_PER	5	2	Score (out of 100)	0.03	Local Time (by Airport)
40	GDPs	8	3	Count	0.47	Zulu (FSM Day)
41	GDP_Length	8	4	Minutes	0.48	Zulu (FSM Day)
42	GDP_Delay	8	5	Minutes	0.52	Zulu (FSM Day)

43	WxRel_MIT	8	10	Count	0.89	Zulu (GMT)
44	Volume_MIT	eliminated	n/a	Count	n/a	Zulu (GMT)
45	Total_MIT	4	8	Count	0.72	Zulu (GMT)
46	Total_Ops	7	1	Count	0.42	UTC
47	Total_Delays	2	1	Count	0.03	UTC
48	Dep_Del	2	9	Count	0.44	UTC
49	Arr_Del	2	11	Count	0.59	UTC
50	Enroute_Del	2	14	Count	0.99	UTC
51	TMS_Del	2	7	Count	0.26	UTC
52	AirCarrier_Del	2	4	Count	0.05	UTC
53	AirTaxi_Del	2	8	Count	0.3	UTC
54	GA_Del	2	12	Count	0.62	UTC
55	Military_Del	2	13	Count	0.97	UTC
56	Wx_Del	2	5	Count	0.07	UTC
57	TermVol_Del	7	2	Count	0.53	UTC
58	CenterVol_Del	7	4	Count	0.68	UTC
59	Equip_Del	eliminated	n/a	Count	n/a	UTC
60	Runway_Del	eliminated	n/a	Count	n/a	UTC
61	Other_Del	7	3	Count	0.62	UTC
62	DelsPer1000Ops	2	2	Count	0.03	UTC
63	AvgDel_Min	2	10	Minutes	0.54	UTC
64	TotDel_Min	2	6	Minutes	0.08	UTC
65	PercOpsDelayed	2	3	%	0.03	UTC

Once the cluster content stabilized and met with the analyst's approval, a representative from each cluster was chosen. Since the algorithm outputs an index for each variable – which represents the strength of association with its cluster – the variable with the strongest association was chosen as the representative. Nevertheless, another variable may be chosen instead for subjective reasons. In one instance we chose "GDP minutes" (total ground delay issued in GDPs) as the representative, even though it was the fifth strongest variable. It was chosen over both "GDP count" (number of GDPs implemented) and "GDP length" (hours of GDP duration summed over all GDPs), because we knew that "GDP minutes" implicitly contains the other two, and is a less discrete indicator of demand-capacity imbalances. (A more discrete variable generally produces a multi-modal distribution, which would tend to cluster the vectors artificially in the next round of analysis.) In all, we overruled the default selection in just this one case.

The clustering process culminated in eight variable bundles and their representatives, which are listed in **Table 11**. Further details on each of the clusters are presented in **Table 12** through **Table 19.**

Table 11. Optimal Variable Cluster Set.

Cluster	Cluster Name	Prominent Variable within Cluster	Members in Cluster
1	Gate Delays	Daily Count of OAG-Based Gate Delays	6
2	Overall Delays	Total Delay Count From OPSNET	14
3	On-time Performance	Daily Total OAG-Based Airport Departure Delay (minutes)	7
4	Traffic Volume	Daily Arrival Count	9

5	Airport Performance Metric	Std Dev of Airport Performance Score (21 ASPM Airports)	3
6	Cancellations	Daily Arrival Cancellations Count	2
7	Volume-related Delays	Total Operation Count From OPSNET	4
8	Weather and GDPs	Total Delay attributed to GDPs (minutes)	11

Table 12. The 6 Variables that form the "Gate Delays" Cluster.

Variable	Variable Description	Ranking
v8	Daily Count of OAG-Based Gate Delays	1
v28	Average Daily Airport Departure Delay (minutes)	2
v29	Average Daily OAG-Based Airport Arrival Delay (Gate In) (minutes)	3
v12	Daily Count of OAG-Based Arrival Delays	4
v10	Daily Count of OAG-Based Airport Departure Delays	5
v23	Average Daily Gate Delay (minutes)	6

Table 13. The 14 Variables that form the "Overall Delays" Cluster.

Variable	Variable Description	Ranking
v47	Total Delay Count From OPSNET	1
v62	Delays per Thousand Operations from OPSNET	2
v65	Percentage of Operations Delayed from OPSNET	3
v52	Total Air Carrier Delay Count From OPSNET	4
v56	Total Weather Related Delay Count From OPSNET	5
v64	Total Minutes of Delay from OPSNET	6
v51	Total Traffic Management Initiative Delay Count From OPSNET	7
v53	Total Air Taxi Delay Count From OPSNET	8
v48	Total Departure Delay Count From OPSNET	9
v63	Avg minutes of Delay from OPSNET	10
v49	Total Arrival Delay Count From OPSNET	11
v54	Total General Aviation Delay Count From OPSNET	12
v55	Total Military Delay Count From OPSNET	13
v50	Total En route Delay Count From OPSNET	14

Table 14. The 7 Variables that form the "On-Time Performance" Cluster.

Variable	Variable Description	Ranking
v16	Daily Total OAG-Based Airport Departure Delay (minutes)	1
v9	Average % OAG-Based On Time Gate Departures (per 1/4 hour)	2
v21	Daily Total OAG-Based Arrival Delay (minutes)	3
v11	Average % OAG-Based On time Airport Departures (per _ hour)	4
v14	Daily Total OAG-Based Gate Delay (minutes)	5
v13	Average % OAG-Based On Time Arrivals (per _ hour)	6
v25	Average Daily Taxi In Delay (minutes)	7

Table 15. The 9 Variables that form the "Traffic Volume" Cluster.

Variable	Variable Description	Ranking

v5	Daily Arrival Count	1
v4	Daily Departure Count	2
v2	Daily OAG Scheduled Departure Count	3
v3	Daily OAG Scheduled Arrival Count	4
v15	Daily Total Taxi Out Delay (minutes)	5
v17	Daily Total Airborne Delay (minutes)	6
v19	Daily Count Of Flights With Block Delay	7
v45	Daily Count of Total MIT Restrictions	8
v18	Daily Total Taxi In Delay (minutes)	9

Table 16. The 3 Variables that form the "Airport Performance Metric" Cluster.

Variable	Variable Description	Ranking
v38	Std. Dev. over 21 ASPM Airports of Daily Airport Performance Score	1
v39	Variance over 21 ASPM Airports of Daily Airport Performance Score	2
v37	Maximum over 21 ASPM Airports of Daily Airport Performance Score	3

Table 17. The 2 Variables that form the "Cancellations" Cluster.

Variable	Variable Description	Ranking
v7	Daily Arrival Cancellations Count	1
v6	Daily Departure Cancellations Count	2

Table 18. The 4 Variables that form the "Volume Related Delays" Cluster.

Variable	Variable Description	Ranking
v46	Total Operation Count From OPSNET	1
v57	Total Terminal Volume Related Delay Count From OPSNET	2
v61	Total Other Delay Count From OPSNET	3
v58	Total Center Volume Related Delay Count From OPSNET	4

Table 19. The 11 Variables that form the "Weather and GDP" Cluster.

Variable	Variable Description	Ranking
27	Average Daily Block Delay (minutes)	1
20	Daily Total Block Delay (minutes)	2
40	Daily GDP Count	3
41	Length of GDP (minutes)	4
42	Total Delay attributed to the GDP (minutes)	5
26	Average Daily Airborne Delay (minutes)	6
24	Average Daily Taxi Out Delay (minutes)	7
31	Average Daily Visibility	8
34	Average Hourly AAR over all airports throughout the day	9
43	Daily Count of Weather-Related MIT Restrictions	10
22	Maximum Reported Wind Speed (Knots)	11

Each cluster was given a name to convey the major theme of the comprising variables. For each day, the eight corresponding aggregate statistics were compiled, with the intent of performing a

Phase II cluster analysis to determine the different “types” of days in the NAS. For instance, the feature vector for February 11, 2001 is shown in **Figure 145**.

Gate Delays	Overall Delays	On-Time Performance	Traffic Volume	Airport Performance Metric	Cancellations	Volume-Related Delays	GDPs
3490 flts.	190 flts.	14,500 min.	20,081 flts.	5.474	471 flts.	47,600 flts.	7,480 min.

Figure 145. Feature vector for February 11, 2001.

5.2.2 Phase II: Clustering to determine the Types of Days in the NAS

In Phase II, a cluster analysis was performed on the NAS feature vectors passed from Phase I. The objective in Phase II was to classify the NAS feature vectors for each day Jan. 2000 through Sept. 10, 2001 into groups that naturally described different types of days in the NAS. For instance, if one of the variables were overwhelmingly bimodal, then the algorithm would tend to break the vectors into two groups corresponding to the two (implicit) distributions given by that variable. Without this crucial step, the multi-modal nature of the NAS feature vector components might render the type-of-day classification meaningless.

In theory, a clustering algorithm could break the feature vectors (data points) into any number of clusters. Each cluster would represent a different type of day in the NAS, and within each cluster, we could define typical and atypical days. But, on an intuitive level, it seemed that the number of types of days in the NAS should be relatively small. For instance, a natural decomposition might be six clusters, resulting from three levels of traffic volume, each with two possible levels of weather conditions.

This time, we used a centroid-based (K-means) clustering algorithm. The overall iterative process was the same as the variable bundling process (**Figure 141**), with two exceptions: (1) the data points are days in the calendar year rather than NAS feature vector variables, and (2) no data points were eliminated. **Appendix E** provides an annotated version of the intermediate results for the type of day clustering process.

First, we performed a relaxed cluster analysis, meaning that we did not interject any subjective biases into the algorithm with a generous upper bound on the maximum number of clusters (20). This resulted in 20 clusters (as we would expect) but only 10 of these had significant membership – many of the other clusters had only 1 or 2 data points in them. So, we ran the algorithm again with the maximum number of clusters set to 10. This drove the singletons back into the major clusters. This time, only 7 of the 10 resulting clusters had significant membership, so we ran the algorithm one more time with the maximum cluster value set at 7. The resulting 7 clusters had memberships of 62, 183, 104, 68, 16, 182, and 4. Each of these is considered significant (at least 2% of the number of data points). The low membership in cluster 7 was a bit unsettling, but examination revealed that these days were statistical outliers, which are often grouped together in a cluster analysis. This means that there were really six major clusters.

Satisfied with the resulting cluster membership counts, we proceeded to investigate which, if any, of the variables had been the primary determinant in dividing the data. (If there were no recognizable pattern, then it would be hard to characterize the clusters as to which types of days they represent.) We used the X-gobi software tool to visually examine the data and clusters from multiple dimensions. In particular, we plotted each of the eight variables against the cluster numbers. A typical scatter plot is shown in **Figure 146**. Each point (m,n) represents a day belonging to cluster n, which had m hours of GDPs run that day. The data points are of course gathered along their respective cluster lines, but there is no recognizable relationship between cluster number and GDP hours.

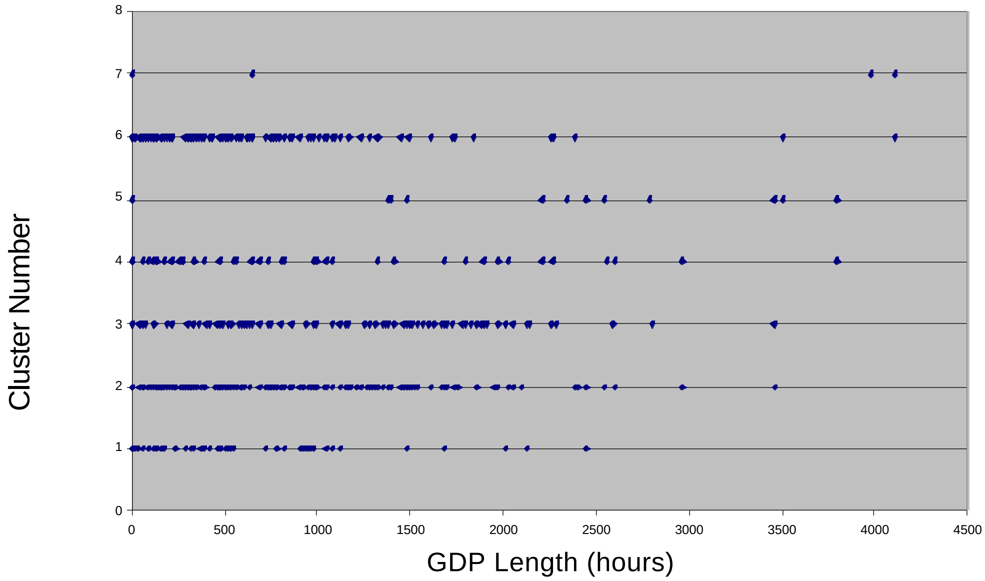

Figure 146. Plot of GDP length vs. Cluster.

In contrast, we found one variable that had an almost perfect relationship with cluster membership: "GDP minutes", which is the number of minutes of ground delay assigned by the FAA during a ground delay program. Each cluster was almost completely characterized by the number of GDP minutes spanned by its members, as seen in **Figure 147**.

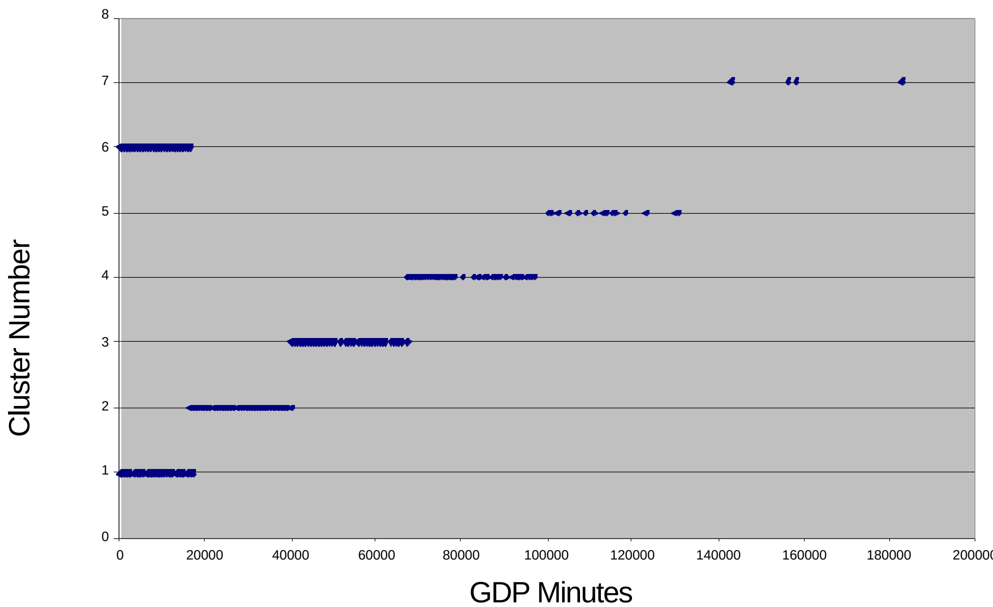

Figure 147. Plot of GDP minutes vs. Cluster.

Aside from a very slight overlap between ranges, the GDP minutes range increases as the cluster number increases. (The notable exception to this association is Cluster 6, which we will discuss shortly.) That is, in terms of the types of days in the NAS, we can now make the following distinction (ordered by GDP minutes):

- All days with NAS Day Type 1 have: 0 < GDP minutes < 17,302
- All days with NAS Day Type 6 have: 0 < GDP minutes < 16,236

- All days with NAS Day Type 2 have: 16,314 < GDP minutes < 40,142
- All days with NAS Day Type 3 have: 40,257 < GDP minutes < 67,269
- All days with NAS Day Type 4 have: 67,139 < GDP minutes < 96,931
- All days with NAS Day Type 5 have:100,020 < GDP minutes < 130,460
- All days with NAS Day Type 7 have:142,770 < GDP minutes < 182,677

Cluster 1 and Cluster 6 clearly share the same GDP minutes range. Cluster 1 overlaps with cluster 2 by just 988 minutes; Cluster 3 overlaps with cluster by just 130 minutes; the remaining are non-overlapping.

Statistically, GDP minutes are the single most important variable to consider of the eight representatives when lumping days by similar characteristics. Intuitively, this means that there are six types of days in the NAS (seven, if one is willing to count the outliers in Cluster 7):

- Type of NAS Day 1: Very low GDP level, with low operations count (cluster 1)
- Type of NAS Day 6: Very low GDP level, with high operations count, (cluster 6)
- Type of NAS Day 2: Low GDP level (cluster 2)
- Type of NAS Day 3: Medium GDP level (cluster 3)
- Type of NAS Day 4: High GDP level (cluster 4)
- Type of NAS Day 5: Very High GDP level (cluster 5)

Variable *v*8, gate delays, exhibited a slight relationship with cluster number, but not nearly as strong as GDP minutes. The same was true for GDP count.

Next, we sought to distinguish the two types of low GDP level days. In particular, why did the cluster algorithm choose to break Cluster 1 into two clusters (1 and 6)? Variable *v46*, which is the number of total operations in the NAS for the day, exhibits a bimodal behavior, as illustrated in **Figure 148**. The algorithm found a more efficient grouping by breaking this cluster into two groups. (A sub-optimal solution would have been to establish 7 levels of GDP minutes, each with a unique range.) We had already seen the bimodal behavior in **Section 4.1.2**, where we directly examined arrival and departure counts. Weekdays (Monday through Friday) tend to have more traffic than weekends (Saturday or Sunday).

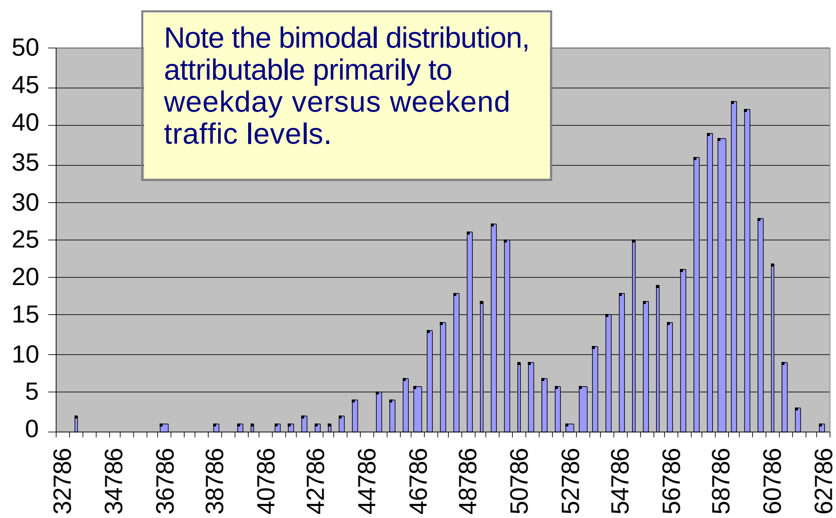

Figure 148. Histogram of total operations count for Jan. 1, 2000 through Sept. 10, 2001.

Traffic volume is a secondary factor in characterizing what type of day it is in the NAS, next to GDP minutes. This is made clear by **Figure 149**. This is a multi-dimensional projection of the data points. The crosses indicate the days with low levels of GDP minutes (Clusters 1 and 6); the other data points are members of the other five clusters. Note that this group is clearly split into two groups by the bimodal nature of the total operations variable (*x*46).

