

DOT/FAA/AM-02/2

Office of Aerospace Medicine
Washington, DC 20591

POWER: Objective Activity and Taskload Assessment in En Route Air Traffic Control

Scott H. Mills
SBC Technology Resources, Inc.
Austin, TX 78759

Elaine M. Pfeiderer
Carol A. Manning
Civil Aerospace Medical Institute
Federal Aviation Administration
Oklahoma City, OK 73125

February 2002

Final Report

This document is available to the public
through the National Technical Information
Service, Springfield, VA 22161.



U.S. Department
of Transportation
**Federal Aviation
Administration**

N O T I C E

This document is disseminated under the sponsorship of the U.S. Department of Transportation in the interest of information exchange. The United States Government assumes no liability for the contents thereof.

Technical Report Documentation Page

1. Report No. DOT/FAA/AM-02/2		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle POWER: Objective Activity and Taskload Assessment in En Route Air Traffic Control				5. Report Date February 2002	
				6. Performing Organization Code	
7. Author(s) Mills, S.H. ¹ , Pfleiderer, E.M. ² , and Manning, C.A. ²				8. Performing Organization Report No.	
9. Performing Organization Name and Address ¹ SBC Tech. Resources, Inc. Austin, TX 78759				10. Work Unit No. (TRAIS)	
12. Sponsoring Agency name and Address Office of Aerospace Medicine Federal Aviation Administration 800 Independence Ave., S.W. Washington, DC 20591				13. Type of Report and Period Covered	
				14. Sponsoring Agency Code	
15. Supplemental Notes This report was prepared in support of approved subtask # AM-B-01-HRR-516.					
16. Abstract Two computer programs, the National Airspace System (NAS) Data Management System (NDMS) and the Performance and Objective Workload Evaluation Research (POWER) program, have been developed to provide a platform for quantifying en route air traffic controller activity and taskload. The NDMS program extracts data produced by en route mainframe computers and encodes the information into database files that provide efficient storage and access. The POWER program calculates specific measures using aircraft positions and controller data entries. The development and use of such measures is important for establishing baseline activity measures and for evaluating modifications to ATC systems. NAS System Analysis Recording (SAR) data were collected from the Jacksonville en route air traffic control center between 8:30-10:30 a.m. and between 12:00-2:00 p.m. (local time) for each of four consecutive days. POWER measures were computed in 30-minute intervals for all active sectors. A Principal Components Analysis (PCA) was conducted to evaluate the current set of POWER variables and provide guidelines for the addition of new measures or the modification of existing ones. PCA with Varimax rotation converged in seven iterations and produced five components with eigenvalues > 1. Cumulatively, the four components accounted for 68.18% of the variability in the data set: Component 1 (Activity) accounted for 26%, Component 2 (Flight Path Variability) accounted for nearly 13%, Component 3 (Objective Workload) accounted for 11%, Component 4 (D-side Activity) accounted for 9%, and Component 5 (Overload) accounted for approximately 8%. Variables comprising the five extracted components provided valuable information about the underlying dimensions of the NAS data set. Additions or modifications that might improve the ability of POWER to describe ATC activity and taskload were identified.					
17. Key Words Air Traffic Control, Taskload Measures, Activity Measures, Principal Components Analysis, System Analysis Recordings, Data Analysis Reduction Tool, National Track Analysis Program, NAS Data Management System				18. Distribution Statement Document is available to the public through the National Technical Information Service; Springfield, VA 22161	
19. Security Classif. (of this report) Unclassified		20. Security Classif. (of this page) Unclassified		21. No. of Pages 22	22. Price

POWER: OBJECTIVE ACTIVITY AND TASKLOAD ASSESSMENT IN EN ROUTE AIR TRAFFIC CONTROL

Two computer programs, the National Airspace System (NAS) Data Management System (NDMS) and the Performance and Objective Workload Evaluation Research (POWER) program, have been developed to provide a platform for quantifying en route air traffic controller activity and taskload. The NDMS program extracts data produced by en route mainframe computers and encodes the information into database files that provide efficient storage and access. The POWER program calculates specific measures using aircraft positions and controller data entries. The development and use of such measures is important for establishing baseline activity measures and for evaluating modifications to ATC systems.

En Route Air Traffic Control

In the continental United States, air traffic between terminal areas is controlled by a network of 20 Air Route Traffic Control Centers (ARTCCs). Collectively, these facilities handle over 40 million flights annually. Each ARTCC has responsibility for a portion of airspace that is divided into discrete sectors, with a single controller or team of controllers working traffic in each sector. A typical sector workstation is

equipped with a radar console, a flight progress strip bay, one or more auxiliary text displays, input devices, and a communications panel (see Figure 1).

Each sector workstation is staffed by one to three controllers, referred to as the R-side (radar), D-side (data), and A-side (associate). The R-side controller operates the radar console and communicates directly with aircraft pilots by radio. The D-side controller manages the flight progress strip bay, performs pre-planning duties, and coordinates with other controllers. The A-side controller provides administrative assistance to the R-side and D-side controllers, including delivering flight strips from the printer to the sector workstation.

The radar console consists of a 20"x20" electronic screen (called the Main Display Monitor, MDM) that displays a map of the sector airspace, aircraft position symbols, and aircraft information tags, called data blocks. The data blocks indicate flight information such as sector ownership, aircraft identity, altitude, ground speed, handoff information, and sometimes destination, which is updated as the information changes. Lists of aircraft in potential conflict, departing aircraft, and other information also appear on the




Figure 1. *En route Sector Workstation*

radar screen. The MDM also features an R-side Computer Readout Device (R-CRD) view (or, window) that displays information such as command entries and system messages for the R-side controller. Similar functionality and information is available on the D-side controller's monitor. A monitor is located on the A-side console as well, although its functions are relatively limited. The flight progress strip bay contains paper flight progress strips that display flight information for individual aircraft. Flight progress strips contain more information about each flight, compared with the data blocks, such as equipment type and planned flight route. Controllers use a keyboard and a pointing device (trackball) to enter commands and information, such as assigned altitude levels, into the system.

Changes in Air Traffic Control

The number of flights handled annually by ARTCCs is projected to rise from over 40 million in the year 2000 to more than 76 million by 2025, an increase of more than 90% (FAA, 1999). To accommodate this increase in air traffic, the FAA is updating and modernizing the air traffic control system with the introduction of new automation systems and related operational procedures. Recently, Display System Replacement (DSR) equipment was installed in all ARTCCs. DSR replaced the older, plan view display radar consoles with modern workstations such as the one shown in Figure 1. The DSR system was designed to be modified and enhanced through software upgrades, including decision support tools (DSTs), and several such modifications are planned for implementation. One DST is the User Request Evaluation Tool (URET) which is currently being evaluated in several en route facilities and is expected to be deployed nationwide. URET provides controllers with enhanced conflict alert and resolution functions as well as an electronic aircraft list. Another DST, Problem Analysis Resolution and Ranking (PARR), will provide automated resolution advisories. As new systems and procedures are added to the en route controller's work domain, it will be important to be able to assess the expected benefits and effects of such changes on controller activity and taskload.