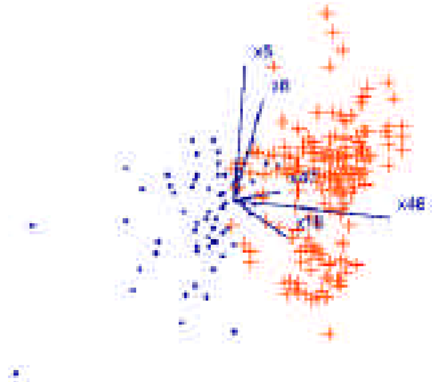

Figure 149. Cluster of low level of GDP minutes (+ symbols).

Having established the six type-of-day clusters, we were able to rank the days within each cluster according to how typical they were, using proximity to the center of the cluster as the criterion. That is, let be the vector created by setting equal to the mean of the k^{th} variable of the NAS feature vector, taken over the vectors in a fixed cluster. Mean vector is the center of the cluster mass. Then the vector closest to was considered to be the most typical day in the cluster. Proximity was defined using a Euclidean-based metric normalized for variance. That is, let and be two eight-dimensional vectors. The weighted distance between them is defined as

where is the variance of the k^{th} variable. Without this normalization, proximity would be skewed toward the variables with the larger values.

Table 20 presents the mean vector for each type-of-day cluster and **Table 21** presents the three closest (most typical) days for each type-of-day cluster. The center of the Type 1 type-day cluster is given by row 1 in **Table 20**. **Appendix H** presents the final rankings for the different types of days in the NAS.

Table 20. Data for the Mean Vector for each Cluster.

Cluster	Count	Variable							
		v8	*v47*	*v16*	*v5*	*v38*	*v7*	*v46*	*v42*
1	62	3737	655	14161	19932	5.722	526.9	46689	6568
2	183	4801	1136	14150	21524	5.455	539.2	54546	28119
3	104	5508	1075	13762	21985	5.229	675.5	54105	51053
4	68	6774	1136	12307	21822	5.145	918.1	54744	79974
5	16	6958	1280	12254	22092	5.183	961.4	56728	112424
6	182	4053	1222	14973	21512	5.801	434.1	56850	5355
7	4	8040	1264	10230	20978	5.339	1008.2	53374	159973

Table 21. Most Typical Day in each Cluster.

Type of Day	Dates			Distances		
Cluster	1st	2nd	3rd	1st	2nd	3rd
1	02-11-01	03-12-00	01-07-01	1443	1799	2643
2	08-03-01	09-05-00	02-17-00	1572	1596	2222
3	10-04-00	06-20-00	06-13-01	2319	2475	3741
4	01-15-01	02-11-00	03-15-01	2930	2935	4103
5	05-22-01	06-16-00	10-27-00	3047	4342	4855
6	07-12-01	05-02-01	03-31-00	1949	2259	2489
7	02-25-01	07-28-00	11-26-00	7035	7780	18252

Figure 150 shows a 2-dimensional representation of the metric being applied within a cluster (see Cluster 2). **Figure 150** also shows how the metric can be used to measure the distance between a cluster center (c_3) and the center of all clusters (U). Within each cluster, the days can then be ranked according to proximity with the center of the cluster. The day whose vector is closest to the center of the cluster is considered the most "typical" day in that cluster.

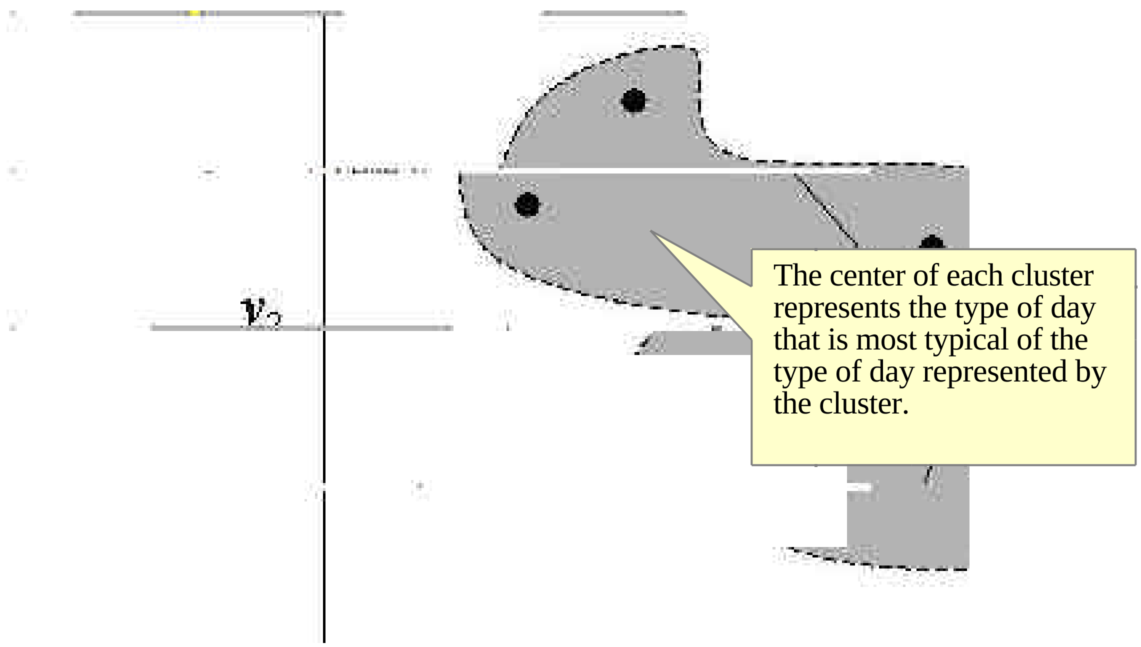

Figure 150. The Euclidean metric is used to measure the distance of a day-vector (point) from the center of its cluster (e.g., Cluster 2) or to measure the distance of a cluster from the center of all clusters (***U***).

5.3 Interpretation of Results

The results of the clustering algorithm showed that variables *v42* and *v46* were the most critical in separating the day vectors into clusters. We investigate why this is and the range of conclusions we are entitled to draw.

In scatter plot format, **Figure 151** shows the clustering of data points by variables *v42* and *v46*. This is a projection of the 8-dimensional day cluster data points onto the *v42*-*v46* plane. Each point represents one day; the vertical coordinate is the number of operations for that day, while the horizontal coordinate is the number of number of GDP minutes. The points are marked and colored by cluster number.

In **Figure 151**, note that the points on the far left (with low or zero GDP minutes) are divided into an upper group and a lower group, clusters 6 and 1, respectively. This is the effect of the bimodal distribution of *v42*, which we have already seen. The separation is designed to alleviate the debate of which of the two modes (low operations or high operations) is more typical, by breaking them into two clusters, thus making it a moot point.

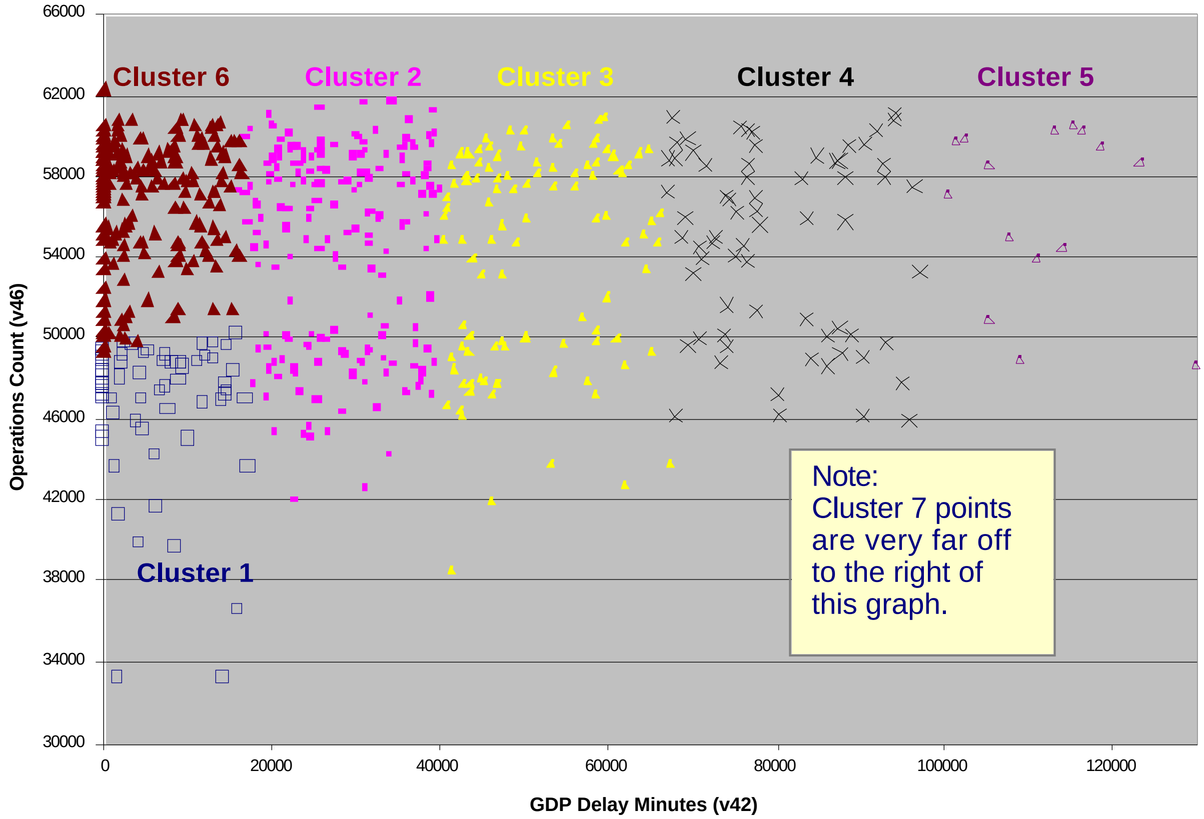

Figure 151. Scatter plot of GDP Delay Minutes vs Operations Count.

The objective of the clustering algorithm is to create a fixed number of clusters such that the variance within each cluster is minimized, while the variance between clusters is maximized. Intuitively, this is the same as identifying concentrations of data in multi-dimensional space. Sweeping left to right in **Figure 151**, the breakdown by GDP minutes forms vertical lines of separation. These separations have been made to reduce the variance in the GDP minutes (*v42*) distribution. In the scatter plot, the horizontal separation of the data seems somewhat arbitrary. This is because the frequency of the points with respect to the horizontal axis is obscured. Consider the frequency distribution of the GDP minutes variable, which is shown in **Figure 152**. The distribution is concave, and heavily skewed to the left. This is the same as saying that in the scatter plot, the number of data points drops off as we move from left to right. The variance of this type of distribution is much greater than that of a classic bell-shaped distribution of equal mass. (In fact, the only way one could rearrange the same mass to have more variance is to evenly divide the mass between the two extreme points.) The intuitive justification for addressing variables with this type of distribution is that they make it the most difficult to answer the question of what is typical. That is, the mean is very far from the mode. (We noted that the other variables in the *v42* bundle tended to have similar shapes.)

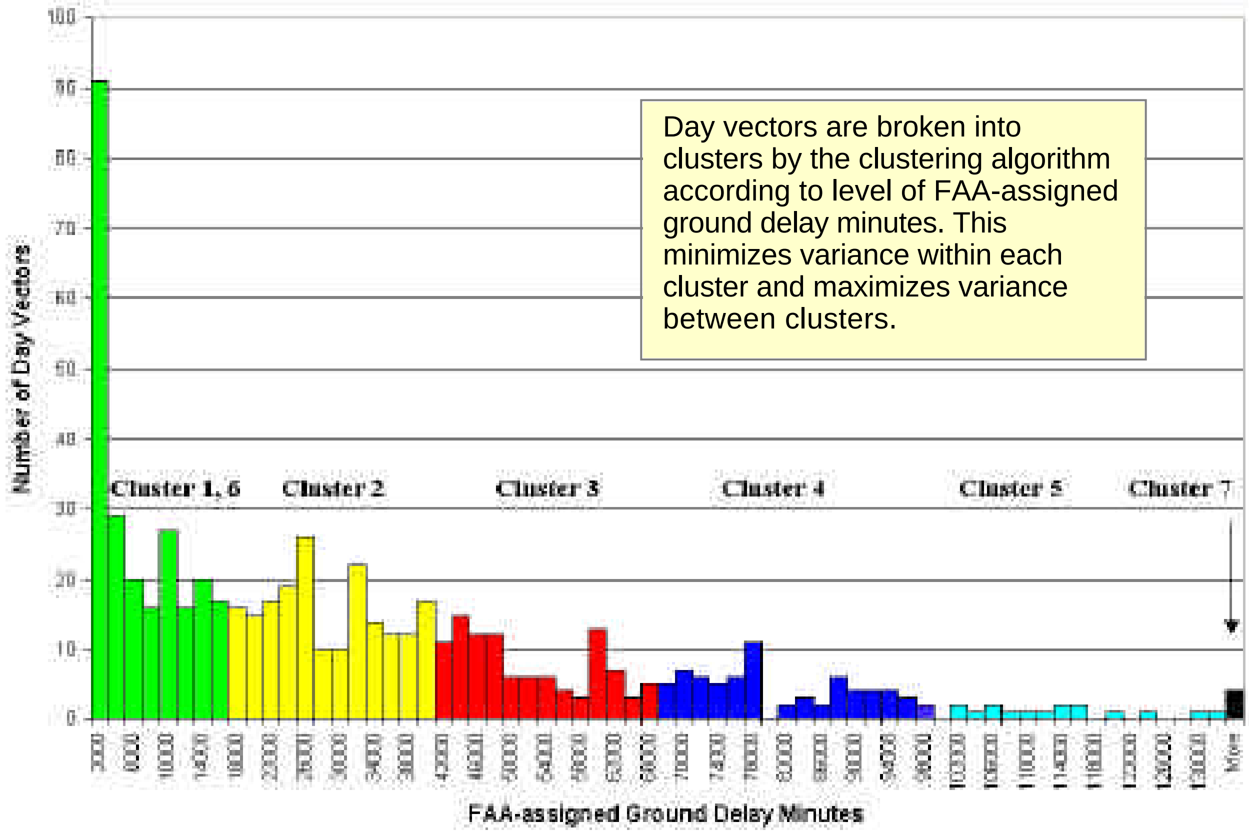

Figure 152. Histogram illustrating the relative cluster locations.

A natural question is why the algorithm didn't continue making vertical separations of clusters 2, 3, 4, and 5? If one were to cluster the points solely on the basis of this scatter plot, a strict separation of the points into an upper group and a lower group seems the most natural. First, as we have pointed out, the greatest payoff is breaking up the horizontal variance, which is not immediately obvious in the scatter plot. A closer examination reveals that the frequency of points drops off as we move from left to right. The magnitude of this drop is not fully appreciated because points on the left are sitting on top of other ones. If one were to draw the points up out of the *v42*-*v46* plane by plotting a third dimension (one of the other variables), then the high variance in *v42* would become more clear. We used the X-gobi software to do this, which requires continuous rotation of data to be visually effective. Second, note that the vertical separation of points in the scatter plot is less pronounced on the right (in Clusters 4 and 5) than it is on the left (Clusters 1 and 6). Overall, the data forms a horseshoe, and the algorithm used an optimization routine to decide where it is no longer profitable to form vertical separations.

The separation of day vectors by variables *v42* and *v46* does not mean that either of these is a more important indication of the state of the NAS than the other six representative variables. It just means that these are the most problematic when trying to determine what a typical day is like. The issue of which of two days respect is more typical breaks down between clusters, but is preserved within clusters.

The end result of this separation process is a set of clusters that are distributed around their geometric centers, in as many dimensions as possible. Confirmation is provided in the distributions shown in **Figure 153** and **Figure 154**.

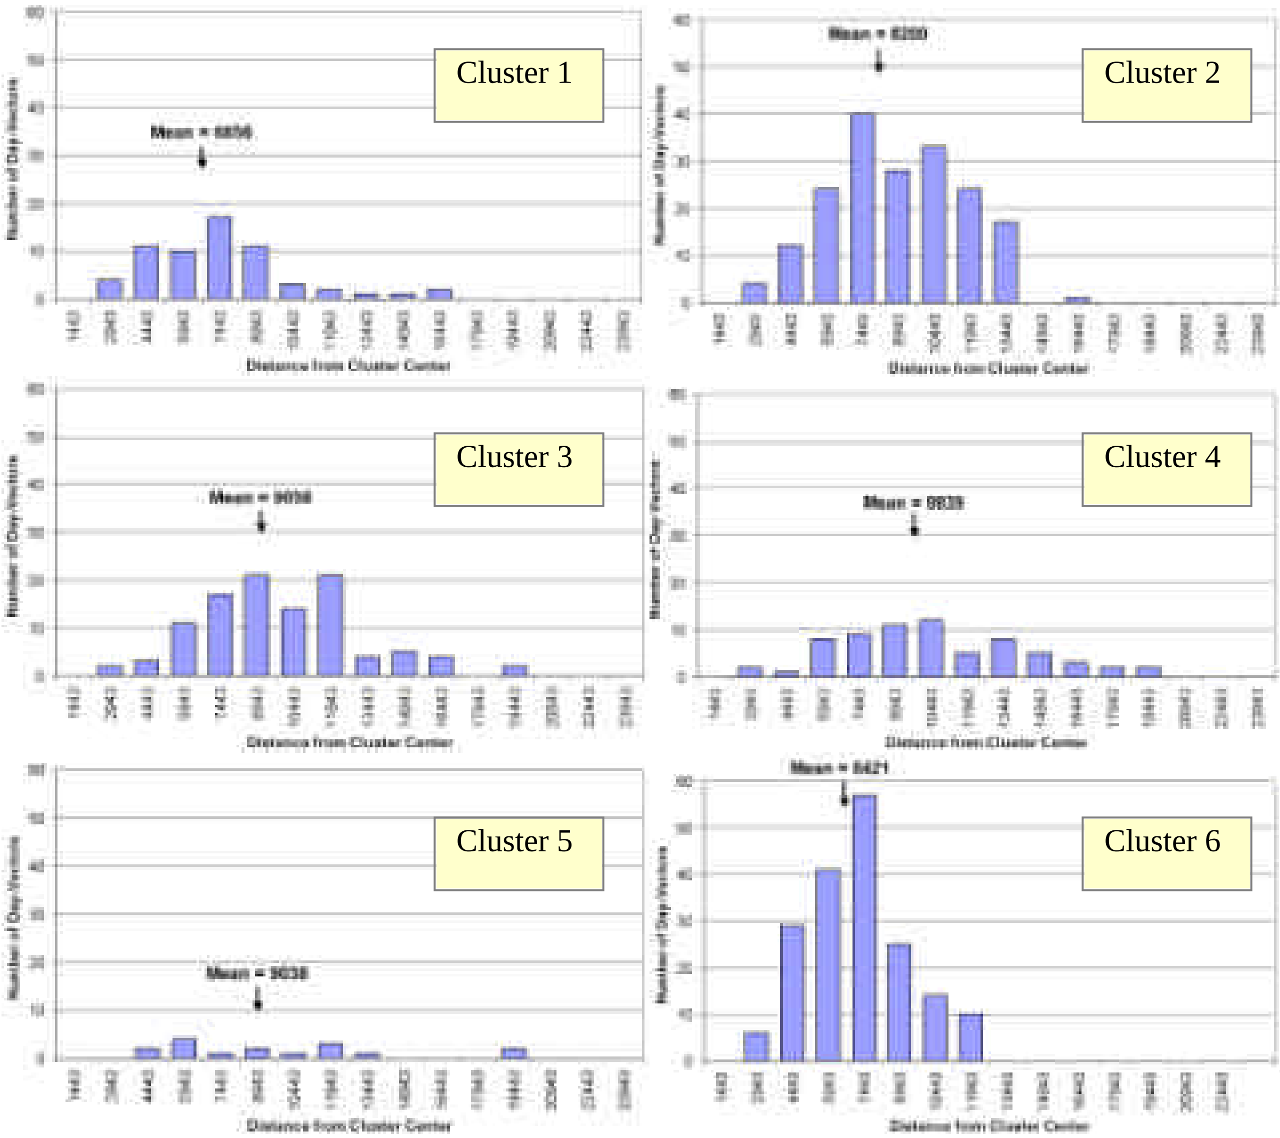

Figure 153. Cluster Variances for 6 Types of Days.

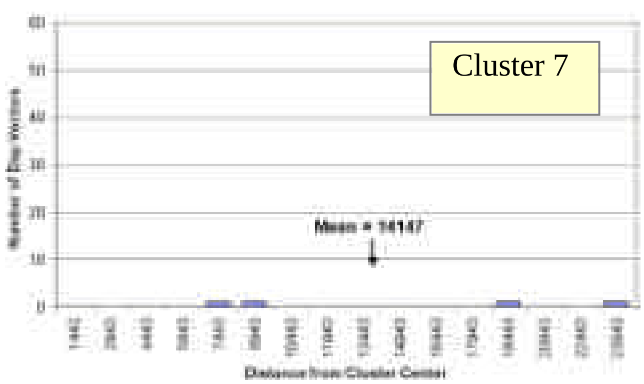

Figure 154. Cluster Variance for the outlier days cluster.

Each of these shows the frequency distribution of cluster member distance from center of cluster. The distributions are bell-shaped, which is directly in line with our intuitive notion of *clustering*.

(The fact that some seem unduly flat is a manifestation of the vertical axis having been stretched to a common size for all clusters.)

We emphasize that the results of the cluster analysis caution us against making comparisons across clusters as to whether one day is more typical that another. Strictly speaking, the point is moot. However, if really pressed to answer this question, one could consider using the cluster memberships as a guide: the cluster with the highest membership is, in some sense, the most typical, and within that cluster, we already know which day is the most typical. This approach is not recommended, since it is fraught with difficulties, however, because it both ignores the results of the cluster analysis and, at the same time, makes use of it.

The lessons learned from the cluster analysis dictate that before a typical day can be selected, certain questions about the nature of the desired day must be addressed. Without this distinction, the concept of typical loses its statistical meaning. These necessary distinctions between different types of days are summarized by the decision tree in **Figure 155**.

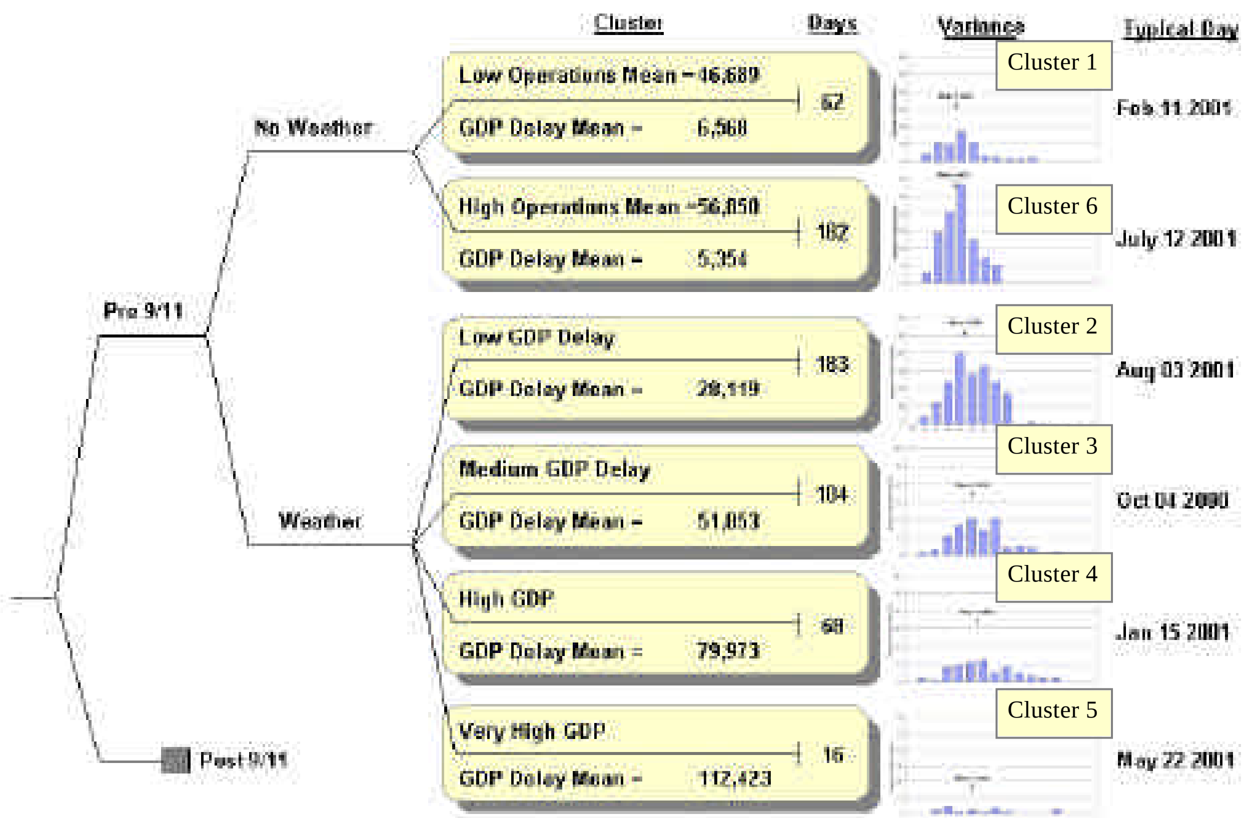

Figure 155. Decision Tree representing the Types of Days in the NAS.

The use of the decision tree is as follows. Reading left to right; a decision must be made at node 1 whether a data set from pre-9/11 or post-9/11 is desired. Since our cluster analysis did not incorporate the post-9/11 data, the user is forced into the upper branch, pre-9/11. (The lower branch has been included for sake of completeness.) At decision node 2, the user decides whether to choose a data set from the collection of no weather days with very low GDP delay minutes (upper branch) or from the collection of weather days with GDP delay minutes ranging low to high (lower branch). The GDP variable is a surrogate for a collection of ground delay and weather related effects. So these branches have been informally dubbed “no weather” days and “weather” days, which are approximate, descriptive terms.

Unfortunately, we were not able to collect any "weather" variables that had much meaning in aviation. For example, variables such as IFR, VFR, and cloud ceilings, do not really help determine what type of day we are having in the NAS. The term "weather" is really a layman's term. What we are really concerned about are meteorological conditions that have an adverse effect on aviation. So, we relied on more indirect indicators of weather, such as ground delay programs and reduced capacity. In some sense, these are the best weather indicators because they are triggered only if there is weather that adversely affects aviation.

Moreover, the main purpose of this main branch is to avoid making comparisons (with respect to which is more typical) between days with essentially no GDPs and days with some (or many) GDPs. We do not claim that there is literally no weather on the days in the upper branch (though it is a reasonable guide and does not hurt to think of it that way).

Suppose that the user has selected the “very low GDP” (no weather) branch. At the next branch, we see that there are two types of “very low GDP” days: those with relatively low operations counts (arrivals and departures), and those with relatively high operations counts. A quantifiable definition of these terms is provided in each cluster box by the mean number of operations for that cluster.

Once the user has decided between these two, there is a unique cluster which houses all days with similar statistical behavior. To the right of the cluster, the most typical day of the cluster has been specified. This would be the optimal data set to consider, meaning that it is most typical. If this day is undesirable for subjective reasons (or if some data elements cannot be collected), then the next most typical day can be selected. We have ranked the days within each cluster by proximity to the Euclidean center of the cluster; higher indexes indicate a more typical day (ranking not shown).

Next, we return to the weather impacted day branch. There are four different types of days to choose from corresponding to days with a low level of GDP delay minutes, a medium level GDP delay minutes, a high level of GDP delay minutes, and a very high level of GDP delay minutes. A quantifiable definition of these terms is provided in each cluster box by the mean number of GDP minutes. As described above, the most typical day in the cluster can be chosen as a representative of that cluster (type of day), or another one can be chosen using the cluster ranking.

The overall interpretation of the results is that there are six types of days in the NAS. Each of those has a most typical day (listed in the far right column) in **Figure 155**. Referring back to **Table 21**, one can see additional days that have very similar statistical behavior to the most typical days.

As for applying these results to the topic of validating NAS simulation, these results indicate that simulation validation sets should consider weather and GDP modeling as a basis for validation data sets. First, if neither weather nor GDP are modeled in the NAS simulation, then the results of our study indicate that there are two types of days that are useful for validating such simulations. They are described by Cluster 1 and Cluster 6. If weather and GDPs are include in the NAS simulation – indeed, if weather is included then GDPs must also included – then, depending on the degree of simulation validation that is desired, there are several choices to be made. One can validate a NAS simulation with modeled weather and GDPs through Cluster 2 through Cluster 5 (with special attention to the lower membership size of Cluster 5). Furthermore, trends may be simulated by comparing pairs of clusters, e.g., (2, 3) vs. (2, 4) vs. (2, 5), each having a difference in magnitude of weather and GDP significance in the validation data sets. For a complete validation of a NAS simulation, simulation developers should validate their NAS simulations with at least one validation run from each type of day in the NAS.

5.4 Linguistic Descriptions

A more intuitive linguistic description of a cluster can be constructed by examination of its center vector. Each component of the center vector ($\mu_1, \ldots, \mu_7$) is mapped into a "high", "medium", or "low" category, by considering its distance from the mean of the variable over the entire data set (over all clusters). That is, let σ_k be the standard deviation of the k^{th} variable, and let M_k be its mean over all data. The categories are set via:

Low:	$\mu_k < M_k - \sigma_k$
Medium:	$M_k - \sigma_k \leq \mu_k \leq M_k + \sigma_k$
High:	$M_k + \sigma_k < \mu_k$

This mapping allows us to give an intuitive description of a given cluster, such as "Low levels of scheduled departures, Medium levels of taxi-in delay", etc. For instance, a cluster can be defined as follows:

- [High, Medium, or Low] Scheduled departures
- [High, Medium, or Low] Minutes of taxi delay
- [High, Medium, or Low] Minutes of block delay
- [High, Medium, or Low] Visibility
- [High, Medium, or Low] Airport performance at the 21 ASPM airports
- [High, Medium, or Low] Weather-related MIT restrictions
- [High, Medium, or Low] Delayed aircraft

While these types of linguistic descriptions may help to understand the clusters and the different types of days in the NAS, we caution that the mapping from variable means to high-medium-low categories will probably not be unique. There may be more than one cluster with the same intuitive description. This means that one cannot work backward from the high-medium-low descriptions to create clusters. In particular, one cannot conclude that two days have "similar behavior" just because their high-med-low descriptions are the same. This could corrupt simulation model validation or demonstration efforts.

6 Special Days in the NAS

This chapter identifies special types of days in the NAS. Weather, special events, special times of the year, such as holidays, and other occurrences affect the NAS by causing increases or decreases in the volume of traffic. Some of these occurrences have virtually no effect at all. These events also affect different parts of the country in different ways.

6.1 Identification of Special Days in the NAS

The following events categorize particular special days in the NAS:

- Severe Weather Days
- Holiday travel (day before a holiday vs. holiday vs. day after a holiday)
- Special event days
- Rare events

Convective weather season is defined as the period from mid-April through September. It is characterized by pop-up thunderstorms. However, severe weather days can occur outside the convective weather season. Holiday travel occurs at different points in the year. Special events are usually yearly events other than holidays, such as the Super Bowl. Rare events include hijackings, hurricanes, major equipment outages, etc. **Figure 156** illustrates the notional way these days partition the year.

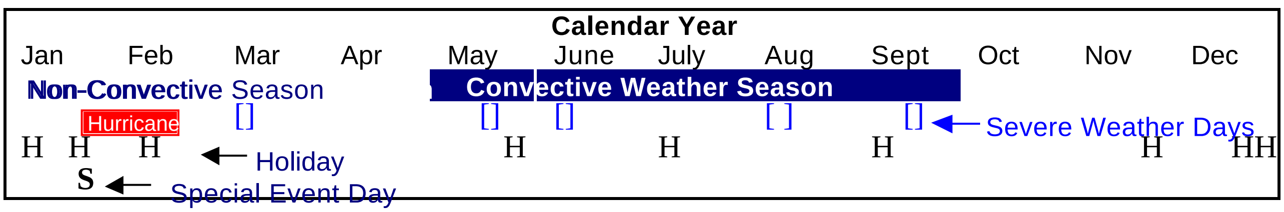

Figure 156. General partitioning of the calendar year based on special days in the NAS.

There is increased air travel around the following **holidays**:

Year 2000	Year 2001	Year 2002	Holiday
Jan. 1	Jan. 1	Jan. 1	New Years Day Holiday
Jan. 17	Jan. 15	Jan. 21	Martin Luther King Holiday
Feb. 21	Feb. 19	Feb. 18	Presidents Day
May 29	May 28	May 27	Memorial Day
July 4	July 4	July 4	Independence Day
Sept. 4	Sept. 3	Sept. 2	Labor Day
Nov. 23	Nov. 22	Nov. 28	Thanksgiving Holiday
Dec. 25	Dec. 25	Dec. 25	Winter Holiday

To address known special days, the ATCSCC has a list of **Special Traffic Management Programs** (STMPs) for special events for one year. This list changes every year as new events arise or a previous year's events do not happen. There is also an online database of events occurring during the current month (dates and which airports are affected). These programs tend to affect mostly smaller airports. The following **special events** may affect total NAS operations:

July 26 – Aug. 1, 2000	Oshkosh Aviation Industry Event (Oshkosh, WI)
Jan. 30, 2000	Super Bowl Football Game (Atlanta, GA)
Jan 28, 2001	Super Bowl Football Game (Tampa, FL)
July 24 – 30, 2001	Oshkosh Aviation Industry Event (Oshkosh, WI)

Feb. 3, 2002	Super Bowl Football Game (New Orleans, LA)
Feb. 8-24, 2002	Winter Olympics (Salt Lake City, UT)
April 7-13, 2002	Sun-n Fun Fly In (Lakeland, FL)
July 23 – 29, 2002	Oshkosh Aviation Industry Event (Oshkosh, WI)

The following special events are also included in the STMP list, but not analyzed in this report:

1. Sun Valley Ski Event
2. Chicagoland NASCAR
3. Brickyard 400
4. Pepsi 400

Such events typically affect the NAS only locally, and we do not include them in our analysis.

The following **rare event days** have special air travel problems associated with them that are tested to identify if they are statistically outside the normal distribution:

- Jan. 31, 2000 Alaska Airlines MD-83 crashes into Pacific Ocean off of Los Angeles, CA
- Sept. 14, 2000 Hurricane Gordon
- Sept. 15, 2000 Tropical Storm Helene
- Oct. 19, 2000 ZLA radar equipment outage resulted in major ground stop
- June 5, 2001 Hurricane Allison (Texas)
- Aug. 2, 2001 Tropical Storm Barry (Florida)
- Sept. 11, 2001 Multiple Hijackings/National Tragedy
- Sept. 14, 2001 Tropical Storm Gabrielle (Florida)
- Nov. 12, 2001 Turbulence forces a tail to break apart on an Airbus A-300 over Belle Harbour, NY
- July 4, 2002 Security Problem related to an shooting at a terminal in LAX

6.2 Analysis of Special Days in the NAS

Figure 157 through **Figure 159** illustrate daily NAS operations for 2000 through 2002 from OPSNET. Significant special days and events are highlighted and the median number of operations is given. The red line is the result of the robust fit regression. The robust fit method is less sensitive to outliers than the least squares method. The median and inter-quartile range was used to determine the measures of central tendency and dispersion, respectively, of the number of operations for the year. These descriptive statistics are more robust than the average or standard deviation. The green lines (top to bottom) represent the median plus the inter-quartile range, the median, and the median minus the inter-quartile range.

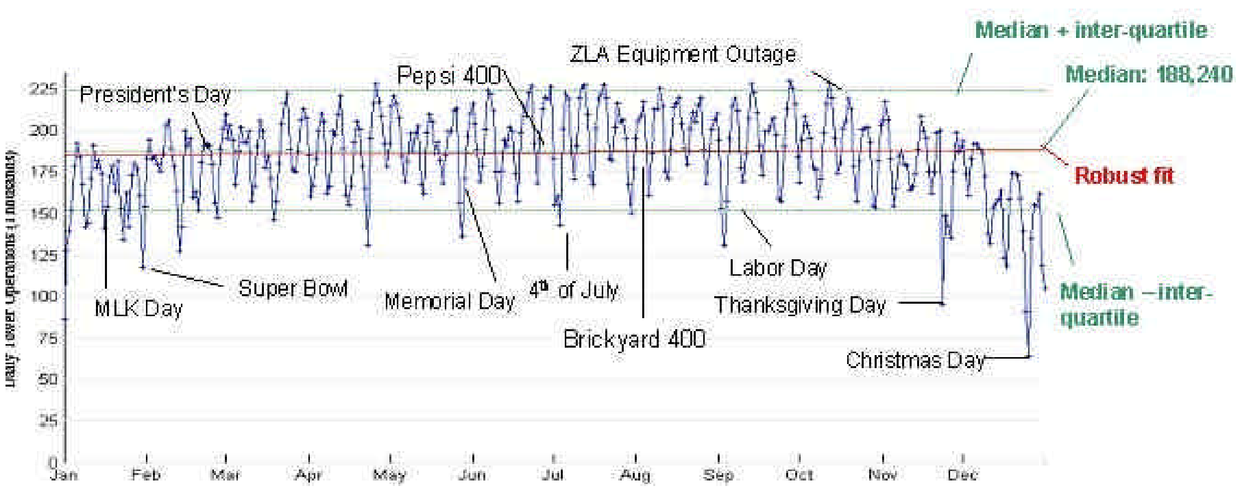

Figure 157. Daily NAS Operations and featured days for 2000.

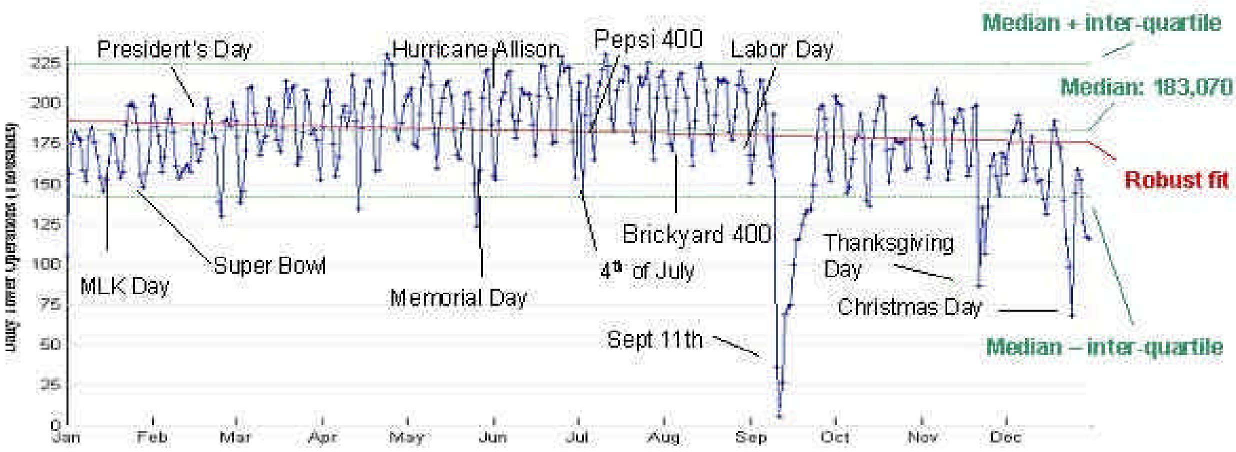

Figure 158. Daily NAS Operations and featured days for 2001.