History of Workload, Taskload, and Complexity Measurement

Several studies have explored sector activity and taskload in various ways using simulation studies (Buckley, DeBaryshe, Hitchner, & Kohn, 1983; Stein,

1985; Mogford, Murphy, and Guttman, 1994; Pawlak, Brinton, Crouch, and Lancaster, 1996). The methodology most often used assumes that many variables affecting activity in an airspace also influence the perceived workload and the objective taskload of the controller. For example, as the number of aircraft in an airspace increases, one might expect the controller to perceive a higher workload level and perform more activities to maintain safe separation and efficient traffic flow, depending on the difficulty or complexity of the ATC situation. The variables examined in such studies have been described with various terms, including workload (the perceived level of effort required to accomplish a task), taskload, (the amount of activity required to accomplish a task), and complexity (the number or combination of elements influencing workload and taskload).

Buckley, DeBaryshe, Hitchner, and Kohn (1983) conducted a study consisting of a series of experiments in an ATC simulation environment and, as a result, identified a set of four general ATC factors (Conflict, Occupancy, Communications, and Delay) and two auxiliary measures that appeared to adequately represent all other ATC measures (Number of Aircraft Handled and Fuel Consumption). The authors recommended the use of these measures for subsequent air traffic simulation studies. In a separate study, Stein (1985) exposed controllers to different levels of airspace activity and concluded that three variables (i.e., Aircraft Count, Clustering, and Restricted Airspace) significantly influenced mental workload. More recently, Mogford, Murphy, and Guttman (1994) used verbal reports from air traffic control specialists and multidimensional scaling to identify a list of 16 factors that contribute to airspace complexity. Finally, Pawlak, Brinton, Crouch, and Lancaster (1996) focused on controllers' strategies and decision-making activities and proposed a list of 15 factors that may influence perceived air traffic complexity.

Simulation studies such as those cited have many advantages, like the ability to construct and manipulate the air traffic scenarios used in experiments. This allows them to design studies to answer specific research questions. Another advantage of simulation studies is that because the participants are not controlling live air traffic, researchers have the freedom to manipulate conditions and measure different variables without disrupting the controllers' task. For example, situation awareness ratings can be collected by freezing a scenario periodically and administering test instruments to participants. This interruption

would not be possible during actual air traffic control. As a result, the internal validity of conclusions based on simulation studies can be maximized.

While simulation studies are valuable tools for investigating many types of research questions, there are also certain disadvantages that limit their usefulness. For example, while it is desirable to use experienced controllers as participants because they are highly trained in the task of ATC, it can be difficult to obtain sufficient numbers of controllers who are available for participation. Another disadvantage can be the expense of maintaining and operating a highly complex ATC simulation facility. Additionally, the participants' knowledge that they are being observed during an experiment may alter their behavior. Finally, as in all experimental research, there is the problem that external validity may be decreased as a result of the manipulation of experimental conditions. Consequently, decreased external validity may limit the extent to which conclusions based on such studies are generalizable to other settings.

In contrast to simulation studies, taskload and activity measures based on routinely-recorded live air traffic data have certain complimentary advantages. First, participants do not have to be recruited or tested, since measures are based on the actual execution of the task. Second, data recording is a routine procedure of every ARTCC and because data are recorded by the ATC computers, participants are not disturbed as they perform their jobs; that is, data collection is unobtrusive. Finally, because conditions are not manipulated by the researcher, the external validity of these studies is maximized. Therefore, conclusions based on the analysis of routinely-recorded data may be more generalizable to other settings than those based on simulator studies.

The most notable disadvantage of studies based on recorded data is that because the experimenter cannot manipulate conditions, the research is observational rather than experimental. This minimizes the studies' internal validity and limits the ability of researchers to draw conclusions about relationships between variables. Another significant disadvantage is that the set of measures used in such research is limited to those that can be derived from the data routinely recorded by the ATC computers.

Other issues related to analysis of recorded ATC data include the requirement for data storage and organization, and proper computation of relevant taskload measures. Recently, however, two software applications, the NAS Data Management System

(NDMS) and Performance and Objective Workload Evaluation Research (POWER), have been developed to address these issues.

System Analysis Recording

As part of the normal operation of ARTCCs, many types of information are routinely recorded by ATC computers. The computer that performs radar and data processing is the IBM 9672-Generation 3, referred to as the HOST computer. This mainframe system receives and organizes radar and flight information and presents it to the controller at the sector workstation. It also accepts commands and information from the controller through various input devices. As these interactions occur, the HOST records relevant activity in the form of SAR data. This information is stored on magnetic tape and normally retained by the ARTCC for at least 15 days for the purpose of system evaluation, although the tapes may be stored longer if they are needed to review incidents or accidents. SAR reports include information about data displayed at the sector workstations or printed on flight progress strips, controller input to the system, and other flight information.

SAR data are written in Jovial, a binary computer language developed to process ATC information on the HOST computer, which is not easily interpretable by humans. However, two data reduction programs, the Data Analysis and Reduction Tool (DART) and the National Track Analysis Program (NTAP) can be used to produce reports of selected subsets of the SAR data. These programs are run on the HOST computer and produce several types of text-based reports.

The DART program produces the Log report, which contains a variety of system messages. These include controllers' keyboard and trackball entries into the system, as well as all information that was sent by the HOST to the radar display and the auxiliary text display, such as data blocks and list items. The Log report also includes records of all flight progress strip messages. In addition to the Log report, DART also produces the Track report, which contains detailed information from the HOST computer's internal radar track database. This includes data such as each tracked flight's heading, speed, altitude, and position. The NTAP program produces the Beacon and Weather reports. These reports indicate the positions of aircraft beacon and weather symbols on the radar display. A partial listing of the contents of Log and Track reports can be found in Table 1. An example of the type of data available can be seen in Figures 2 and 3.

Table 1. Log and Track Reports – Partial Contents

LOG Report	TRACK Report
Controller Entries	Projected Flight Position
Data Block Contents	Flight Altitude, Heading, and Speed
Auxiliary Text Display Messages	Controlling Sector
Flight Progress Strip Information	Flight Assigned Altitude

```

-----
1 LOG SORT=AID,TIME NAS ID=JX11E015 NAS DATE=09/06/98 CO
! TIME !AID CID/MID!DEV DID !MSG TYP! MSG CONTENT
-----
!112801.7!ASH5504 431!DRC DARC !O FPLAN! FLIGHT MODIFICATION ALT = 100 BCN =
!112801.7!ASH5504 431!HSP HSP !O FPDAT! 1128 TLH AM *431 ASH5504 T/BE02/A 5567 279 TLH P
!112801.8!ASH5504 431!NAS ZCX !I DA ! X976 DA J315
!112801.8!ASH5504 431!NAS ZCX !O AF ! J315 AF ASH5504/431 TLH PNS 06 TLH 09 100 10 TLH..
!112803.9!ASH5504 431!RFD TLHA1!O ACPT! ACCEPT AM ASH5504/431
!112807.5!ASH5504 431!DCR 12D !O RS UP! ASH5504 431 REMOVE * 1*
! ! ! ! ! STRIPS
!112808.8!ASH5504 431!RFD TLHA1!O DEP ! ASH5504 5567 TLH TLH PFN PNS
! ! ! ! ! 1
! ! ! ! ! T/BE02/A P1130
! ! ! ! !
! ! ! ! ! 431 100
-----

```

Figure 2. Log File Excerpt

```

TRACK AIRCRAFT ID = AAL2003 COMPUTER ID = 835 TRACK ID = 065
TIME :MRI TMI ND TC M TMS CN HOI AALT RALT ABC RBC : PREDX PREDY SPD HDG PVX
:ASI PRZ SID RS AHI AHDT BETA DDR DDT : TPF TNF PFX EXS DELSP
CPT : TOC TSB REPA PVB : CLUTR- CONTR DRI
-----
112808.0 :FLT FRE 00 DB P XTE NS FNS 00350 00350 3763 3763: 0523.7 0087.7 00405 326 -0022
:NUL NFZ RAT 16 YES 00000 00000 : 00000 NON 00040
042015.0 : 6297 00358 00000 3763: 0,0,0,0 0,0,0,0 NO
-----
112814.0 :FLT FRE 00 DB P XTE NS FNS 00350 00350 3763 3763: 0523.3 0088.3 00404 327 -0022
:NUL NFZ RAT 16 YES 00000 00000 : 00000 NON 00040
042021.0 : 6298 00423 00000 3763: 0,0,0,0 0,0,0,0 NO
-----
112820.0 :FLT FRE 00 DB P XTE NS FNS 00350 00350 3763 3763: 0523.0 0089.0 00404 327 -0022
-----

```

Figure 3. Track File Excerpt

The reports produced by DART and NTAP can be used to review controller and system activity for a variety of purposes. An example of this process is the method used to review the occurrence of operational errors and incidents with the Systematic Air Traffic Operations Research Initiative (SATORI) program (Rogers & Duke, 1993). SATORI uses information from DART and NTAP reports to graphically recreate ATC incidents on a computer screen. It is used at ARTCCs in an attempt to understand the combination of events that contribute to operational errors and deviations.

Although the data in the LOG and TRACK reports can be used to review air traffic control incidents, the format in which they are created does not provide a practical platform for exploring and computing taskload and activity measures. Unfortunately, DART and NTAP produce very large text-based reports that consist of tables of information designed to be reviewed manually by humans. That is, they are not designed to be efficiently processed by computers. These reports contain large amounts of redundant information (such as formatting characters and table headings) and must be electronically accessed sequentially. This creates difficulties with regard to storage space and processing time.

NAS Data Management System (NDMS)

The NDMS program was developed to provide an optimal platform for ATC activity and taskload research. The program transforms the information in DART and NTAP reports into organized database files that can be accessed rapidly by computer programs. It provides access that allows researchers to investigate the unique characteristics of SAR data and to subsequently develop appropriate methods for calculating measures. An example is the detection of interim altitude assignments, which can be accomplished by scanning specific fields of data block records. Another example is the calculation of handoff latency. This requires searching both data block records and track records. The format of NDMS database files allows computer programs to perform these types of operations quickly and efficiently.

A problem with DART and NTAP reports in their original format is that they require large amounts of electronic storage space. For example, the reports for a single 24-hour period from the Los Angeles ARTCC recorded in January, 2000 require approximately 6.5 gigabytes of computer storage space. This storage issue is rendered inconsequential by using the NDMS system since it reduces the storage space requirements for DART and NTAP output significantly. Once

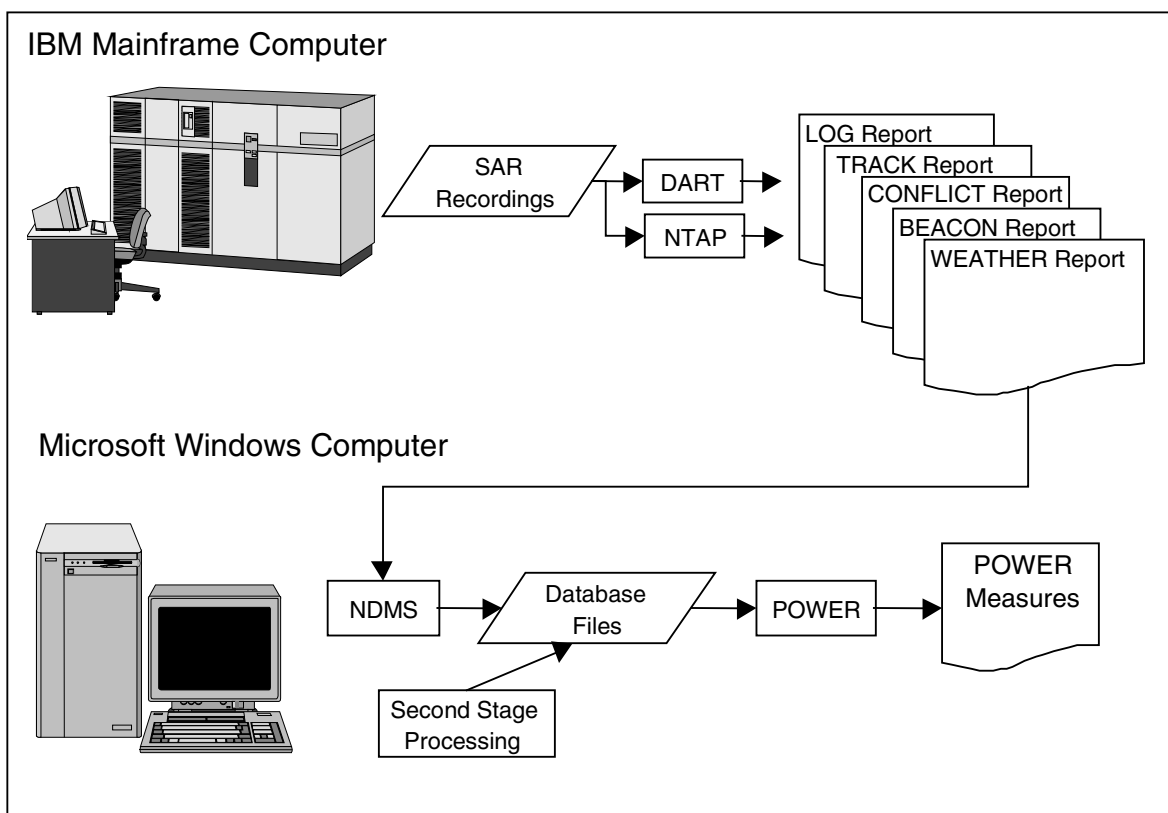


Figure 4. *Data Flow Activity*

Table 2. POWER Measures

Variable	Description
Number of Aircraft Controlled	Count of aircraft controlled during defined time period.
Maximum Number of Aircraft Controlled Simultaneously	Maximum number of aircraft being controlled at one time.
Number of Handoffs (by type — multiple measures)	Number of times control of an aircraft was transferred to or from this sector.
Handoff Latency (by type, e.g., initiate, accept — multiple measures)	Average time from initiation of handoff to acceptance.
Number of Entries (by type — multiple measures)	Count of entries to the computer system made by the controller.
Number of Entry Errors (by type — multiple measures)	Count of controller entries that were rejected by the system because of errors (e.g. acceptance of a handoff for aircraft already under control of the sector).
Pairs of Aircraft in Conflict	Count of pairs of aircraft classified by the computer as being in potential trajectory conflict.
Assigned Altitude Changes	Count of number of times controlled aircraft's assigned altitudes were changed by the controller.
Interim Altitude Changes	Count of number of times temporary altitude assignments were entered into the computer system.
Heading Variation	Average standard deviation of heading changes across flights
Speed Variation	Average standard deviation of speed changes across flights
Altitude Variation	Average standard deviation of transponder reported altitude changes across flights
Control Duration	Average time individual aircraft were controlled by this controller.

converted by NDMS, the output requires only 10 to 15% of the space it required in its original format. NDMS accomplishes this by storing the data in binary database files instead of the original text files that contain mostly formatting characters like spaces and lines.