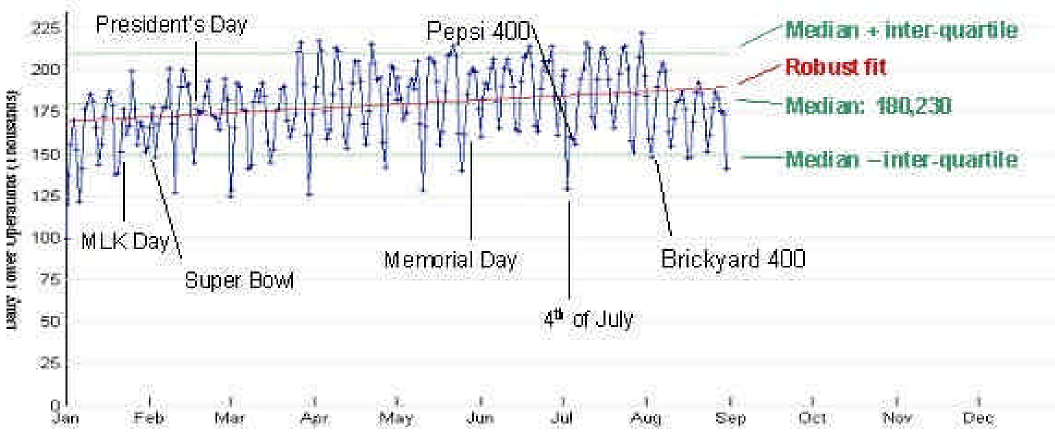

Figure 159. Daily NAS Operations and featured days for 2002.

The median number of operations is 188,240 for 2000, 183,070 for 2001, and 180,230 for 2002 (first 8 months).

6.2.1 Severe Weather Days

Our discussions with ATCSCC traffic flow specialists revealed that they view the convective weather season (May - September) as having very different behavior from days in the non-convective weather season. Convective weather (e.g., thunderstorms) is highly unpredictable, and highly disruptive. Even when the area that will be hit by convective activity can be accurately predicted, there is no way to know exactly where and when storm cells will pop up. This makes it hard to predict capacity of the affected regions of airspace. In the case of convective weather season, we let the cluster analysis of **Chapter 5** decide if the data should be partitioned. This was one of the breaks that naturally occurred.

Severe weather events that occurred in the time period of this study had little or no effect on overall NAS operations. For example, Hurricane Allison had no obvious effect on operations. The hurricane hit Texas hard but did not affect the overall operations of the NAS. This means the number of operations affected was too small to notice in the statistics of our study. A closer study of flights in and around Fort Worth and Houston Centers will probably reveal decreased operations. There was likely a plan in place when this event occurred and the affects of the storm were localized. This probably contributed to there not being any notable effect on NAS-wide operations.

Winter storms are also known about in advance, but, occasionally, these storms can hit an area harder than planned. Note that winter storms are characteristically different from summer storms in that winter storms are rarely convective. Also, winter storms tend to be at lower altitudes, and aircraft often navigate over such storms.

6.2.2 Holidays

For each observed year, the operations on holidays are low but increase on the surrounding days. Passengers are usually at their destination on the actual holiday and therefore the total number of operations falls on those days. ATCSCC specialists noted that Thanksgiving travel is more predictable than Christmas travel because Thanksgiving always falls on a Thursday and travel tends to occur between the Wednesday before and the following Sunday. Not only is there an increase in travel the day after Thanksgiving, but also there is a decrease on the Saturday after followed by another significant increase on Sunday (recall **Figure 157** and **Figure 158**). Thanksgiving operations decrease by approximately 50% on the holiday and increases approximately 50% the day after. Christmas sees a more gradual decrease in operations beginning 4 to 5 days before the holiday. Operations increased more than 100% the day after Christmas. These are considered as special travel days.

For holidays such as Martin Luther King, Memorial Day, and Labor Day, the valleys in the graph occur on the days before the actual holiday, a Sunday. These holidays always fall on Mondays. Sunday is a light travel day in general so it is fitting that the total number of operations is low. The graphs above show similar behavior between 2000, 2001, and 2002 operations around these holidays, i.e. peaks and valleys tend to occur at the same point in time. There is an increase in travel the Thursday and Friday before the holiday, on the holiday, and the day after the holiday. Observe that operations on these holidays fell well below the median.

6.2.3 Special Events

Special events like the Super Bowl, NASCAR, and Oshkosh have different effects on the NAS. Note that a STMP may be pre-emptively avoiding congestion for these events. Furthermore, these events affect the NAS locally around the location of these events, and thus, they may not affect (statistically) the entire NAS. Operations on Super Bowl Sunday fall well below the median while Pepsi 400 Saturday sees almost average operations. The Sunday of the Brickyard 400 finds slightly above average operations. The Oshkosh Event takes place over a week; the number of operations appears to follow the usual trend of NAS volume for that week. For these events, it must be determined whether or not the actual events are the reason for the slow traffic or because it is a Sunday or Saturday, which normally sees lower operations numbers.

6.2.4 Rare Events

Rare events have different effects on total operations. Some are planned for, such as hurricanes and tropical storms (covered in the section on **Severe Weather Days**). Other rare events are unexpected, such as equipment failures and hijackings. Within the time period of this study, an equipment outage, the result of a computer software upgrade, occurred in Los Angeles Center on 19 October 2000. There were two outages that day: 6:50 am to 8:30 am and 9:00 am to 10:30 am. The outage had no statistically significant effect on total operations for the day. There was backup equipment being used during these outages also, so this could contribute to the minimal impact on operations. There was actually a 19% increase in operations from the yearly average.

The events of September 11, 2001 had an overwhelming effect on operations that day and several days after. The hijackings took place in the morning (East Coast Time), which is the beginning of daily operations throughout the NAS. A NAS-wide GS was issued along with en route flights being forced to land. Thousands of flights never got underway. Operations decreased 97% between September 10 and 12, 2001.

6.3 Analysis Results for Special Days in the NAS

In order to determine which days had a significant impact on operations, special days and events that fell outside of one inter-quartile range were considered to have greatly impacted operations. Special days and events such as Martin Luther King Day and the Super Bowl make a significant impact on NAS operations. There is a noticeable increase or decrease in operations around these days. As stated earlier, Martin Luther King Day, Labor Day, and Memorial Day behave similarly. These holidays take place on Mondays and weekend travelers tend to extend their trips one day and return on the holiday. Therefore, operations increase on the actual holidays and a few days before the holidays. President's Day has almost no impact on operations.

Thanksgiving and Christmas operations also behave similarly in that they greatly impact the NAS. The days before see just above average operations. Operations considerably decrease on the actual holiday and significantly increase the days after.

Rare events that are specific to small regions of the country also do not greatly affect operations. Examples are hurricanes and equipment outages. Obviously, a rare event such as September 11th drastically affected operations. The effect is seen on the day of the event and several days after. Special events, such as racing events and Oshkosh, have little or no effect on operations.

Table 22 summarizes all of the above findings. The data is classified as follows:

- **High:** Data falls 1 or more standard deviations from the median.
- **Medium:** Data falls between .5 and 1 standard deviations from the median.
- **Low:** Data falls .5 or less standard deviations from the median.

A "High" score means operations significantly increased in a positive or negative direction and a "Low" score means operations barely increased or decreased. A score of "NO Effect" means there was no apparent effect on operations. Recall that although the actual day may have no effect, the days surrounding an event may be affected. To summarize, the most significant Special Days in the NAS are from the Sept. 11, 2001 tragedy, the Super Bowl, and from the following re-occurring holidays: 4th of July, Thanksgiving, and Christmas.

Table 22. Summary of Special Days.

Day	Probable Overall Effect on NAS (High, Medium, Low, NO EFFECT)
MLK Holiday	Medium
Super Bowl	**High**
President's Day	NO Effect
Memorial Day	Medium
Hurricane Allison	NO Effect
4th of July	**High**
Pepsi 400	NO Efftect
Oshkosh	NO Effect
Labor Day	Low
Sept. 11th	**High**
ZLA Outage	NO Effect
Thanksgiving	**High**
Christmas	**High**

Finally, we note the relationship between the holidays studied as special days in the NAS and the cluster analysis in **Chapter 5** that yielded six types of days in the NAS. We investigated type of day the holidays tended to be classified as. There were eight holidays in year 2000, and six holidays in year 2001 (Thanksgiving and Christmas occurred after the end of our data window, September 11, 2001), for a total of 14 holidays. **Figure 159** shows the distribution of holidays within the type-of day clusters. Since holidays tend to be low traffic, quiet days in aviation, we would expect that the holidays would tend to fall in the "Very Low" GDP level bins. Indeed, this is the case. 10 of the 14 holidays fell in the "Very Low" GDP level bins. This is not surprising, since traffic tends to be lower on holidays, and lower traffic days tend to have few GDPs. The other 4 holidays are spread over the "Low", "Medium", and "High" level bins. There were no holidays in the Very High type of day bin.

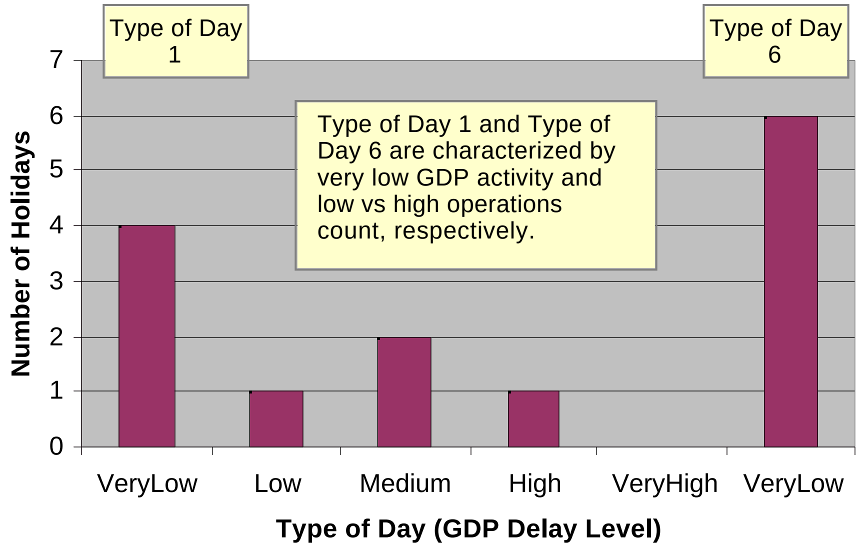

Figure 160. Holidays are typically in the two types of days characterized by low GDP levels.

7 Conclusions and Recommendations

This chapter states the conclusions and some recommendations.

7.1 Conclusions

The first and foremost conclusion of this study is that there is no single day of the year which could be described as a “typical” day in the NAS. We found that one must select the “type” of day in the NAS first before identifying the most “typical” day of that type. Hence, we identify a total of six “typical” days in the NAS, one for each of six representative “types” of days in the NAS.

We have concluded that a day in the NAS is described by a set of 8 key variables that constitute an optimal NAS feature vector. We reached this point by considering 65 NAS variables in our analysis, which statistically clustered into 8 major bundles, each bundle with a single representative variable. These variables represent the 8 variables that constitute the “optimal” feature vector for the NAS, as identified in **Table 23**.

Table 23. The Optimal NAS Feature Vector Variables.

Optimal NAS Feature Vector Variable	Description of Variable
Gate Delays	Daily Count of OAG-Based Gate Delays
Overall Delays	Total Delay Count From OPSNET
On-time Performance	Daily Total OAG-Based Airport Departure Delay (minutes)
Traffic Volume	Daily Arrival Count
Airport Performance Metric	Std Dev of Airport Performance Score (21 ASPM Airports)
Cancellations	Daily Arrival Cancellations Count
Volume-related Delays	Total Operation Count From OPSNET
Weather and GDPs	Total Delay attributed to GDPs (minutes)

The number of minutes of ground delay assigned in a GDP is the most prominent variable in characterizing the different types of days in the NAS. With the exception of "blue sky days", once the number of GDP minutes is known, a determination of how typical a day is can readily be made by comparing it to other days with similar GDP minutes. A weighted Euclidean metric (normalized for variance) was used to rank each day within a cluster as most typical to least typical; days closest to this center of the cluster were considered most typical.

Our study indicates that there are six types of days in the NAS (seven, if one is willing to count the outlier cluster):

- Type of NAS Day 1: Very low GDP level, with low operations count (cluster 1)
- Type of NAS Day 6: Very low GDP level, with high operations count, (cluster 6)
- Type of NAS Day 2: Low GDP level (cluster 2)
- Type of NAS Day 3: Medium GDP level (cluster 3)
- Type of NAS Day 4: High GDP level (cluster 4)
- Type of NAS Day 5: Very High GDP level (cluster 5)

NAS Days 1 and 6 were characteristic of no weather days, whereas Days 2 through 5 were all weather impacted.

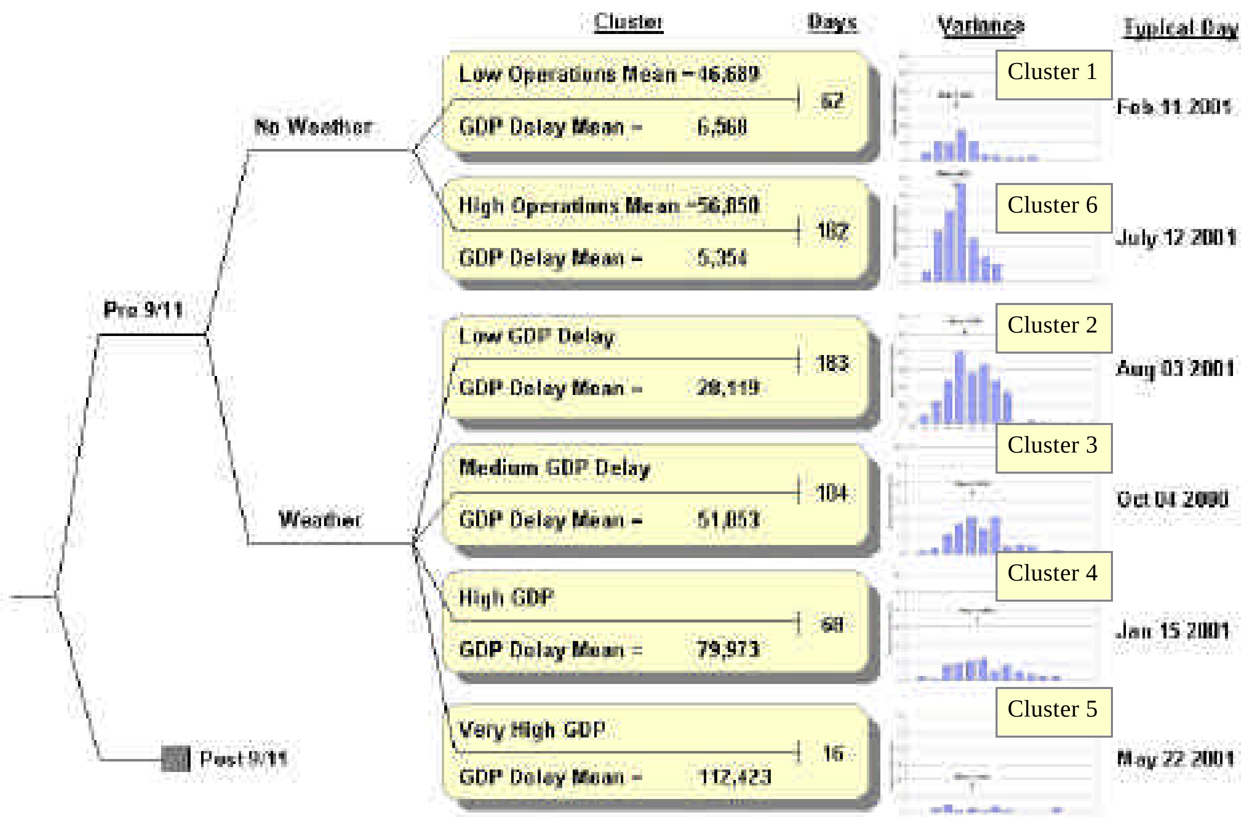

Figure 161. Decision Tree representing the Types of Days in the NAS.

The necessary distinctions between these different types of days are summarized by the decision tree, as illustrated in **Figure 161**. The use of the decision tree is as follows. Reading left to right; a decision must be made at node 1 whether a data set from pre-9/11 or post-9/11 is desired. Since our cluster analysis did not incorporate the post-9/11 data, the user is forced into the upper branch, pre-9/11. (The lower branch has been included for sake of completeness.) At decision node 2, the user decides whether to choose a data set from the collection of no weather days with very low GDP delay minutes (upper branch) or from the collection of weather days with GDP delay minutes ranging low to high (lower branch). The GDP variable is a surrogate for a collection of ground delay and weather related effects. So these branches have been informally dubbed "no weather" days and "weather" days, which are approximate, descriptive terms.

Of the holidays, Independence Day, Thanksgiving and Christmas holidays greatly impact the NAS. The Super Bowl also impacts the NAS in a similar manner to a holiday. Operations considerably decrease on the actual holiday and significantly increase the days after. Our study does not suggest that holidays need to be separated as a separate type of day in the NAS, most holidays fall within one or two of the clusters identified in the type-of-day results.

Rare events typically impact only small regions of the country, hence, they do not greatly affect operations. Examples are hurricanes and equipment outages. Obviously, a rare event such as September 11th drastically affected operations. Special events, such as racing events and Oshkosh, have little or no effect on the aggregate statistics of NAS operations.

Finally, there is an important complication during the course of this study; namely, data integrity. Certain data sources were plagued with missing records, typographical errors, incorrectly formatted entries, and poor documentation. This posed a challenge to overcome, as much of the analysis required both well-formatted and complete data sets to be of value. A fair amount of effort was required to cleanse the data, and this entailed developing software routines that would revise inconsistent records in most circumstances.

7.2 Recommendations for Conducting NAS Simulation Validations

In this study, we focused on determining a small, manageable set of data that can be used to validate NAS-wide simulations. This was accomplished by analyzing historical datasets to determine the variable dependencies and to determine the most statistically significant variables that constitute a minimal NAS feature vector size. Our analysis suggests that validations of low fidelity NAS-wide simulations should mainly focus on the 8 variables that are identified in the optimal NAS feature vector. This recommendation will potentially reduce the total quantity of data analyzed in validating a low fidelity NAS simulation. We did not investigate this issue with respect to medium and high fidelity simulations, so we refrain from making a recommendation for validating those types of simulations. When higher fidelity is added to a NAS simulation, more than just the aggregate statistics should be considered for the validation. Additionally, one must note that our recommendation assumes that there is no other variable independent of the 8 variables in the optimal NAS feature vector important to a NAS simulation validation. Our recommendation is that NAS simulation validations should consider at least those elements that constitute the optimal NAS feature vector, and if not possible, to attempt to select a variable from the same cluster set as a substitute.