Because the NDMS system was developed using the Visual Basic programming language, it allows researchers to quickly modify its encoding logic to accommodate the frequent changes that occur in NAS data as the NAS software is updated. In addition, ARTCCs may add unique software “patches” or modifications to the NAS software at their location that allow the programs to adapt to individual characteristics of that particular ARTCC. This results in minor differences in DART and NTAP output between facilities. Nevertheless, NDMS can easily be modified to accommodate these output differences and can, therefore, be used to process data from any en route facility.

Another advantage of the NDMS program is that it allows for second-stage processing of data. To efficiently calculate several activity measures, new data structures that incorporate information from different DART reports are needed. NDMS uses second-stage processing to build these reports after the first-stage encoding has been performed.

To summarize the data flow activity, ATC information is first obtained and collected on SAR recordings (see Figure 4). The SAR tapes are then processed by the HOST computer using the DART and NTAP programs. DART and NTAP produce the LOG, TRACK, CONFLICT, BEACON, and WEATHER reports. These files are then processed by NDMS, which runs on a Microsoft Windows computer. NDMS produces database files which are then processed by the POWER program to produce measures that can be analyzed using statistical analysis packages.

POWER Measures

The POWER application uses NDMS database files as input to produce a set of activity and taskload measures calculated for a specified time period and airspace. The selected airspace may be an individual sector or an entire en route facility. The set of measures produced can be modified or added to by researchers to address specific questions about ATC activity and taskload in specific situations. A list of POWER measures is shown in Table 2. For a complete description of POWER measures and the procedures used to derive them, see the reference provided in Appendix A.

Applications of POWER

POWER will allow for the development of baseline measures of controller activity and taskload for en route ATC. These baselines will be useful for evaluating the effects of changes in equipment and procedures used by controllers. For example, the effects of introducing DSR could be evaluated by comparing pre-DSR baseline POWER measures and post-DSR POWER measures. Additionally, as new enhancements are added to the ATC system with software upgrades, the associated changes in controller activity can be evaluated with POWER measures.

Simulation studies of en route ATC can also benefit from the use of POWER measures as objective indicators of ATC activity and taskload. For instance, POWER could be used during simulation studies to evaluate the effectiveness of proposed software or procedural changes on controller activity or taskload. Although some problems with external validity would certainly exist, the collection of POWER measures is less intrusive than other methods of assessment (e.g., subjective over-the-shoulder ratings, and self-reported workload). Thus, the use of POWER would enhance external validity in the simulated environment.

Finally, POWER can be used as a research tool for developing new ATC metrics that will be needed to further assess the modernization of the ATC system. An example of developing such a metric relates to the concept of dynamic density. Dynamic density is defined as “the projected workload and safety of future traffic in an operational environment” (Radio Technical Commission for Aeronautics [RTCA], 1995). RTCA, Inc. is an organization that addresses requirements and technical concepts for aviation and functions as a Federal Advisory Committee. This organization has emphasized the need for a method of assessing dynamic density so that sectors could be dynamically reconfigured, thereby increasing the capacity and operational efficiency of the NAS. Research with POWER measures could provide valuable knowledge about the relative contributions of different variables to dynamic density.

POWER measures are constantly evolving with these applications in mind. A preliminary analyses is needed to identify further required evolutions of the measures. Principal Components Analysis (PCA) is a statistical technique used to identify and describe complex constructs that may not be directly observable. Components extracted by PCA can contribute to our understanding of a phenomenon by consolidating variables into parsimonious groups. Moreover, the results of such an analysis might provide insight into the development of new measures or the modification of existing ones.

Method

En Route Data From Jacksonville ARTCC

SAR data were collected from the Jacksonville ARTCC between 8:30-10:30 a.m. and between 12:00-2:00 p.m. for each of four consecutive days: 6/8/1998 to 6/11/1998. DART and NTAP reports were produced by a Host computer from the SAR data, and the resulting files were processed by the NDMS program. For the analyses reported here, POWER measures were computed in 30-minute intervals for all active sectors. The number of sectors that were active varied slightly, with an average of 29 active sectors ($SD = 2$) within any given interval. This produced a total of 913 observations for each POWER measure.

The 30-minute interval was chosen for the PCA because it was small enough to produce a sufficient number of observations for the analysis, yet large enough to reduce the risk of ceiling effects in duration measures (i.e., Control Duration and handoff latencies). Control Duration, in particular, is susceptible to these effects because control durations tend to be longer than handoff latencies. When flight durations cross processing intervals, the recorded durations are artificially shortened. This risk increases as processing intervals become shorter. For example, in Figure 5 Flights A and C have 30-minute durations. Flight B has a total duration of 35 minutes. If POWER processing of this sample were conducted using a 15-minute interval, the average Control Duration for

Interval 1 would be 10 minutes; for Interval 2 would be 15 minutes; and for Interval 3 would be 10 minutes. If a 30-minute interval were chosen instead, the average control duration for Interval 4 would be 25 minutes; if a 45-minute interval were chosen, the average control duration for Interval 5 would be 31.7 minutes. Although the computed Control Durations are accurate with respect to the individual intervals, changing the length of the interval will artificially change the value of the resulting measure. Thus, the choice of the appropriate processing interval is important to obtaining meaningful values for the measures.

For these data, the length of the interval was determined by computing Control Duration for all active sectors from one hour of Jacksonville center data (9:00-10:00 a.m. local) using multiple time intervals (i.e., 9:00:00 to 9:04:59 [5 min.], 9:00:00 to 9:09:59 [10 min.], 9:00:00 to 9:14:59 [15 min.], . . . 9:00:00 to 9:59:59 [1 hour]). A single value for Control Duration, averaged over all sectors at the facility, was also computed. Data points for all intervals were then plotted and visually examined. For a few sectors, the relationship was linear (i.e., the durations gradually increased as processing intervals became longer). However, in 22 of the 28 active sectors, the average Control Duration reached asymptote between the 20- and 40-minute processing intervals and the value of facility-wide Control Duration reached asymptote at the 30-minute processing interval. Thus, a 30-minute interval was chosen for all subsequent analyses.

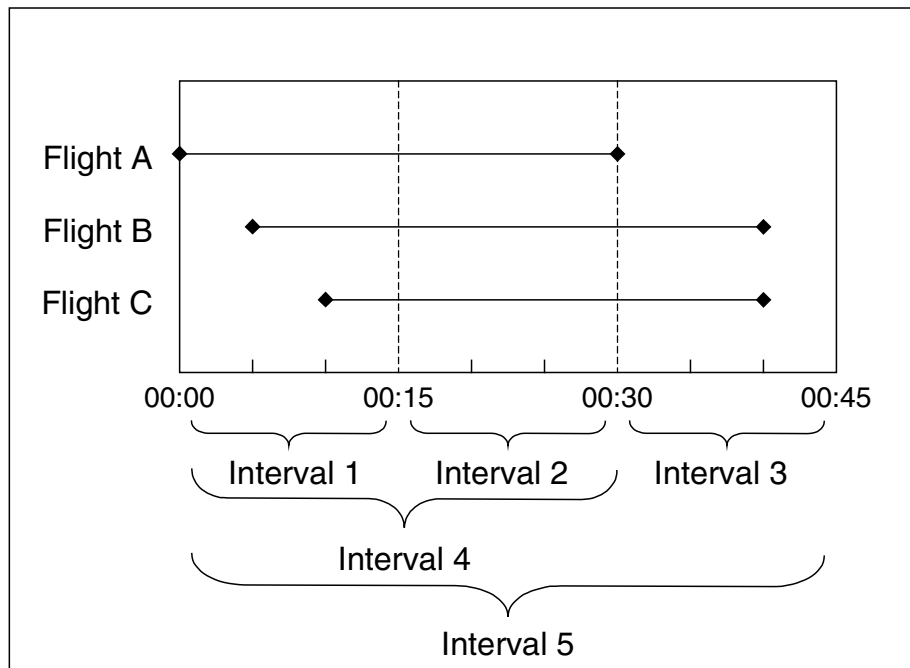


Figure 5. Sample Flight Durations Relative to POWER Intervals

Results

Descriptive Statistics

Several POWER measures computed from this data sample had zero or near-zero incidence. For example, no Conflict Alert List Directives or Immediate Alert Summaries were sent to any of the sectors. Furthermore, none of the Jacksonville controllers made Conflict Alert Suppression entries or Hold requests. Thus, these variables were excluded from further analysis.

Table 3 shows means and standard deviations for the POWER measures that occurred at least once during the data samples. Because specific data entry measures (i.e., Number of Pointouts, Route Display Entries, Track Reroute Entries, Start Track Entries, Data block Offset Entries, and Strip Request Entries) are a subset of general entry counts (i.e., R-side and D-side Entries) they cannot be used in conjunction with the general measures. Therefore, specific entry counts were excluded from the analysis.

After eliminating several variables (as previously described), the following variables were included for further analysis: Number of Aircraft Controlled; Maximum Number of Aircraft Controlled Simultaneously; Counts of Assigned and Interim Altitude Changes, Counts of Handoffs Initiated and Accepted; Latency to accept Handoffs and latency with which Initiated Handoffs are accepted; Control Duration; Pairs of Aircraft in Conflict, Numbers of R-side and D-side Data Entries; Numbers of R-side and D-side Data Entry Errors, and Heading, Speed, and Altitude Variation.

Principal Components Analysis

Prior to the analysis, the Kaiser-Meyer-Olkin (KMO) measure of sampling adequacy was examined to test whether partial correlations among the variables were small (which is desirable in PCA). KMO values of .6 and above are required for a good solution. A KMO of .78 was produced by the set of variables selected. SPSS (10.0.7 for Windows) procedure FACTOR was employed to perform a Principal Components Analysis with Varimax rotation. The rotation converged in seven iterations and produced five components with eigenvalues greater than 1. These components accounted for 68.18% of the variability in the data set. The rotated component matrix is provided in Table 4.

Discussion

The amount of variability accounted for by the five extracted components (slightly more than 68%) was disappointing, but not entirely unanticipated. On the other hand, none of the selected variables failed to load on at least one of the components, and all had a loading of .30 or greater. (Note that in orthogonal rotation, loadings represent the correlations between a variable and a component. Variables with stronger loadings are generally considered to be more representative of a component's underlying processes). Moreover, most of the components were readily interpretable.

Component 1 –Activity

With an eigenvalue of 4.46, Component 1 accounted for about 26% of the variability in the data. The variables comprising this component (shown in Table 4) relate to Activity. The number of Radar controller and Radar Associate (D-side) data entries are straightforward activity measures that relate to the number of commands entered. Handoff Initiates and Accepts, the Number of Aircraft Controlled, the Maximum Number of Aircraft Simultaneously Controlled, and Interim Altitude Changes all relate to aircraft activity in and around the sector. Handoff Accepts involve accepting the transfer of control for an aircraft entering the sector.

The fact that D-side Entries had a loading of only .37 on this component does not necessarily mean D-side Entries are less indicative of activity: Active sectors are always staffed by a radar controller, but not all sectors are worked by a control team. The reduced prevalence of D-side Entries would tend to weaken the association.

Component 2 –Flight Path Variability

Component 2 had an eigenvalue of 2.33 and accounted for about 14% of the variability in the data. The variables comprising this component (shown in Table 4) relate to Flight Path Variability. The component is defined by Average Heading, Speed, and Altitude Variation, although the number of Pairs of Aircraft in Conflict and number of Interim Altitude Changes are also related.

Table 3. Descriptive Statistics for non-zero POWER measures (n = 913)

Variable	Mean	S.D.
* Number of Aircraft Controlled	17.90	6.78
* Maximum Number of Aircraft Controlled Simultaneously	8.27	3.15
* Assigned Altitude Changes	1.75	1.65
* Interim Altitude Changes	10.47	9.41
* Number of Handoff Accepts	11.71	5.05
* Handoff Accept Latency (in seconds)	101.27	60.07
* Number of Handoff Initiates	11.47	5.38
* Handoff Initiate Latency (in seconds)	84.61	48.80
* Control Duration (in seconds)	516.52	158.65
* Pairs of Aircraft in Conflict	1.06	1.08
* Number of R-side Entries	86.29	35.94
* Number of R-side Entry Errors	1.59	1.73
* Number of D-side Entries	5.97	8.25
* Number of D-side Entry Errors	.56	1.39
* Heading Variation (in degrees)	.86	.30
* Speed Variation (in knots)	1.01	.34
* Altitude Variation (in feet/100)	.38	.33
Number of Pointouts	2.40	2.64
Route Display Entries	.40	.97
Track Reroute Entries	.13	.43
Start Track Entries	.18	.59
Data block Offset Entries	.19	.82
Strip Request Entries	.11	.39

* Included in Principal Components Analysis

Table 4. *Principal Components Analysis Rotated* Component Matrix*

Variable	Component**				
	1	2	3	4	5
Number of R-side Controller Entries	.90				
Number of Handoff Initiates	.92				
Number of Handoff Accepts	.91				
Number of Aircraft Controlled	.91				
Interim Altitude Changes	.60	.33	-.43		
Heading Variation		.86			
Speed Variation		.84			
Altitude Variation		.62			
Control Duration			.84		
Maximum Number of A/C Controlled Simultaneously	.68		.58		
Handoff Accept Latency			.57		
Pairs of Aircraft in Conflict		.53			.34
Number of D-side Entry Errors				.91	
Number of D-side Entries	.37			.78	
Number of R-side Entry Errors					.75
Latency to Accept Initiated Handoffs			.39		-.39
Assigned Altitude Changes					.59