The data requirements listed in **Chapter 2** demonstrate that the NAS is a very complex system with very many variables that describe it. A very small subset of these variables was studied in our analysis, and of those, the minimal set of variables was determined to define the optimal NAS feature vector. This approach is open to speculation when a new variable that was not in the original set of 65 variables is introduced. While engineering judgment was used to select a set of 8 variables that most likely characterize the behavior of the NAS, we were limited to variables that are available in historical datasets. Thus, our conclusions are limited to what can be said about how the 65 variables relate to the 8 variables of the optimal NAS feature vector. Caution must be taken when considering new variables outside the set of 65 variables in this study. In such a case, we recommend that a small scale study be performed to test if the new variable is dependent on one or more of the dominant variables in the optimal NAS feature vector. If the new variable is dependent, then it is not recommended to add the new variable to the validation dataset. If the new variable is independent, then engineering judgment should be used to determine if the new variable should be included in a NAS simulation validation.

Our type-of-day analysis indicated that NAS simulation validation sets should consider weather and GDP modeling as a basis for validation data sets. First, if neither weather nor GDP are modeled in the NAS simulation, then the results of our study indicate that there are two types of days needed to validate such simulations. They are embodied by Cluster 1 and Cluster 6. If weather and GDPs are include in the NAS simulation – indeed, if weather is included then GDPs must also included – then, depending on the degree of simulation validation that is desired, there are several choices to be made. One can validate a NAS simulation with modeled weather and GDPs through Cluster 2 through Cluster 5 (with special attention to the lower membership size of Cluster 5). Furthermore, trends may be simulated by comparing pairs of clusters, e.g., (2, 3) vs. (2, 4) vs. (2, 5), each having a difference in magnitude of weather and GDP significance in the validation data sets. For a complete validation of a NAS simulation, simulation developers should validate their NAS simulations with at least one validation run from each type of day in the NAS.

7.3 Recommendations for Future Research

In this study, performance statistics were collected and assessed on a NAS-wide level. An area of future study would be to apply a geographical component to the study. Questions of interest are:

- Can a region of the country serve as an indicator of overall NAS behavior? For instance, if delays are high in the northeast, does this mean that delays are high all over the NAS? Can we collect performance and delay statistics strictly in one region (e.g., the northeast) to assess the overall condition of the NAS?
- In terms of performance metrics, what is the most natural decomposition of the NAS into regions? Does this decomposition coincide with the ARTCCs?
- Are there any local anomalies (e.g., in weather and delays) severe or noteworthy enough to be significant drivers of NAS-wide statistics?
- Our reduction of the size of the NAS feature vector to a set of 8 variables greatly simplified the amount of data that was analyzed in the final cluster analysis. However, we still gather statistics over the ASPM-50 airports whenever possible. This leads us to the following questions: What is the smallest number of airports whose performance is a reasonable surrogate for NAS-wide airport performance? And which airports are these?
- Are there days when a small, local weather disturbance causes big problems? For example, can a small isolated storm over Chicago, IL cause NAS-wide problems? To what degree does fog in San Francisco, CA affect the NAS?

A major theme which was absent – and which could be taken up in future research – is the concept of a cause-and-effect chain. For instance: A, B, and C cause D; B and D cause E; A, B, C, and E cause F. With such a network, we might be able to identify that certain variables may be treated as dependent variables in one instance, but as independent variables in another.

8 References

[A73] Anderberg, M. R., *Cluster Analysis for Application*. New York: Academic Press, 1973.

[C01] Cherniavsky, E. A., "Aviation System Performance Metrics: Airport Utilization," MITRE-CAASD document, Project Number 02001207-20, April 2000.

[FAA98] FAA, *Air Traffic Control*, Tech. Document 7110.65, Washington, DC, Feb. 1998.

[FAA99] FAA, Documentation for Airport Utilization Metrics, Office of Aviation Policy and Plans, FAA, Washington, DC, December, 1999.

[FAA02] FAA, *Documentation for the Aviation System Performance Metrics (ASPM) Actual versus Scheduled Metrics*, Office of Aviation Policy and Plans, FAA, Washington, DC, May, 2002.

[G99] Gordon, A.D., *Classification*, Chapman and Hall, 1999.

[H75] Hartigan, J.A., *Clustering algorithms*, Wiley, New York, 1975.

[H76] Harman, H. H, *Modern Factor Analysis*. Chicago University Press, 1976.

[KCH97] Klopfenstein, M., Cook, D., and Hsu, R., *An Initial Look at the Value of Dynamically Contributed Airline Schedule Data to Air Traffic Management*, Metron, Inc., August, 1997.

[KGL99] Klopfenstein, M., Gallus, D., Lambert, C., Bonham, D., Smith, P., Chapman, R., and Obradovich, J., *En Route User Deviation Assessment*, Metron, Inc., November, 1999.

[No94] Nolan, M.S., *Fundamentals of Air Traffic Control*, Wadsworth Pub. Comp., Belmont, CA, 1994.

[R71] Rand, W. M., “Objective criteria for the evaluation of clustering methods,” *Journal of the American Statistical Association*, Vol. 66, pp. 846-850, 1971.

[Ra02a] Raytheon Company, *Phase 2 Modeling System Requirements Document*, Technical Report Deliverable CTOD 7.15-1, Contract NAS2-00015, June, 2002.

[Ra02b] Raytheon Company, *Phase 2 Airspace Concept Evaluation System (ACES) Assessment and Validation Plan*, Technical Report, Contract NAS2-00015, July, 2002.

Appendix A: Notation for Statistics

Figure 162 illustrates the notation that is implicitly used throughout this report to describe the statistics within graphics.

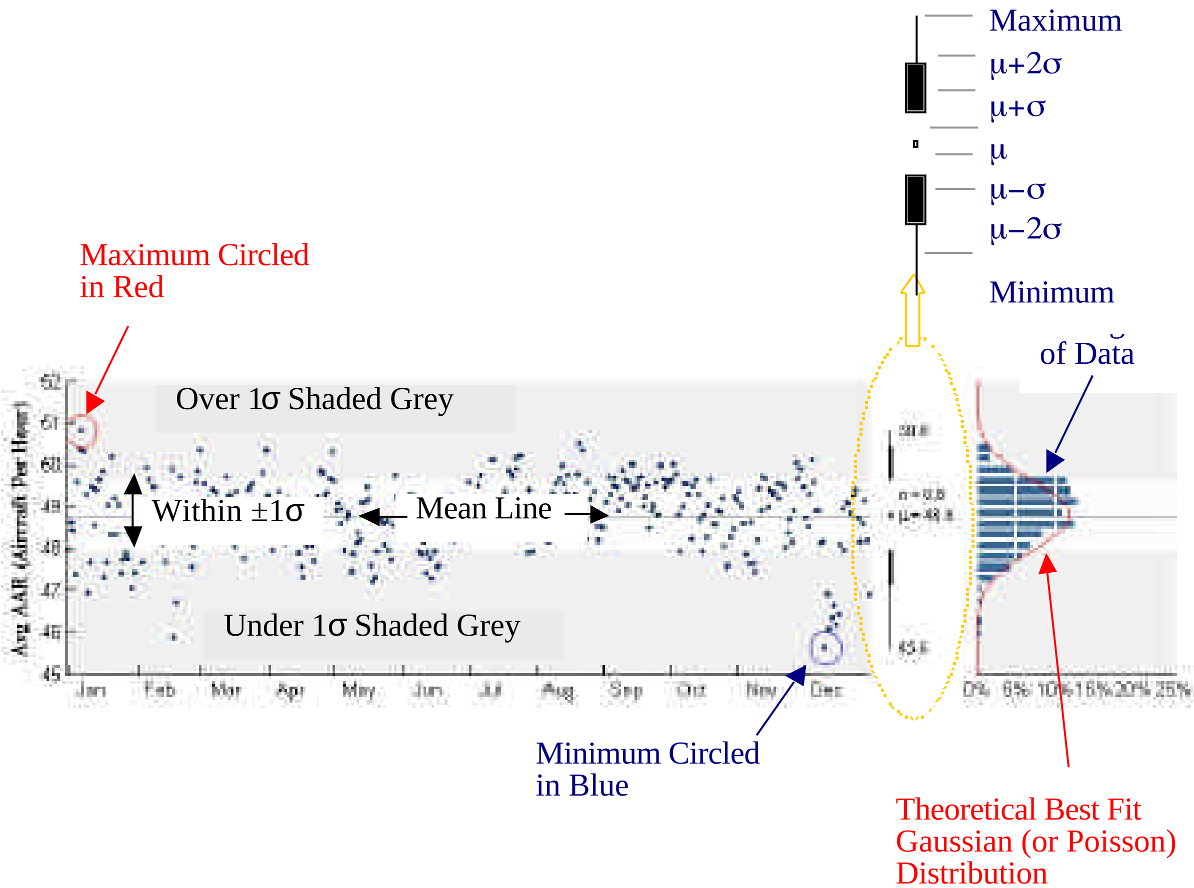

Figure 162. Notation used in statistical plots.

Appendix B: Fleet Mix

This appendix investigates fleet mix in terms of user class and weight class data for the year 2001. Insufficient data were available for a study of year 2000 data, so these variables were not included into the cluster analysis of this report. These data are presented here for additional understanding of the NAS. Overall, fleet mix stays relatively constant from day to day, which eliminates the fleet mix as a useful variable for the cluster analysis. Furthermore, the fleet mix data could not be used in the cluster analysis due to missing data.

Figure 163 depicts the user class totals for 2001. User classes are defined as Air Cargo/Freight, Air Taxi, Commercial, General Aviation (GA), Military, and Other. The details of these user class statistics are given in **Figure 164** through **Figure 169**. The "Other" aircraft is assigned to aircraft, which were not classified in the POET ETMS database. There are 49 days worth of data missing from the user class data, excluding the military class. The military class is missing 162 days. There are some outliers that are inexplicable at this point in time.

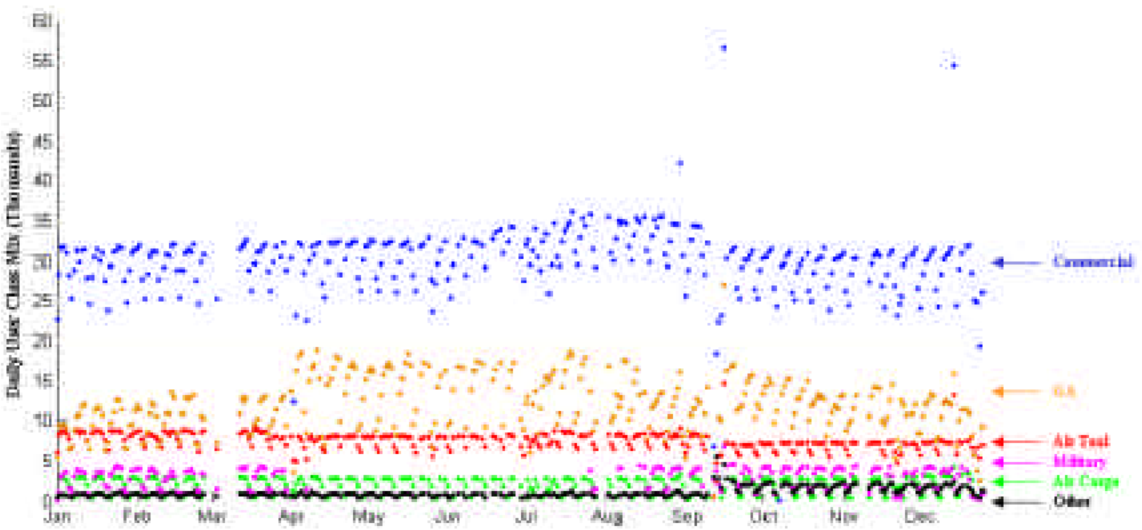

Figure 163. User Class Totals for the year 2001.

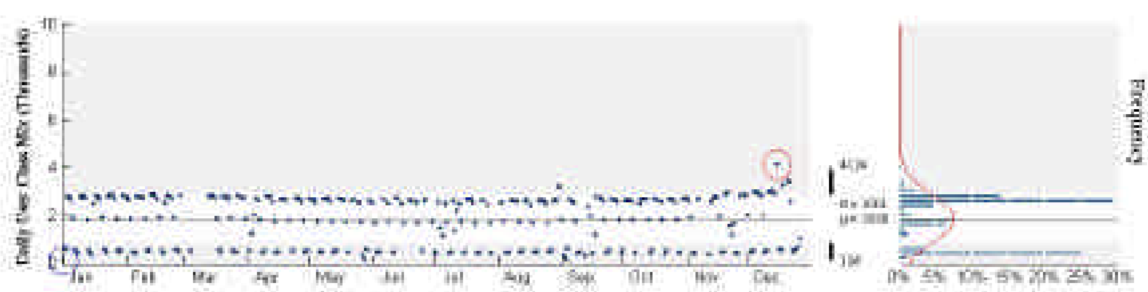

Figure 164. Air Cargo User Class Totals for the year 2001.

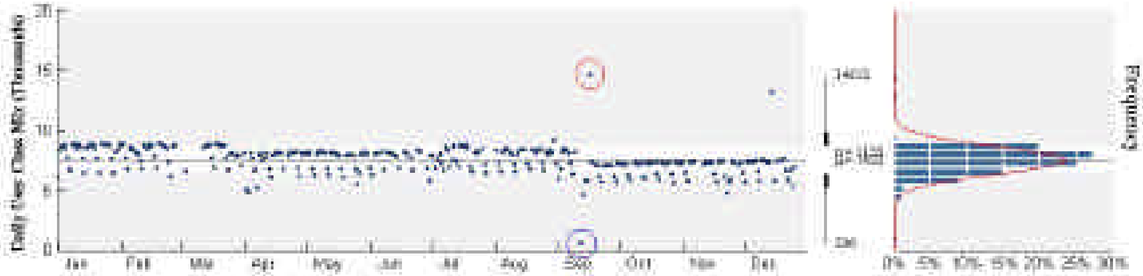

Figure 165. Air Taxi User Class Totals for the year 2001.

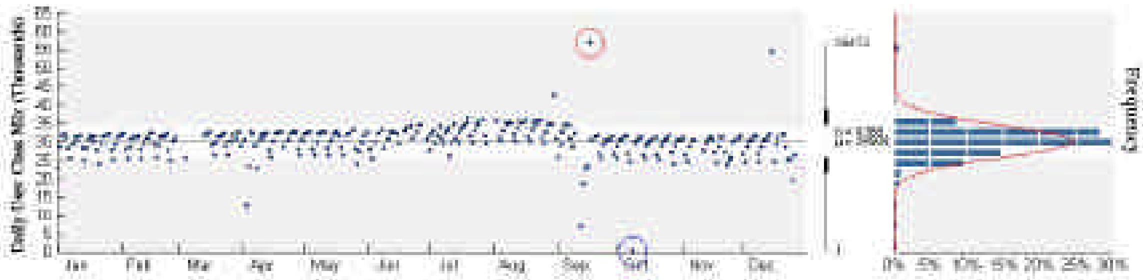

Figure 166. Commercial User Class Totals for the year 2001.

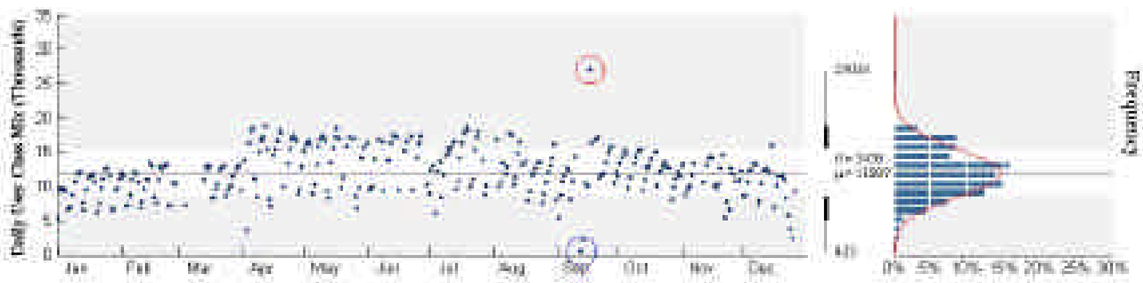

Figure 167. GA User Class Totals for the year 2001.

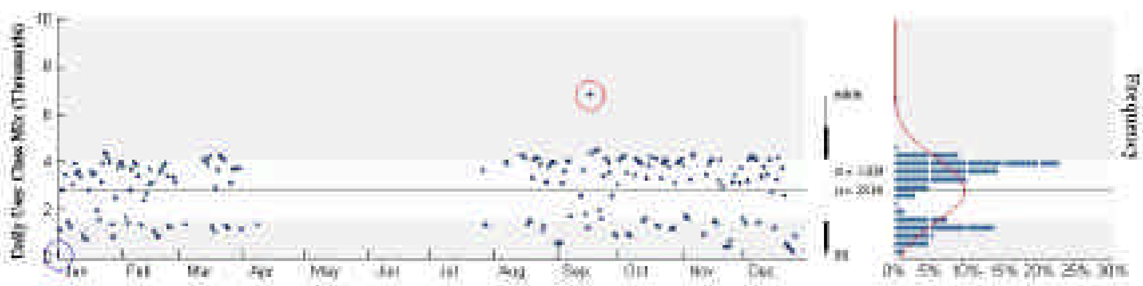

Figure 168. Military User Class Totals for the year 2001.

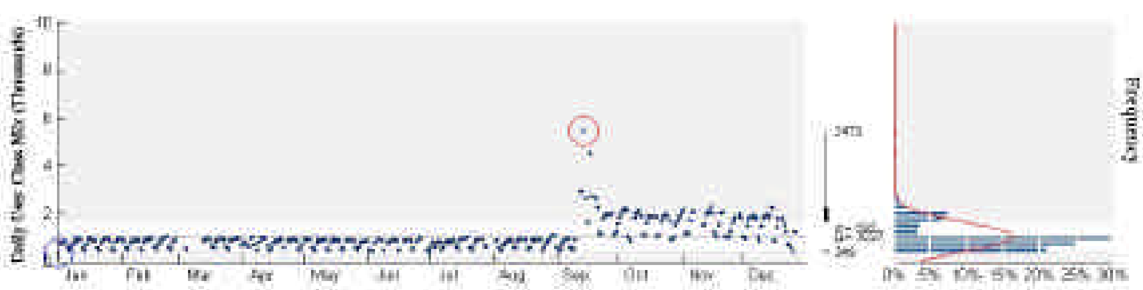

Figure 169. Other User Class Totals for the year 2001.

Figure 171 depicts the weight class (separation class) totals for 2001. Weight class totals are defined as Heavy, Large, and Small **Figure 171** through **Figure 174** illustrates detailed plots for the classes. For weight class, when there is no record, a Null weight class is assigned. There are 50 missing days of data from the weight class data.