* Varimax rotation with Kaiser normalization

** Loadings < .30 not shown

However, there is an inherent difficulty in accurately interpreting this factor. As can be seen in Table 3, the three variables – Heading, Speed, and Altitude Variation – have limited variability. All three measures represent the average standard deviation of changes across flights. It is doubtful that, in their present form, they are sufficient to describe aircraft movements because the distribution of standard deviations is calculated from incremental differences, rather than actual “changes.” For instance, as an aircraft begins to change its speed the difference is recorded as 3 knots (from 280 to 283). By the next update, the aircraft might have reached 287 knots (a difference of 4 knots). At the next update, the change is complete and the aircraft levels off at 290 knots (a difference of 3 knots). The actual speed change was 10 knots, but this would not be reflected by computing the standard deviation of the differences. In the effort to measure heading, speed, and altitude changes, it was exceedingly difficult to establish parameters that eliminated error variance (i.e., natural deviations in real data) and still retain the actual changes. The current variables were computed in an attempt to circumvent such difficulties. Unfortunately, this method also hides pertinent information within the error variance and produces measures that are of limited usefulness. Thus, while Component 2 describes an underlying communality between three variables that may describe some aspect of aircraft movement, it may also only reflect a similarity in computational methods. The lower loadings of the other two variables with this component could indicate some other aspect of changes in aircraft flight paths, or it may reflect the non-normality of their corresponding distributions.

Component 3 – Objective Workload

After rotation, Component 3 had an eigenvalue of 1.92 and accounted for about 11% of the variability in the data. The variables comprising this component (listed in Table 4) reflect objective workload in that they represent the controllers’ reactions to the events to which they were exposed. Generally, workload not only refers to controllers’ reactions to events, but also to their perception of the effort involved in managing those events. However, we cannot make inferences about the subjective experience of these controllers based on the available data. Therefore, this component has been given the interpretive label *Objective Workload*.

The “marker” variable for Component 3 is Control Duration, which represents the average amount of time aircraft are under a sector’s control. This variable

relates to workload because the longer an aircraft is in the sector the longer the controller must attend to it. Maximum Number of Aircraft Controlled Simultaneously is indicative of workload as well, since the more aircraft controlled simultaneously, the more often the controller must assess potential conflicts and other problems. Handoff Accept Latency may also be indicative of workload since it takes longer to accept handoffs from another sector when a controller is busy. Likewise, Interim Altitude Changes are generally avoided when the controller is busy because of the amount of data entry required to perform them, hence the negative relationship of this variable with others comprising this component. Finally, Latency to Accept Initiated Handoffs reflects workload because the sector controller must attend to aircraft in handoff status until he/she is certain the handoff has been accepted.

One of the more interesting aspects of Component 3 has to do with the combination of Control Duration and Maximum Number of Aircraft Controlled Simultaneously. Together, they constitute a gross measure of traffic density. These variables roughly correspond to average sector flight time and peak traffic count, which are used to compute density at en route centers (see FAA, 1984, Appendix 1) to estimate required staffing standards. The fact that elements of Component 3 might be an indirect measure of density suggests that a more direct measure might improve the set of POWER variables. The density formula (FAA, 1984) was developed because it was considered impractical to manually compute the average number of aircraft controlled each minute, but this would be extremely simple to calculate with a minor revision to the POWER processing code. However, the average number of aircraft under the sector’s control does not convey any information about the proximity of the aircraft. If traffic characteristics are a contributor to workload (as proposed by Mogford et al., 1994 and Pawlak et al., 1996) then it might also be advantageous to include one or more proximity measures in the POWER suite. The information necessary to compute such measures is readily available and so the addition is feasible.

Component 4 – D-side Activity

Component 4 had an eigenvalue of 1.56 and accounted for about 9% of the variability in the data set. The variables included in the component may be considered to describe D-side activities, and thus, Component 3 was labeled *D-side Activity*. The variables included in this component (listed in Table 4) are D-side data entries and errors. It must be remembered when interpreting this component that sectors

are often staffed by only one person instead of a team of controllers. When an R-side controller is working alone, he or she must move to the D position to make certain data entries that cannot be entered on the R-side workstation. The only way to determine whether a sector is being worked by more than one person is to examine sector staffing records, which are recorded in electronic SISO (Sign In Sign Out) logs. Without this information it is impossible to accurately identify which sectors are actually being worked by D-side controllers. Unfortunately, SISO data were not available for the Jacksonville Center data set. However, large SAR data sets from the Kansas City (ZKC), Los Angeles (ZLA), and Washington (ZDC) centers are currently being processed, and corresponding SISO data are also available. With this additional information it will be possible to investigate the stability, correlates, and implications of a *D-side activity* component.

Component 5 – Overload

Component 5 had an eigenvalue of 1.33 and accounted for about 8% of the variability in the data set after rotation. This component has been tentatively labeled *Overload* because, as shown in Table 4, the variable with the highest loading was R-side Entry Errors (.75). Assigned Altitude Changes might also be indicative of overload because workload can be higher in transition sectors where altitude changes are made more often. The time measured by the Latency to Accept Initiated Handoffs variable is the time it takes another controller to accept a handoff initiated in the current sector. Perhaps the increased workload required to attend to whether another controller has accepted a handoff contributes to overload as well. Finally, Pairs of Aircraft in Conflict may be indicative of overload because attending to more conflict alert notifications (which are often not an indication of a real conflict) requires time that might be better spent on other activities. Whereas the variables that loaded on the component seem to suggest overload, most of the loadings were small and so interpretation of this component is somewhat ambiguous.

Conclusions

POWER measures were developed to provide a platform for quantifying en route air traffic controller activity and taskload. The development and use of such measures is important for establishing baseline activity measures and for evaluating the effects of modifications to ATC systems. Success depends on the selection of variables that are, in combination, sufficient to comprehensively describe the ATC environment.

The value of conducting the PCA using these data, albeit restricted, is that the five components extracted suggested possible additions or modifications that might improve the ability of the POWER measures to describe air traffic controller activity and taskload. For example, the lack of a relationship between the aircraft dynamics measures and other measures of controller and aircraft activity suggests that Average Heading, Speed, and Altitude Variation may not measure what they were intended to measure (i.e., they currently represent the standard deviation of incremental differences rather than actual changes). Because of the results of the PCA, it is apparent that additional measures of aircraft dynamics (e.g., counts and amounts of actual changes, duration of changes, etc.) may be more effective measures of variability in aircraft movements.

The results of the PCA contribute to our understanding of the POWER measures in other ways. For example, the pattern of variable loadings on Component 3 (Objective Workload) suggested that there might be an element related to aircraft density or proximity; traffic characteristics that are not being measured by the current set of POWER variables. Presently, a new measure of proximity is being developed and will be added to the POWER suite. The amount of variability explained by the POWER measures before and after inclusion of the new variable will be tested in the upcoming baseline (pre-DSR) study.

The existence of Component 4 (D-side Activity) and aspects of Component 1 (Activity) also made it clear that, in the future, separate analyses should be conducted for sectors staffed by individual controllers and those staffed with control teams. Because corresponding SISO data have also been collected for the three centers involved in the baseline study, we will be able at that time to conduct separate analyses for sectors staffed by individual controllers and those with control teams.