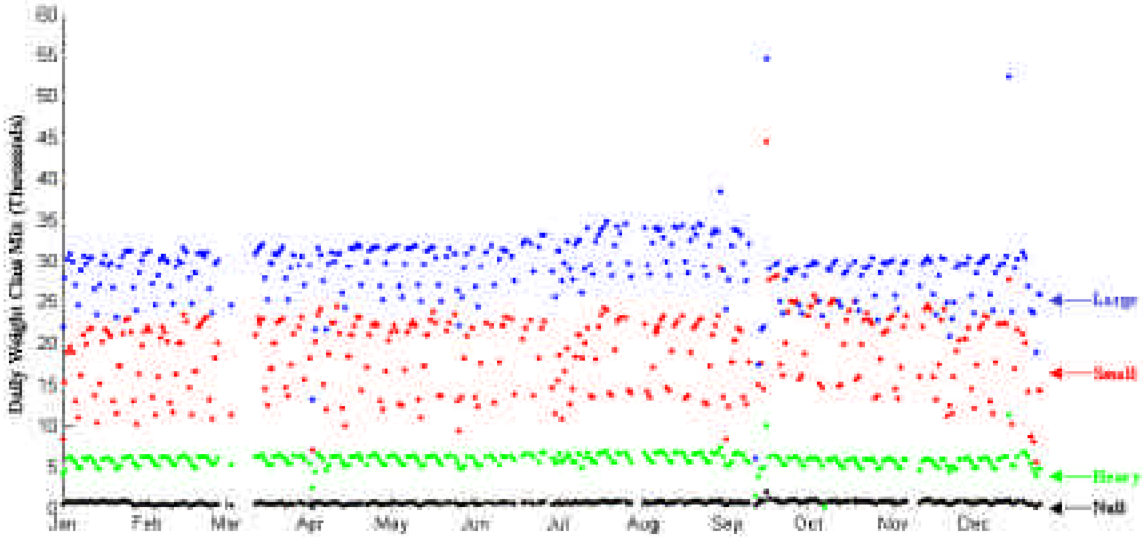

Figure 170. Weight Class Totals for the year 2001.

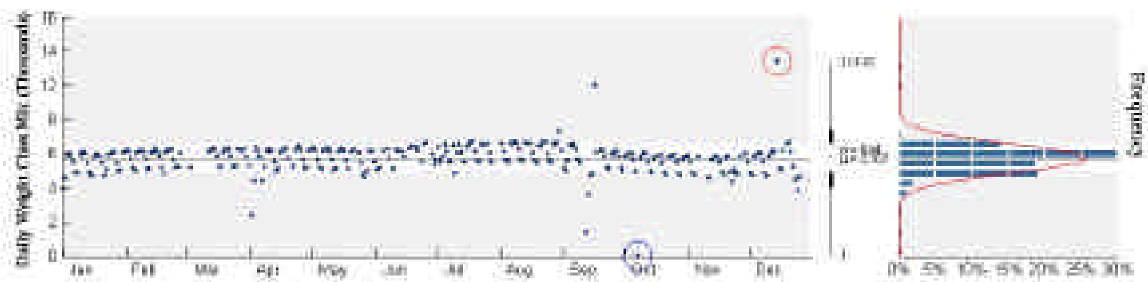

Figure 171. Heavy Weight Class Totals for the year 2001.

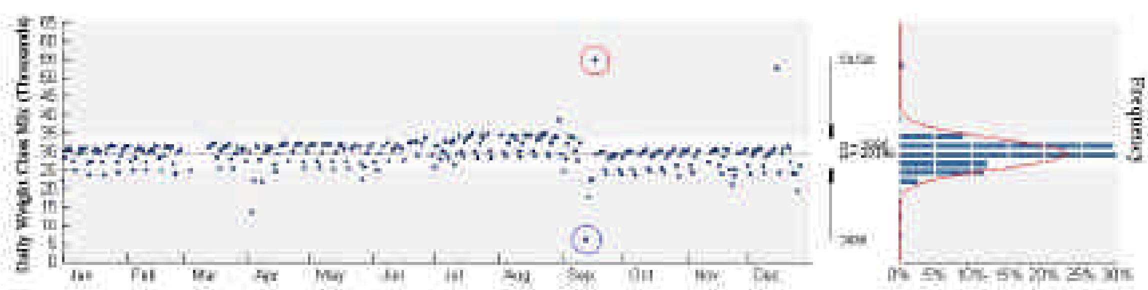

Figure 172. Large Weight Class Totals for the year 2001.

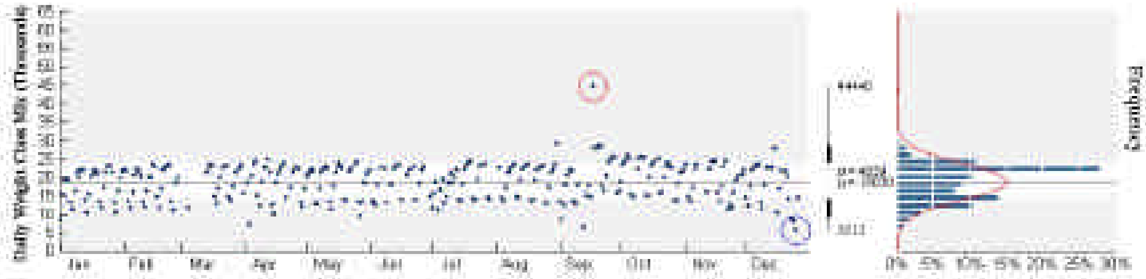

Figure 173. Small Weight Class Totals for the year 2001.

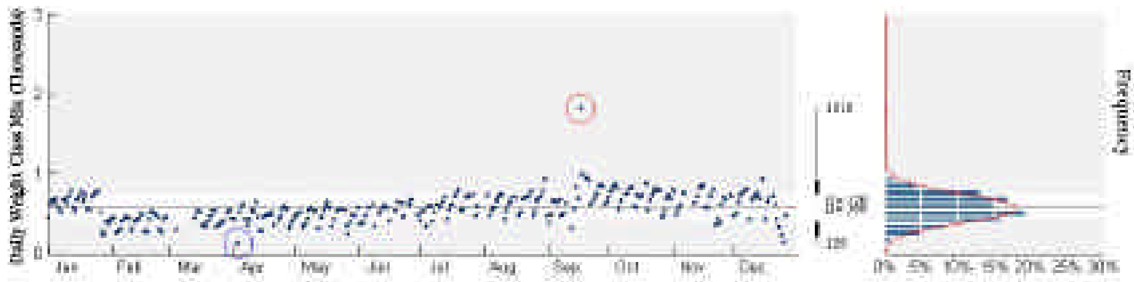

Figure 174. Null Weight Class Totals for the year 2001.

Appendix C: Cancellations and Weather

This appendix investigates the relationship between cancellations and weather data.

The Post Operations Evaluation Tool (POET)-archived ETMS database was used to query data from 21 major airports for January 29, 2002. The cancellation statistics may differ from those reported in ASPM, since ASPM records cancellations for 50 major airports. However, the 21 major airports in this plot are the largest hub airports in the NAS. Note that there are two groups of departure cancellations, one around 15 hours and one about 2 hours prior to departure (see **Figure 175**). Although the airlines may submit cancellations days or even weeks before departure, Volpe does not synchronize the airline CDM data with ETMS data until 15 hours before departure time. This explains the large spike in the cancellation data at 15 hours prior to departure, which accounts for approximately 38% of the cancellations. These plots also show that flights are often cancelled at or after scheduled departure times.

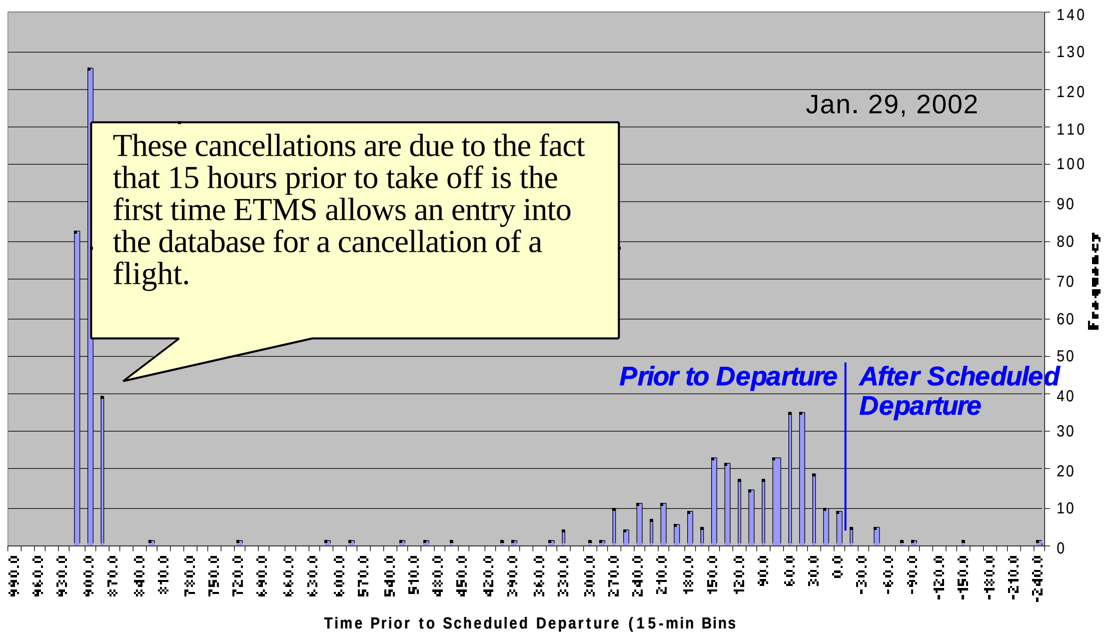

Figure 175. Cancellation times relative to scheduled take off time grouped by 15-minute bins.

Furthermore, one can track the cancellation statistics day by day to see how the statistics of the mean and standard deviation differ. To illustrate this point, the POET archived ETMS database was used to query data from 21 major airports for the week of January 28 - February 1, 2002. **Figure 176** through **Figure 181** illustrate the cancellation statistics for this time period. This set of cancelled flights does not include those flights that are considered "cancelled but flew", or regular controlled flights. A "cancelled but flew" flight is a flight that received a cancellation message and still flew. This happens if a flight was cancelled but ETMS received an activation message within a certain time of the predicted departure time. An airline will issue cancellation messages when substituting flights during a GDP. Also, flights that are diverted to an alternate destination will receive a cancellation message. These are examples of flights that received cancellation messages but flew. Note that the peak number of cancellations for this period is highly related to the location of the weather activity.

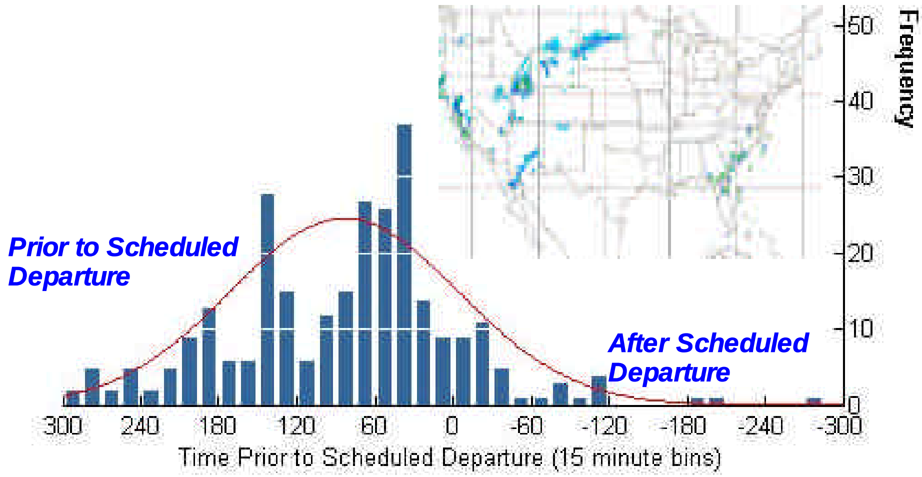

Figure 176. Cancellation times relative to scheduled take off time for Jan. 28, 2002.

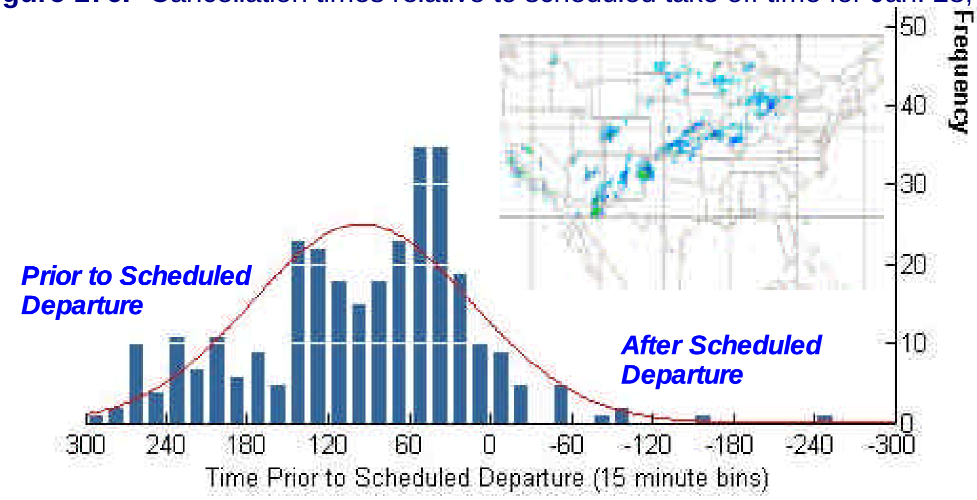

Figure 177. Cancellation times relative to scheduled take off time on Jan. 29, 2002.

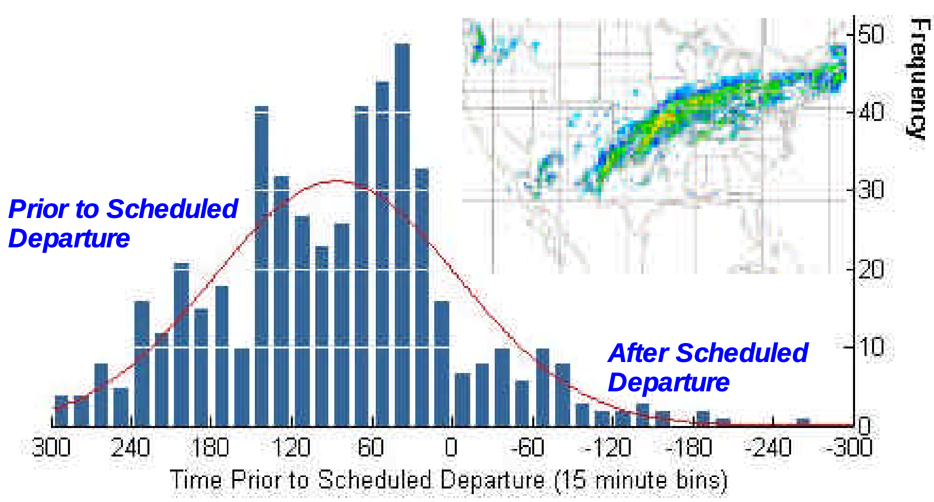

Figure 178. Cancellation times, relative to scheduled take off time for a heavy weather day, Jan. 30, 2002.

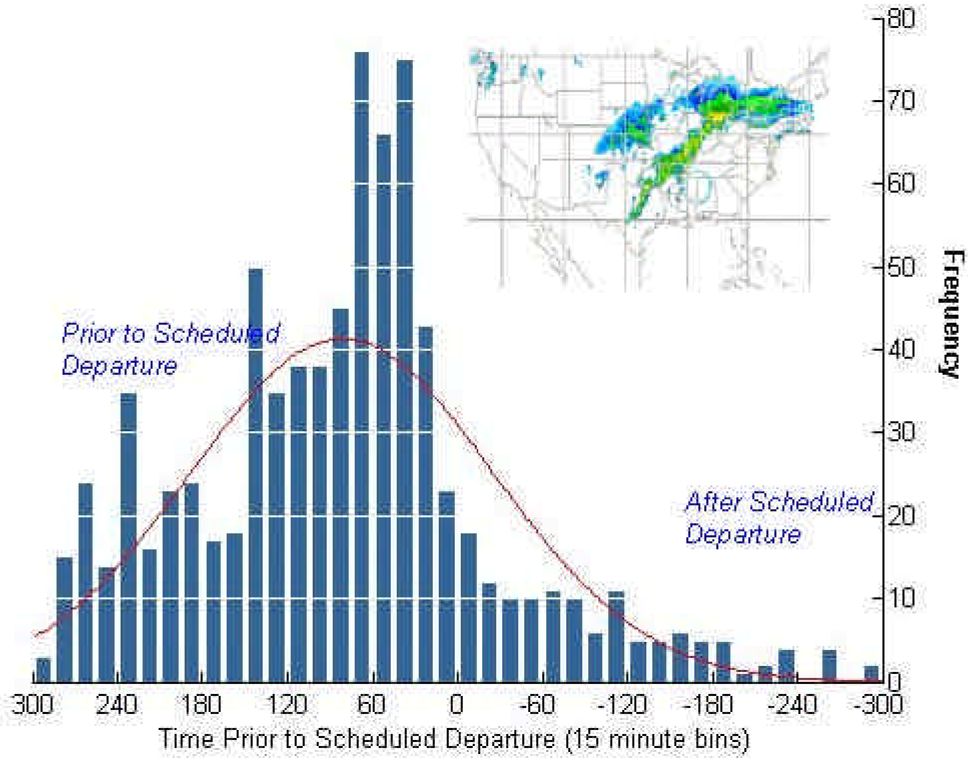

Figure 179. Cancellation times, relative to scheduled take off time for a heavy weather day, Jan. 31, 2002.

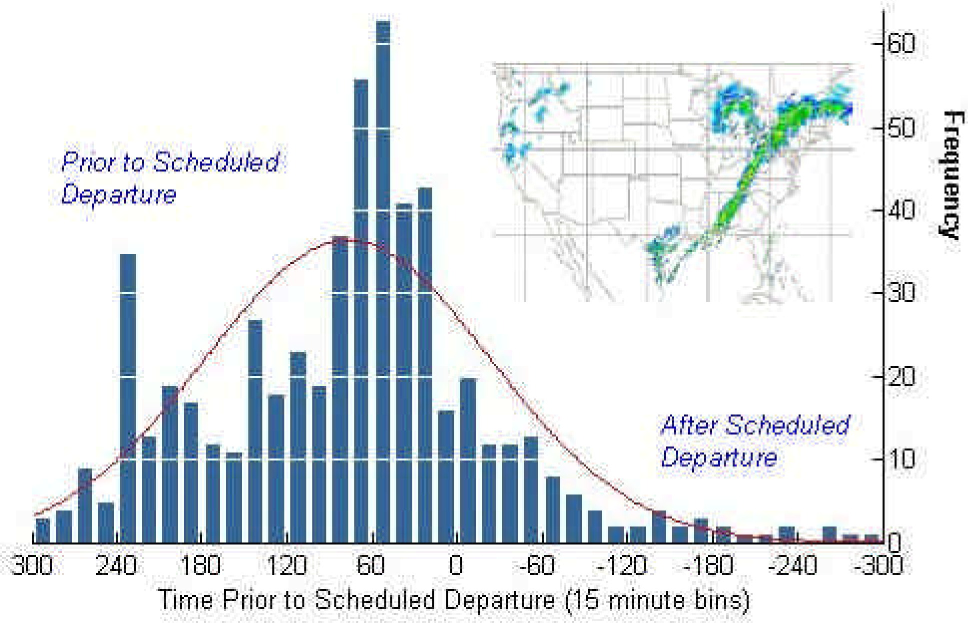

Figure 180. Cancellation times, relative to scheduled take off time for a heavy weather day, Feb. 1, 2002.

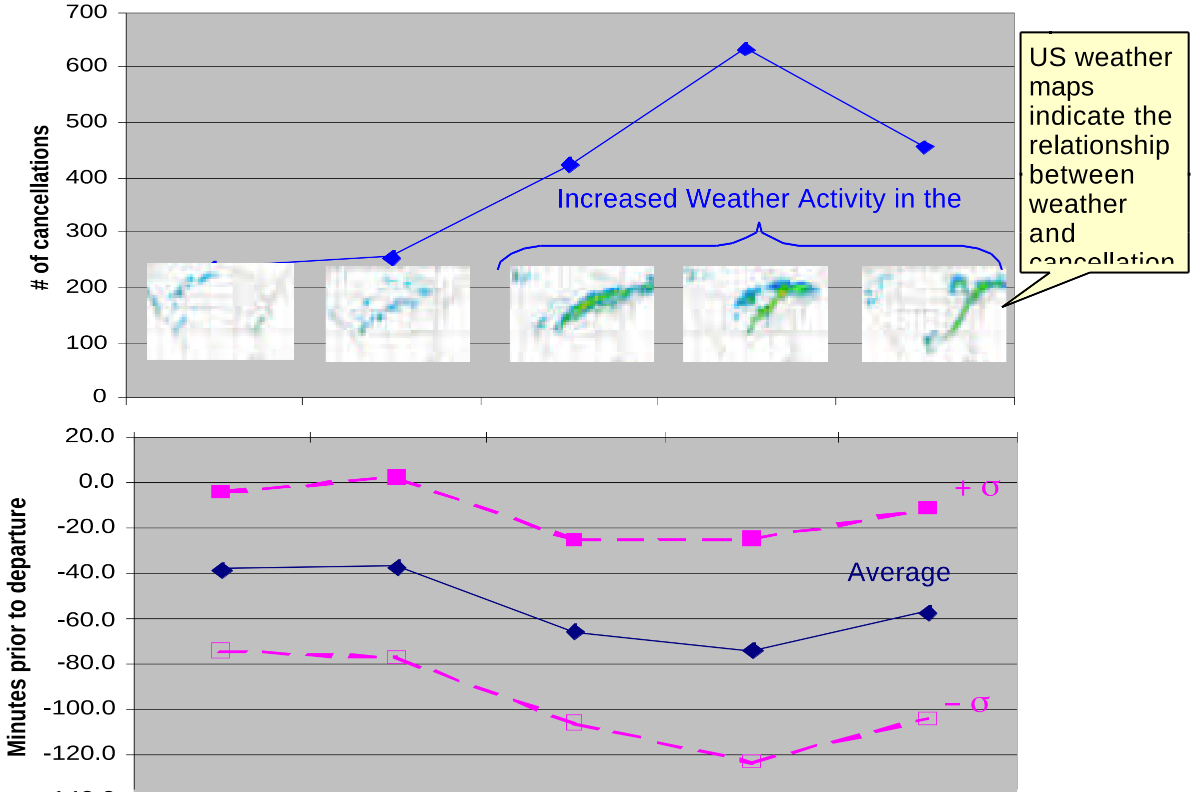

Figure 181. Cancellation statistics as they vary from Jan. 28 – Feb. 1, 2002.

Appendix D: Variable Bundling Process

The following table provides the process used for variable bundling.

Preprocessing:		Data restricted to before September 11, 2001. This removes the influential effects of outliers, which we observed on or just after September 11. There is some argument for just excluding a few months worth of data just after September 11, but we decided to err on the conservative side.	
Analysis Iteration	**Bundle Count**	**Observation**	**Variable Eliminated**
1	12	Variables *v32* and *v33* (IFR and VFR) have a perfect correlation (correlation coefficient = 1). By definition, these variables sum to 1 on each day.	*v33*
2	11	Variables *v47* and *v48* (Volume-related MIT restrictions, Total MIT restrictions) are highly correlated with cc = 0.83. One of *v47* and *v48* can be eliminated. We chose *v47* on the grounds that it is given implicitly in *v48*.	*v47*
3	10	Variables *v62* (Total Equipment Related Delay Count from OPSNET) and *v63* (Total Runway Related Delay Count from OPSNET) are in the same bundle, but are not strongly associated with the bundle. They are most likely in the same bundle because they each have a large number of outliers (e.g., days 322, 334, and 342). We decided to keep this bundle, since it explains about 1/40 of the variation in the data set.	
Note		We wish to get a bundle count to 7 or less. So, we force a reduction in bundle count. We hope that bundle 10 is eliminated, since it is not a very good indicator of NAS conditions.	
4	9	Variables *v62* and *v63* have moved into a delay-related bundle. We would like to get the bundle count down to 7 or less. This is justified by the fact that the gain in explanation of variance from the existence of the bundles 8 and 9 is less than 2%. Bundles 8 and 9 are adding 1/22 of total variation.	
5	7	The membership count of each bundle is now 5% or more. A rule of thumb in bundle analysis is that no bundle should contain less than 5% of the data points. However, we note that bundle 1 has too many data points; it contains a lot of delay statistics and a lot of GDP statistics. So, we request 8 bundles.	
6	8	Bundle 1 broke into two bundles, which is what we felt should happen on an intuitive level. The delay statistics formed one bundle, while the GDP statistics formed another.	
		We note that *v50* (Total En route Delay Count from OPSNET) is very weakly in bundle 1; it's association is so weak that it really could be in any bundle. We chose to put it into bundle 2, with the other delay-related statistics.	
		The variable *v30* (Cloud Ceilings) is weakly in bundle 3. A plot of this variable over time reveals that is has no apparent pattern. So, we eliminate *v30* from the study.	*v30*
		Bundle 2 is mostly OPSNET "Total delay" information. Seems OK.	
		Bundle 3 is pretty much airport/gate/departure delays. Seems OK.	
		Bundle 4 variables share a common attribute, traffic volume. This is good.	

		Bundle 5 is airport performance metrics (4 members). Move this later. We performed a centroid bundle analysis. Again a weak association results. IFR could be its own bundle. Decide later. No connection between CEILING and IFR, or anything else for that matter. We checked all pairs via scatter plot. Keep this as its own bundle.	
		Bundle 6 consists of cancellation statistics and two of the airport performance scores (*v35* and *v36*). Direct examination in X-gobi reveals strong association. We eliminate these two variables because they are virtually constant over time, and therefore add no value to a description of the NAS. Now, bundle 6 is just cancellation statistics.	*v35*, *v36*
		Eliminate *v32* (IFR) from the study on the grounds that is does not really belong in any of the bundles and that it would add no value to NAS description by forming its own bundle. Also, it had no correlation with the cloud ceilings variable, which was also eliminated. Weak association with any of the bundles was confirmed with a centroid bundle analysis.	*v32*
		Bundle 7 is mostly center and volume delays. Note that equipment and runway delays again have very weak associations. No home exists for these variables. We decide to eliminate *v62*, *v63*.	*v62*, *v63*
		GDP count, GDP length and GDP delay minutes are strongly related (correlation coefficients over 0.83).	
		We noticed that clusters 2 and 3, (overall delays and on-time performance, respectively.) should perhaps be the same clusters, since they are complements of each other. We plotted one representative variable (*v21* and *v50*) against the other and found it to be a data cloud, meaning that there is no correlation. This is true even when outliers are removed. So, we are comfortable with these being two different clusters. It would be nice to know, at some point, why these variables are so unrelated.	
Final Results			
		8 variable bundles are formed, each with a distinct statistical and intuitive characteristic. A representative variable was chosen for each bundle. The variable was the one with the strongest association with the bundle. Also, a name is given to each bundle based on the most common characteristic in that bundle. We note that in each bundle, the second or third most strongly associated variable could also serve as the representative, since their association was almost as strong as the most strongly associated variable.	

Appendix E: Type-of-Day Cluster Analysis Process

The following table provides the process used for the type-of-day clustering.

Preprocessing:		Data restricted to before Sept. 11 2001. We have already bundled the variables (using cluster analysis) into 8 bundles, and selected a representative variable for each bundle. This creates an 8-dimensional vector for each day. These are the data points upon which we now perform cluster analysis.
		Based on our review of each variable and our intuition, we break the data into six clusters. This number comes from a break into weather vs. non-weather days, then each of these branches breaks into three clusters: high volume, medium volume, low volume.
		First, do a relaxed cluster analysis (before forcing any breaks).
		We note that daily schedule arrival count is tri-modal, with breaks at 20,500 and 22,000 flights. On the other hand, daily arrival count is bimodal with a break at 21,000. We choose daily arrival count, because it has a stronger association in the traffic volume cluster. So, we force a traffic volume break at 21,000 flights.
		Note: why are *v46* and *v5* in different variable bundles? The scatter plot is very spread out.
Analysis Iteration	**Cluster Count**	**Observation**
1	**20**	Sizes are 2, 1, 14, 21, 1, 3, 13, 24, 63, 1, 1, 96, 83, 92, 8, 80, 21, 25, 33, 37. 10 clusters have sizes 15 and above (15 chosen arbitrarily; about 2%).
2	**10**	Recluster with a max. of 10 clusters. Cluster counts are now 1, 98, 60, 26, 3, 7, 170, 113, 81, 60. We will now try for 7 clusters, because only 7 of these have significant membership.
3	**7**	Recluster with a max. of 7 clusters. Former small ones got lumped together. 62, 183, 104, 68, 16, 182, and 4 are the cluster sizes.
		We note that *v57* (terminal volume related delay) varies as a super-linear function of *v47* (total operations from OPSNET), with increasing variance. This makes sense, since delays increase super-linearly with increasing operations; variance will increase quadraticaly as well.
		We note that *v56* (total weather related delays) and total delays are virtually the same. This just means that most delays are attributable to weather.
		Cluster 1 has (62 elements) has mostly Sunday, some Monday and Tuesday, but no Wed or Thu. This is sort of a Sunday (low volume) cluster.
		Cluster 2 has no apparent day-of-week pattern; uniform over day of week
		Cluster 3 is strangely low on Saturdays, yet generally high or extreme values.
		Cluster 4 is low on Saturday, but has many other days; this is high or extreme volume.

		Cluster 5 (15 elements) has high or medium volume but too low count to tell, maybe holidays; 8/15 days are Saturdays (medium volume). Don't know what to make of it. Cluster with 15 elements: Th 4/20/00, Tu 6/13/00, Th 7/27/00, Fr 10/6/00, Fr 10/27/00, Mo 11/6/00, Fr 12/15/00, We 3/21/01, Fr 4/6/01, Mo 5/21/01, Tu 5/22/01, Fr 6/22/01, Fr 8/10/01, Mo 8/13/01, Fr 8/31/01.
		Cluster 6 is very high on Saturday and very low on Sunday; the rest random (Monday a bit low too). Appears to be a Saturday (medium volume) cluster.
		Cluster 7 elements (only 4 of them); can't tell about day of week; elements are: Fr 7/28/00, Su 11/26/00, Su 02/25/01, Th 4/12/01.
		Discovery! Clusters are grouped by GDP minutes. Check other representative variables to see if it happens with any others.
		Same for *v16*.
		v47 (total delay count): no association with cluster variable
		v46 have some effect but not much. For *v46*, we check the first cluster.
		v5 is dead.
		v38 is dead.
		v8 has a similar pattern to *v42*, but clusters have more overlap. This is because *v8* (gate delays) is strongly related to *v42* (ground delays in a GDP).
		Note: *v56* (total weather delays) doesn't work either!
		Note: *v43* (weather related MITs) doesn't work either!
		Try re-clustering without *v42* present. We got 8 clusters with significant membership using a max. of 20. We got 7 clusters with significant membership using a max. of 10. We got 7 clusters with significant membership using a max. of 7. Now, we no longer have clustering by *v42*. This means *v42* is a natural clustering variable. We note that *v8* has a little structure (like before) but not distinct. No other dominant variable has structure to it.
Final Results		
		There are 6 types of days in the NAS; each of these is primarily characterized by its level of GDP minutes, with the caveat that there is some overlap between these clusters. That is, a day with level X of GDP minutes could be in one of two clusters. So, the other variables 7 (out of 8) are still needed to fully characterize a day. We are still investigating which (if any) subset of the other seven variables distinguishes these ambiguous cases.
		Note: we actually had 7 clusters in the final analysis, but cluster 7 had only 4 outlier data points, so we disregard it.
		Note: this means that the same type of analysis will most likely work on the entire data set, that is, Jan. 2001 to date, rather than chopping off the data at Sept. 10, 2001.
		One can characterize a day in the NAS by adding it to the existing data set of observations and using a cluster analysis. Knowledge of just GDP minutes on a day is not sufficient, though it would narrow the possible clusters down to 2.

Appendix F: Airport Identifiers

This appendix identifies 3 letter airport identifiers (not all are referred to in this report).

Code	Airport Name	Location
ABY	SOUTHWEST GEORGIA REGIONAL	ALBANY, GEORGIA
APF	NAPLES MUNI	NAPLES, FLORIDA
ATL	WILLIAM B HARTSFIELD ATLANTA INTL	ATLANTA, GEORGIA
BFI	BOEING FIELD/KING COUNTY INTL	SEATTLE, WASHINGTON
BFL	MEADOWS FIELD	BAKERSFIELD, CALIFORNIA
BNA	NASHVILLE INTERNATIONAL	NASHVILLE, TENNESSEE
BOS	GENERAL EDWARD LAWRENCE LOGAN INTL	BOSTON, MASSACHUSETTS
BUF	GREATER BUFFALO INTL	BUFFALO, NEW YORK
BWI	BALTIMORE-WASHINGTON INTL	BALTIMORE, MARYLAND
CLE	CLEVELAND-HOPKINS INTL	CLEVELAND, OHIO
CLT	CHARLOTTE/DOUGLAS INTL	CHARLOTTE, NORTH CAROLINA CAROLINA
CRP	CORPUS CHRISTI INTL	CORPUS CHRISTI, TEXAS
CVG	CINCINNATI/NORTHERN KENTUCKY INTL	COVINGTON/CINCINNATI, OH,KENTUCKY
CYYZ	LESTER B. PEARSON INTL	TORONTO,ONT,CANADA
DEN	DENVER INTL	DENVER,COLORADO
DET	DETROIT CITY	DETROIT,MICHIGAN
DTW	DETROIT METROPOLITAN WAYNE COUNTY	DETROIT,MICHIGAN
EUG	MAHLON SWEET FIELD	EUGENE,OREGON
EWR	NEWARK INTL	NEWARK,NEW JERSEY
FNT	BISHOP INTERNATIONAL	FLINT,MICHIGAN
HRL	VALLEY INTL	HARLINGEN, TEXAS
IAD	WASHINGTON DULLES INTERNATIONAL	WASHINGTON, DIST. OF COLUMBIA
IAH	HOUSTON INTERCONTINENTAL	HOUSTON, TEXAS
JAN	JACKSON INTERNATIONAL	JACKSON, MISSISSIPPI
JFK	JOHN F KENNEDY INTL	NEW YORK, NEW YORK
LAS	MC CARRAN INTL	LAS VEGAS, NEVADA
LAX	LOS ANGELES INTL	LOS ANGELES, CALIFORNIA
LBB	LUBBOCK INTL	LUBBOCK, TEXAS
LGA	LA GUARDIA	NEW YORK, NEW YORK
MCI	KANSAS CITY INTL	KANSAS CITY, MISSOURI
MCO	ORLANDO INTL	ORLANDO, FLORIDA
MDT	HARRISBURG INTERNATIONAL	HARRISBURG, PENNSYLVANIA
MDW	CHICAGO MIDWAY	CHICAGO, ILLINOIS
MEM	MEMPHIS INTL	MEMPHIS, TENNESSEE
MFE	MC ALLEN MILLER INTL	MC ALLEN, TEXAS
MGM	DANNELLY FIELD	MONTGOMERY, ALABAMA
MHT	MANCHESTER	MANCHESTER, NEW HAMPSHIRE
MKE	GENERAL MITCHELL INTERNATIONAL	MILWAUKEE, WISCONSIN
MLI	QUAD-CITY	MOLINE, ILLINOIS
MSP	MINNEAPOLIS-ST PAUL INTL	MINNEAPOLIS, MINNESOTA
MYR	MYRTLE BEACH INTL	MYRTLE BEACH, SOUTH CAROLINA
ORD	CHICAGO O'HARE INTL	CHICAGO, ILLINOIS
PDX	PORTLAND INTL	PORTLAND, OREGON
PHX	PHOENIX SKY HARBOR INTL	PHOENIX, ARIZONA
PSP	PALM SPRINGS REGIONAL	PALM SPRINGS, CALIFORNIA
RKD	KNOX COUNTY REGIONAL	ROCKLAND, MAINE
RSW	SOUTHWEST FLORIDA INTL	FORT MYERS, FLORIDA
SEA	SEATTLE-TACOMA INTL	SEATTLE, WASHINGTON
SFO	SAN FRANCISCO INTL	SAN FRANCISCO, CALIFORNIA
SGF	SPRINGFIELD-BRANSON REGIONAL	SPRINGFIELD, MISSOURI
SJC	SAN JOSE INTERNATIONAL	SAN JOSE, CALIFORNIA
SNA	JOHN WAYNE AIRPORT-ORANGE COUNTY	SANTA ANA, CALIFORNIA
SYR	SYRACUSE HANCOCK INTL	SYRACUSE, NEW YORK
YKM	YAKIMA AIR TERMINAL	YAKIMA, WASHINGTON

Appendix G: BTS-Defined Hub Airports

The BTS (www.bts.gov) defines Large, Medium, and Small Hub airports to be the following.

Large Hubs (29 airports)

- Atlanta, GA
- Baltimore, MD
- Boston, MA
- Charlotte, NC
- Chicago, IL
- Cincinnati, OH
- Dallas/Ft. Worth, TX
- Denver, CO
- Detroit, MI
- Honolulu, HI
- Houston, TX
- Las Vegas, NV
- Los Angeles/Burbank/Long Beach, CA
- Miami/Ft. Lauderdale, FL
- Minneapolis/St. Paul, MN
- New York, NY
- Newark, NJ
- Orlando, FL
- Philadelphia, PA/Camden, NJ
- Phoenix, AZ
- Pittsburgh, PA/Wheeling, WV
- Portland, OR
- St. Louis, MO
- Salt Lake City, UT
- San Diego, CA
- San Francisco/Oakland, CA
- Seattle/Tacoma, WA
- Tampa/St. Petersburg/Clearwater/Lakeland, FL
- Washington, DC

Medium Hubs (31 airports)

- Albuquerque, NM
- Anchorage, AK
- Austin, TX
- Buffalo & Niagara Falls, NY
- Cleveland, OH
- Columbus, OH
- El Paso, TX

- Fort Myers, FL
- Hartford/Springfield/Westfield, CT
- Indianapolis, IN
- Jacksonville, FL
- Kahului, HI
- Kansas City, MO
- Louisville, KY
- Memphis, TN
- Milwaukee, WI
- Nashville, TN
- New Orleans, LA
- Oklahoma City, OK
- Omaha, NE
- Ontario/San Bernardino/Riverside, CA
- Providence, RI
- Raleigh/Durham, NC
- Reno, NV
- Sacramento, CA
- San Antonio, TX
- San Jose, CA
- San Juan, PR
- Tucson, AZ
- Tulsa, OK
- West Palm Beach/Palm Beach, FL

Small Hubs (54 airports)

- Albany, NY
- Allentown/Bethlehem/Easton, PA
- Amarillo/Borger, TX
- Atlantic City, NJ
- Baton Rouge, LA
- Birmingham, AL
- Boise, ID
- Brownsville/Harlingen/San Benito, TX
- Cedar Rapids/Iowa City, IA
- Charleston, SC
- Charlotte Amalie, St. Thomas, VI
- Colorado Springs, CO
- Columbia, SC
- Corpus Christi, TX
- Dayton, OH
- Des Moines, IA
- Fairbanks, AK

- Fayetteville, AR
- Grand Rapids, MI
- Green Bay/Clintonville, WI
- Greensboro/High Point/Winston-Salem, NC
- Greenville/Spartanburg, SC
- Guam, GU
- Gulfport/Biloxi, MS
- Harrisburg/York, PA
- Hilo, HI
- Huntsville, AL
- Indio/Palm Springs, CA
- Islip, NY
- Jackson/Vicksburg, MS
- Kailua-Kona, HI
- Knoxville, TN
- Lexington/Frankfort, KY
- Lihue, HI
- Little Rock, AR
- Lubbock, TX
- Madison, WI
- Manchester/Concord, NH
- Midland/Odessa, TX
- Moline, IL
- Myrtle Beach, SC
- Norfolk/Virginia Beach/Portsmouth/Chesapeake, VA
- Pensacola, FL
- Portland, ME
- Richmond, VA
- Rochester, NY
- Sarasota/Bradenton, FL
- Savannah, GA
- South Bend, IN
- Spokane, WA
- Syracuse, NY
- Valparaiso, FL
- White Plains, NY
- Wichita, KS

Appendix H: Data for Types of Days in the NAS

The following table presents the final clusters for the different types of days of the NAS.