In future POWER research we will continue to examine the combination of variables that make up the POWER measures and seek ways to improve their ability to describe workload and taskload. In addition, we will examine information about geographic and traffic characteristics of sectors in different facilities and compare patterns of POWER measures in similar and dissimilar sectors. We eventually plan to conduct additional validation research in a simulated environment with the goal of further examining the relationship between POWER measures and subjective measures of workload.

References

- Buckley, E.P., DeBaryshe, B.D., Hitchner, N., & Kohn, P. (1983). *Methods and measurements in real-time air traffic control system simulation*. (Report No. DOT/FAA/CT-83/26). Atlantic City, NJ: Federal Aviation Administration Technical Center.
- Federal Aviation Administration (1984). *Establishment and validation of en route sectors*. (FAA Order No. 7210.46). Washington, DC: Author.
- Federal Aviation Administration (1997). *National Airspace System (NAS) Architecture, Version 2.5*. Office of System Architecture and Program Evaluation (ASD-1). Washington, DC: Author.
- Federal Aviation Administration (1999). FAA Long-Range Aerospace Forecast. [HTTP://api.hq.faa.gov/default.htm](http://api.hq.faa.gov/default.htm).
- Galushka, J., Frederick, J., Mogford, R., & Krois, P. (1995). *Plan view display baseline research report*. (Report No. DOT/FAA/CT-TN95/45). Atlantic City, NJ: Federal Administration Technical Center.
- Mogford, R.H., Murphy, E.D., & Guttman, J.A. (1994). Using knowledge exploration tools to study airspace complexity in air traffic control. *The International Journal of Aviation Psychology*, 4(1), 29-45.
- Pawlak, W.S., Bowles, A., Goel, V., & Brinton, C.R. (1997). *Initial evaluation of the Dynamic Resectorization and Route Coordination (DIRECT) System*. NASA Final Contract Report No. NAS2-97057.
- Pawlak, W.S., Brinton, C.R., Crouch, K., & Lancaster, M. (1996). *A framework for the evaluation of air traffic control complexity*. Proceedings of the AIAA Guidance Navigation and Control Conference, San Diego, CA.
- Rodgers, M.D., & Duke, D.A. (1993). SATORI: Situation Assessment Through Re-creation of Incidents. *The Journal of Air Traffic Control*, 35(4), 10-14.
- RTCA (1995). *Final Report of RTCA Task Force 3: Free Flight Implementation*. Washington, DC: RTCA Incorporated.
- Stein, E.S. (1985). *Air traffic controller workload: An examination of workload probe*. (Report No. DOT/FAA/CT-TN84/24). Atlantic City, NJ: Federal Aviation Administration Technical Center.
- Tabachnick, B.G. & Fidell, L.S. (1989). *Using multivariate statistics* (2nd ed.). New York: Harper Collins.

APPENDIX A

Performance and Objective Workload Evaluation Research (POWER) Reference

The following is a description of the POWER software system and the procedures used to derive ATC activity measures from FAA System Analysis Recording (SAR) data. The first step in the process is a reduction of the SAR data into reports generated by the Data Analysis Reduction Tool (DART) and National Track Analysis Program (NTAP). These reports can only be generated by a National Airspace System (NAS) Host computer. Due to the size and format of the DART and NTAP text files, it would be impractical to use these reports in their raw form. Therefore, the NAS Data Management System (NDMS) was developed to organize this information into Microsoft Access databases. Preliminary organization of the raw data provides a dual advantage: It decreases the size of the data to be stored, which in turn increases the speed of POWER processing. Although POWER measures are computed exclusively from DART reports, NTAP reports are also processed and stored by NDMS.

NDMS organizes the messages from DART and NTAP reports into hour-long databases. These databases consist of content-specific tables. Beacon and Weather data from the NTAP reports are stored in one-minute tables (e.g., BCN_01 contains data from minutes 00 through 01, WTH_01 contains weather information from minutes 0 through 01, etc.). LOG input and output messages are parsed by message type (e.g., LOG_O_FPL contains flight plan messages output by the system). Track data are organized by individual flights, distinguished by both AID and CID (e.g., AAL1234_5678). Message fields are parsed and stored within separate columns in the tables. This format facilitates computer processing of the POWER measures.

Once NDMS processing is completed, the Certify program is run on the hour-long data files. The Certify program inserts an aircraft type reference table (TYPEREF) into each database that lists aircraft type and equipment information for all AIDs, derived from flight progress strip messages. (If no flight progress strip messages are available, the AID is entered into the database, but the type

and equipment fields remain blank.) The Certify program then compares the aircraft type designation with a resource database that provides additional data (i.e., manufacturer, average climb and descent rate, etc.) and writes this information to a table (TACTYPE) in the hour-long database. The Certify program also compresses the data files to reduce storage space requirements.

The next step is the Daytrack program. As its name suggests, the Daytrack program compiles hourly Track information for each aircraft into “day-long” TK tables. These tables are written to a separate file. Hourly files are labeled in a ddmmyyhh format with an extension that corresponds to the three-letter facility identifier. For example, a file labeled 12289813.zkc contains data from 13:00:00 to 13:59:59 (ZULU) recorded on 12/28/98 at the Kansas City en route facility. The file created by the Daytrack program that contains all available hour-long data for the day would be labeled 122898DR.zkc. The “DR” of the Daytrack files is a non-numeric two-character identifier that makes the day-long databases easily distinguishable from the numeric hour-long ones.

POWER Measures

Number of Aircraft Controlled

This value represents the total number of aircraft controlled by any given sector or facility during a specified POWER interval. Controlling sector information is derived from the TRACK file produced by the DART report. POWER compiles a temporary list of controlled aircraft for any given interval using the CN (controlling sector) and DGTIM (digital time) fields from the TRK tables (i.e., tables that contain Track data for each flight). The list is used to calculate the total number of controlled aircraft. Aircraft do not have to be controlled by the sector for the entire interval. Any aircraft that is controlled by the sector at any time within the interval is included. Table A1 contains a sample list of aircraft for the POWER interval 12:10:00 to 12:29:59 in which the number of aircraft controlled by Sector 16 equals five.

Table A1. Sample List of Controlled Aircraft

AID	CN	START	STOP
DAL422	16	12:10:08	12:11:50
EJA157	16	12:15:14	12:27:56
AAL1661	16	12:15:26	12:30:56
AWE726	16	12:18:38	12:20:38
AAL61	16	12:21:08	12:25:14

Maximum Number of Aircraft Controlled Simultaneously

This value represents the maximum number of aircraft under simultaneous control within a specified POWER interval. The same list used to calculate the total number of controlled aircraft is used to calculate the maximum number of aircraft under simultaneous control. POWER checks the number of aircraft under a sector's control for each minute of a given POWER interval. Using the list of aircraft in Table A1, the maximum number of aircraft under simultaneous control in Sector 16 during the POWER interval 12:00:00 to 12:29:59 equals three. This was calculated by first checking the number of aircraft controlled from 12:00:00 to 12:00:59. As no aircraft in the list were controlled by the sector, a value of 0 was retained for comparison with the number of aircraft that were controlled from 12:01:00 to 12:01:59, and so on. From 12:10:00 to 12:10:59, DAL422 was controlled by the sector. Therefore, the stored value would be replaced with 1. From 12:15:00 to 12:15:59, there were two aircraft controlled by the sector (i.e., EJA157 and AAL1661). This value was greater than the stored value of 1, and so the stored value was replaced. From 12:18:00 to 12:18:59, there were three aircraft controlled by the sector (i.e., EJA157, AAL1661, and AWE726) and the stored value was again replaced. Because AWE726 was no longer under the sector's control by the time the sector assumed control of AAL61, the stored value remained unaltered. The maximum number of aircraft controlled is equal to the stored value that remains after all minutes of the interval have been evaluated.