Date	Day of Week	Cluster	Dist from Center
21101	0	1	1443.14
31200	0	1	1799.19
10701	0	1	2643.26
71501	0	1	2667.22
21801	0	1	2998.41
10800	6	1	3013.58
70201	1	1	3014.89
20700	1	1	3027.82
31801	0	1	3080.46
70300	1	1	3502.60
42901	0	1	3617.61
90400	1	1	3692.51
52800	0	1	3701.97
51500	1	1	3901.94
31300	1	1	4346.41
12900	6	1	4645.67
12200	6	1	5029.69
60400	0	1	5068.79
20401	0	1	5118.92
100300	2	1	5183.18
52801	1	1	5226.19
62501	1	1	5522.28
31101	0	1	5646.19
20600	0	1	5849.82
52701	0	1	5875.83
43001	1	1	6170.60
30501	1	1	6240.71
60500	1	1	6313.01
61100	0	1	6349.66
111200	0	1	6603.86
52001	0	1	6679.98
62401	0	1	6727.57
32600	0	1	6739.60
40201	1	1	6751.78
50601	0	1	6772.90
61001	0	1	6784.99
11500	6	1	6982.36
100800	0	1	7134.57
41501	0	1	7136.00
43000	0	1	7136.38
12801	0	1	7189.89
72400	1	1	7279.92
82700	0	1	7525.68
62601	2	1	7643.08
90201	0	1	7820.11
12201	1	1	7846.35
122400	0	1	7982.82
11600	0	1	8075.07
22800	1	1	8093.23
40301	2	1	8249.71
12500	2	1	8258.47
51301	0	1	8494.31
13100	1	1	8621.06
50800	1	1	9083.79
10101	1	1	9142.84
42300	0	1	10193.43
32501	0	1	10460.39
22300	3	1	11683.66
123100	0	1	12140.49
12301	2	1	14619.57
10100	6	1	15926.52
13000	0	1	16223.98
80301	5	2	1572.28
90500	2	2	1595.86
21700	4	2	2222.09
60800	4	2	2839.22
111400	2	2	3003.41
112800	2	2	3243.90
82600	6	2	3347.75
120600	3	2	3984.36
52901	2	2	3984.85
70500	3	2	4010.60
12800	5	2	4011.78
81200	6	2	4093.16
40400	2	2	4313.71
50900	2	2	4327.68
71401	6	2	4368.60
80300	4	2	4394.14
53100	3	2	4458.51
41800	2	2	4562.76
41100	2	2	4645.63
62200	4	2	4757.46
30101	4	2	4864.70
60101	5	2	5378.51
70800	6	2	5395.08
91800	1	2	5412.43
62800	3	2	5432.21
41500	6	2	5440.44

51700	3	2	5463.60	112200	3	2	7308.79
41301	5	2	5473.78	121100	1	2	7319.10
101700	2	2	5507.98	12101	0	2	7359.28
71701	2	2	5536.31	81801	6	2	7375.21
62301	6	2	5603.69	20701	3	2	7421.28
40500	3	2	5608.17	82300	3	2	7557.45
21501	4	2	5637.94	82801	2	2	7667.60
92400	0	2	5643.83	102800	6	2	7679.09
72801	6	2	5668.08	81600	3	2	7704.03
12701	6	2	5754.52	31901	1	2	7763.59
102100	6	2	5821.41	22200	2	2	7769.30
70301	2	2	5854.07	31500	3	2	7821.88
71500	6	2	5857.76	63001	6	2	7833.63
40700	5	2	5909.96	42200	6	2	7835.18
30601	2	2	5962.83	10200	0	2	7843.26
83001	4	2	5964.20	50701	1	2	8063.81
91400	4	2	5991.79	81800	5	2	8097.34
101000	2	2	6076.16	52401	4	2	8186.04
70200	0	2	6148.44	82800	1	2	8197.03
71601	1	2	6182.99	42401	2	2	8247.00
91300	3	2	6255.54	80401	6	2	8248.16
110300	5	2	6276.73	20801	4	2	8257.21
60201	6	2	6296.52	10901	2	2	8310.85
80501	0	2	6333.23	120300	0	2	8324.37
32901	4	2	6376.92	100100	0	2	8405.73
10301	3	2	6377.03	11200	3	2	8474.81
90300	0	2	6403.61	40200	0	2	8737.22
12501	4	2	6462.15	32000	1	2	8743.64
31600	4	2	6495.19	102900	0	2	8815.80
92700	3	2	6501.10	110200	4	2	8861.05
72001	5	2	6514.21	110500	0	2	8878.08
12400	1	2	6520.13	42001	5	2	8904.36
42600	3	2	6522.07	22101	3	2	8919.87
81300	0	2	6547.66	51100	4	2	9051.89
103100	2	2	6586.65	93000	6	2	9052.65
41600	0	2	6594.97	120400	1	2	9068.28
42201	0	2	6628.05	42700	4	2	9093.74
111500	3	2	6630.49	32800	2	2	9169.75
70900	0	2	6662.26	51101	5	2	9234.49
81101	6	2	6678.13	32601	1	2	9367.54
111900	0	2	6731.96	112500	6	2	9436.53
51400	0	2	6772.37	90600	3	2	9437.89
50700	0	2	6799.81	41801	3	2	9517.62
90900	6	2	6853.00	120900	6	2	9523.07
81901	0	2	6935.07	30800	3	2	9671.15
22600	6	2	6949.08	90800	5	2	9679.31
91100	1	2	6970.13	22801	3	2	9680.00
60200	5	2	6993.28	31000	5	2	9727.03
21900	6	2	7079.60	113000	4	2	9732.46

40901	1	2	9751.84	100700	6	2	12501.94
50500	5	2	9856.51	100200	1	2	12635.91
101500	0	2	9873.32	70801	0	2	12642.68
72301	1	2	9997.78	40800	6	2	12724.32
11900	3	2	10031.45	52501	5	2	12725.82
12300	0	2	10089.36	72901	0	2	13051.21
11401	0	2	10095.29	31100	6	2	13055.53
61900	1	2	10103.06	111700	5	2	13399.78
12401	3	2	10121.58	21400	1	2	13414.61
91500	5	2	10129.57	122300	6	2	15638.21
81000	4	2	10136.43	100400	3	3	2318.57
102300	1	2	10264.06	62000	2	3	2474.90
82900	2	2	10278.51	61301	3	3	3741.03
101900	4	2	10326.14	12100	5	3	4251.73
22000	0	2	10357.56	11400	5	3	4346.61
73001	1	2	10361.15	11101	4	3	4488.71
61101	1	2	10440.00	103000	1	3	4578.34
80800	2	2	10453.65	100900	1	3	4621.75
90801	6	2	10578.79	30801	4	3	4837.39
22001	2	2	10633.44	12000	4	3	4967.12
102200	0	2	10658.25	11100	2	3	4997.93
21201	1	2	10689.07	120700	4	3	5028.88
92300	6	2	10746.40	32201	4	3	5565.37
72201	0	2	10775.70	112700	1	3	5818.93
101100	3	2	10788.68	110800	3	3	5846.94
72000	4	2	10990.97	20300	4	3	5942.67
82601	0	2	11076.53	61601	6	3	5978.33
101200	4	2	11088.22	90901	0	3	6169.97
30901	5	2	11167.72	102000	5	3	6363.07
11601	2	2	11240.98	31700	5	3	6431.29
71901	4	2	11365.53	112900	3	3	6453.47
120800	5	2	11409.23	92800	4	3	6456.72
20900	3	2	11523.92	120100	5	3	6464.14
90700	4	2	11565.70	90100	5	3	6474.96
10801	1	2	11581.11	30600	1	3	6489.54
51000	3	2	11627.53	10300	1	3	6528.49
61700	6	2	11644.12	90701	5	3	6530.09
52100	0	2	11651.52	33001	5	3	6738.29
52200	1	2	11732.38	81700	4	3	6950.31
32700	1	2	11775.21	110700	2	3	7048.48
30701	3	2	11839.67	20501	1	3	7065.64
40101	0	2	11963.07	11801	4	3	7175.10
22700	0	2	12117.67	91001	1	3	7234.84
122600	2	2	12134.00	102600	4	3	7761.46
22100	1	2	12172.03	52301	3	3	7853.54
80900	3	2	12191.29	121900	2	3	8068.21
82101	2	2	12220.22	40900	0	3	8076.14
71801	3	2	12230.96	60301	0	3	8176.39
83000	3	2	12466.56	30700	2	3	8187.53

83100	4	3	8246.15	31201	1	3	11856.35
30200	4	3	8271.71	62500	0	3	12012.68
92600	2	3	8347.67	102500	3	3	12129.17
81601	4	3	8374.92	22500	5	3	13011.43
30201	5	3	8397.86	32001	2	3	13434.10
42301	1	3	8615.83	70101	0	3	13454.43
40300	1	3	8722.22	31601	5	3	13651.45
62101	4	3	8731.69	62001	3	3	14190.03
71000	1	3	8782.13	92500	1	3	14629.41
22601	1	3	8820.92	90401	2	3	14888.50
30500	0	3	8851.68	11300	4	3	15182.30
72501	3	3	8867.25	122900	5	3	15338.45
91000	0	3	8869.14	13001	2	3	15596.21
53101	4	3	8925.36	122200	5	3	16326.16
112100	2	3	8936.13	12901	1	3	18382.00
32301	5	3	9065.76	11501	1	3	19203.09
31301	2	3	9079.30	21100	5	4	2929.79
121000	0	3	9221.24	31501	4	4	2935.34
82201	3	3	9240.02	11201	5	4	4103.50
82001	1	3	9290.60	71001	2	4	4618.27
82701	1	3	9440.67	32100	2	4	4728.06
92000	3	3	9565.72	11000	1	4	4822.26
52500	4	3	9896.52	61401	4	4	5326.55
40801	0	3	9931.01	51900	5	4	5381.01
60601	3	3	10118.10	72600	3	4	5568.16
122800	4	3	10119.48	11700	1	4	5619.76
51401	1	3	10143.69	60501	2	4	5700.26
110100	3	3	10175.93	101800	3	4	5995.19
61701	0	3	10289.70	22401	6	4	6024.52
121600	6	3	10463.27	100500	4	4	6132.53
111600	4	3	10474.41	91900	2	4	6271.86
82501	6	3	10561.08	21601	5	4	6357.65
10900	0	3	10642.05	62100	3	4	6515.77
62700	2	3	10721.12	41101	3	4	6619.24
30300	5	3	10727.81	21301	2	4	6954.97
30900	4	3	10739.37	71700	1	4	7164.53
62600	1	3	10743.67	10400	2	4	7547.03
81500	2	3	10770.68	80700	1	4	7853.52
12601	5	3	10802.94	121800	1	4	7944.59
31401	3	3	10940.49	21401	3	4	8235.24
111000	5	3	10941.08	21300	0	4	8254.00
112000	1	3	11006.21	80600	0	4	8277.92
30401	0	3	11048.48	121400	4	4	8336.99
10501	5	3	11127.52	92100	4	4	8499.92
102400	2	3	11159.05	80200	3	4	8626.19
40600	4	3	11272.75	81400	1	4	8790.96
61800	0	3	11440.02	60600	2	4	8859.22
40701	6	3	11733.46	73000	0	4	9023.93
21901	1	3	11772.57	41700	1	4	9096.48

61500	4	4	9178.18	110600	1	5	19036.21
22201	4	4	9370.21	83101	5	5	7758.08
61201	2	4	9444.00	71201	4	6	1949.01
61200	1	4	9550.00	50201	3	6	2259.49
22400	4	4	9613.21	33100	5	6	2489.50
61501	5	4	9773.72	81501	3	6	2586.24
101600	1	4	10075.72	72500	2	6	2914.57
122000	3	4	10303.67	10500	3	6	2940.65
122100	4	4	10319.13	71101	3	6	2984.81
11001	3	4	10438.12	71900	3	6	3090.75
81201	0	4	10657.24	10700	5	6	3129.38
61400	3	4	11009.44	71300	4	6	3161.04
91200	2	4	11214.12	72401	2	6	3168.04
41900	3	4	11266.72	41200	3	6	3230.61
110900	4	4	11616.11	42500	2	6	3283.36
71600	0	4	12040.22	42601	4	6	3391.95
52400	3	4	12056.82	32200	3	6	3473.79
60401	1	4	12072.76	82301	4	6	3498.12
51800	4	4	12748.21	10401	4	6	3515.84
72601	4	4	12783.21	32401	6	6	3632.29
20901	5	4	13115.74	41901	4	6	3679.30
52600	5	4	13144.70	81401	2	6	3693.82
11901	5	4	13288.17	70601	5	6	3784.96
62900	4	4	13580.76	80801	3	6	3906.33
73100	1	4	13788.82	31701	6	6	3908.36
81100	5	4	13829.78	90601	4	6	3927.97
121300	3	4	14070.42	82500	5	6	4028.16
80201	4	4	14906.11	72701	5	6	4031.44
121700	0	4	14950.42	30400	6	6	4032.58
92200	5	4	15262.53	22701	2	6	4089.80
42100	5	4	15603.20	60701	4	6	4197.00
71400	5	4	16602.00	80101	3	6	4217.28
31900	0	4	16623.45	71301	5	6	4221.61
21800	5	4	18109.20	120200	6	6	4320.54
80100	2	4	18144.70	21600	3	6	4348.79
52201	2	5	3047.10	32701	2	6	4355.82
61600	5	5	4342.42	53001	3	6	4364.02
102700	5	5	4855.32	21001	6	6	4697.10
81001	5	5	5063.66	62801	4	6	4756.33
121500	5	5	5313.92	73101	2	6	4788.12
61300	2	5	5491.63	10601	6	6	4793.47
40601	5	5	6543.01	51501	2	6	4864.21
52101	1	5	8780.11	52601	6	6	5008.87
81301	1	5	9470.80	52700	6	6	5017.53
42000	4	5	10777.07	70100	6	6	5021.99
62201	5	5	10849.36	82901	3	6	5044.40
72700	4	5	11932.38	12700	4	6	5083.53
32101	3	5	12790.13	50200	2	6	5092.59
100600	5	5	18549.40	60700	3	6	5100.40

12600	3	6	5123.74	42501	3	6	6306.79
82400	4	6	5164.11	11301	6	6	6342.20
51300	6	6	5182.47	71800	2	6	6362.39
50400	4	6	5218.88	51901	6	6	6384.81
51201	6	6	5245.02	50901	3	6	6409.23
81701	5	6	5263.93	70400	2	6	6430.30
111100	6	6	5271.48	61000	6	6	6471.84
112300	4	6	5347.82	50600	6	6	6500.65
50501	6	6	5403.08	62300	5	6	6521.56
20200	3	6	5440.37	33000	4	6	6531.06
50300	3	6	5469.67	21701	6	6	6533.75
70501	4	6	5488.74	82401	5	6	6556.92
41701	2	6	5594.78	20100	2	6	6577.50
51701	4	6	5656.58	33101	6	6	6591.85
80701	2	6	5660.55	42800	5	6	6608.32
60901	6	6	5679.28	40501	4	6	6633.01
11701	3	6	5690.87	90301	1	6	6636.00
70700	5	6	5712.07	112400	5	6	6647.24
62400	6	6	5721.92	22301	5	6	6649.33
20201	5	6	5734.62	41401	6	6	6665.39
12001	6	6	5746.21	62701	3	6	6680.95
32801	3	6	5779.28	20800	2	6	6694.80
71200	3	6	5781.93	11800	2	6	6709.46
60100	4	6	5786.05	101400	6	6	6749.00
71100	2	6	5815.95	61901	2	6	6850.93
81900	6	6	5865.02	70600	4	6	6888.72
72200	6	6	5885.20	31400	2	6	6928.61
32300	4	6	5892.58	90200	6	6	6934.61
13101	3	6	5914.72	70901	1	6	6953.06
10600	4	6	5945.28	20500	6	6	6955.67
50401	5	6	5979.30	42900	6	6	6964.48
51801	5	6	5985.89	52000	6	6	7004.94
53000	2	6	5990.11	41000	1	6	7022.81
22900	2	6	6006.62	111300	1	6	7184.94
60300	6	6	6007.32	80400	5	6	7189.13
32500	6	6	6019.52	20601	2	6	7220.09
32400	5	6	6038.22	111800	6	6	7356.14
52300	2	6	6041.00	42701	5	6	7549.37
50301	4	6	6067.86	40401	3	6	7590.46
41001	2	6	6084.42	42101	6	6	7674.87
50801	2	6	6145.49	91700	0	6	7754.03
40100	6	6	6147.36	72300	0	6	7856.97
62901	5	6	6157.54	80500	6	6	7918.78
30301	6	6	6219.75	51600	2	6	7988.12
110400	6	6	6225.91	90501	3	6	8015.60
20101	4	6	6230.24	51601	3	6	8100.28
20301	6	6	6261.27	72101	6	6	8242.07
70401	3	6	6305.87	30100	3	6	8244.63
42801	6	6	6306.18	52900	1	6	8245.67

31800	6	6	8251.53
21500	2	6	8264.49
101300	5	6	8398.14
50100	1	6	8442.36
60801	5	6	8524.53
82000	0	6	8547.57
92900	5	6	8563.30
51001	4	6	8603.60
41601	1	6	8613.07
82200	2	6	8622.33
41300	4	6	8837.74
21000	4	6	8873.38
50101	2	6	8939.00
80601	1	6	9056.11
63000	5	6	9131.25
61801	1	6	9183.89
82100	1	6	9319.08
10201	2	6	9521.35
51200	5	6	9565.49
72100	5	6	9645.86
21200	6	6	9646.40
60900	5	6	9686.87
32900	3	6	9740.07
31001	6	6	9742.37
122700	3	6	9751.40
72900	6	6	10082.08
20400	5	6	10176.79
41400	5	6	10480.19
80901	4	6	10758.22
122500	1	6	10803.28
70701	6	6	11190.97
120500	2	6	11274.07
42400	1	6	11401.58
91600	6	6	11445.69
90101	6	6	11537.97
121200	2	6	11635.42
123000	6	6	11738.51
22501	0	7	7035.14
72800	5	7	7779.92
112600	0	7	18251.88
41201	4	7	23519.51