Control Duration

This value represents the average time aircraft were controlled (in seconds) within a specified POWER interval. At the time POWER stores a temporary list of AIDs for a given sector, it also

stores the time at which the sector took control of the aircraft and the last recorded time before the aircraft left the sector's control. For example, in Table A1 Sector 16 assumed control of DAL422 at 12:10:08 and maintained control until 12:11:50. The control duration for DAL422 would be 102 seconds. The same calculation is made for all aircraft controlled by Sector 16 during the interval. Control duration for any given POWER interval is equal to the mean of the durations of all aircraft controlled by the sector within the POWER interval. Control time occurring before or after the POWER interval is not included in the calculations.

Heading Variation

This value represents the average standard deviation of heading changes across flights. For each flight, POWER calculates heading differences and stores all changes that do not exceed a specified value (The default value is 12°, but this value may be set manually prior to the POWER run) into an array (or, temporary list). POWER then calculates the standard deviation of the distribution of differences and sends this information to a second array. When standard deviations have been collected for all flights, POWER computes the mean of the distribution.

Speed Variation

This value represents the average standard deviation of speed changes across flights. For each flight, POWER calculates differences in speed and stores all changes that do not exceed a specified value (default value is 30 knots, but this value may be set manually prior to the POWER run) into a temporary array. POWER then calculates the standard deviation of the distribution of differences and sends that information to a second array. When standard deviations have been collected for all flights, POWER computes the mean of the distribution.

Altitude Variation

This value represents the average standard deviation of transponder reported altitude changes across flights. For each flight, POWER calculates differences in altitude and stores all changes that do not exceed a specified value (default value is 10,000 feet, but this value may be set manually prior to the POWER run) into a temporary array. POWER then

calculates the standard deviation of the distribution of differences and sends that information to a second array. When the standard deviations have been collected for all flights, POWER computes the mean of the distribution.

Handoff Count (total)

This value represents the total number of handoffs occurring within this sector/facility. A handoff is defined as a change in the CN (controlling sector) field of the TRK tables in which at least one of the involved sectors (either the initiating sector or the accepting sector) was this sector/facility.

Handoff Count (valid)

This value represents the total number of handoffs for which a corresponding initiate message can be located. At the time a change in the CN field is detected, POWER stores the information in an array. Table A2 contains a sample of elements in this array. These include the aircraft identifier of the aircraft being handed off (AID), the initiating sector from which control was assumed (CN From), the sector that assumed control (CN To), the last recorded time the aircraft was under the (CN From) sector's control, and the time the change was recorded (Time of Change). Once this information is collected for each item, POWER searches information recorded in DART log files for an initiate/accept message corresponding to the change (Time of Initiate/Accept). If no initiate message is found, the handoff is excluded from the count.

Handoff Latency

This value represents the average time between initiation and acceptance of valid handoffs, regardless of whether the aircraft was entering or leaving the sector. POWER determines initiate and accept times from information displayed in Field E of the data block tag. In some cases, Field E indicates that the data block was "busy" at the time the handoff was accepted and the earliest Field E accept time is less accurate than the time of change noted in the CN field of the TRACK tables. In these instances the time of change in the CN field is substituted and used for the latency calculation.

Handoff Accept Count (for sector)

This value represents the sum of the number of valid handoffs accepted by this sector.

Handoff Accept Latency (for sector)

This value represents the mean latencies between the time a handoff is initiated by a previous sector and the time the handoff is accepted by this sector.

Handoff Initiate Count (for sector)

This value represents the sum of all valid handoffs initiated by this sector.

Handoff Initiate Latency (for sector)

This value represents the mean latencies between the time a handoff is initiated by this sector and the time a handoff is accepted by the next sector.

Table A2. *Sample of Data Used to Determine Valid Handoffs and Handoff Latencies*

AID	CN From - CN To	Time of Change	Time of Initiate/Accept
AAL1661	NS - 16	12:15:20 - 12:15:26	12:13:53 - 12:15:26
AAL61	15 - 16	12:21:02 - 12:21:08	12:20:29 - 12:21:08
AAL61	16 - 17	12:25:14 - 12:25:20	12:22:50 - 12:25:20
AWE726	15 - 16	12:18:32 - 12:18:38	12:17:27 - 12:18:33
AWE726	16 - 17	12:20:38 - 12:20:44	12:19:55 - 12:20:44
DAL422	16 - 17	12:11:50 - 12:11:56	12:10:02 - 12:11:56
EJA157	17 - 16	12:15:08 - 12:15:14	12:14:11 - 12:15:14
EJA157	16 - NS	12:27:56 - 12:28:02	12:26:08 - 12:28:02

Conflict Alert Pairs

Conflict alert pairs are calculated from output messages in the LOG files of the DART reports (message type LOG_O_CA) that record all conflict alert pairs sent to the High Speed Printer (HSP). Each unique conflict alert pair is counted once during the analysis epoch. For individual position processing, all pairs with at least one aircraft currently under the sector's control are counted. For entire facility processing, POWER counts all pairs with at least one aircraft controlled by a sector within the facility. Conflict alert pairs with both aircraft controlled by a sector (or sectors within the facility) are counted only once.

Conflict Alert – Zero

Conflict Alert – One

Conflict Alert – Two

These values are derived from Conflict Alert List Directive (LOG_O_CALDR) output messages recorded in the DART log files. When the initial conflict alert directive is identified (determined by the DISPLAY value recorded in the message content), POWER stores the information in an array. Elements of this array include: the time the message was transmitted, AIDs of the aircraft involved, controlling sectors for both aircraft, and the device to which the information was sent. Unique messages meeting criteria for Conflict Alert Zero, One, and Two are calculated from items in the array. The criteria for these variables are as follows:

Conflict Alert – Zero

The number of conflict alert messages sent to this sector in which no aircraft were controlled by this sector. Comparable facility counts include messages sent to any sector in which neither aircraft were controlled by a sector within the facility.

Conflict Alert – One

The number of conflict alert messages sent to this sector in which one aircraft was controlled by this sector. Comparable facility counts include messages sent to any sector in which only one aircraft was controlled by a sector within the facility.

Conflict Alert – Two

The number of conflict alert messages sent to this sector in which both aircraft were controlled by this sector. Comparable facility counts include messages sent to any sector in which both aircraft were controlled by sectors within the facility.

Conflict Alert – Total

This value represents the sum of all conflict alert messages (i.e., Conflict Alert – Zero, Conflict Alert – One, and Conflict Alert – Two.)

Conflict Alert – Suppresses

This value represents the total number of conflict alert blink suppressions initiated by this sector (or, sectors within this facility). Conflict alert suppresses are derived from the message content of Conflict Alert List Directive (LOG_O_CALDR) output messages recorded in the DART log report. Information from the messages are stored in an array that contains the AIDs and unique Conflict Alert identification number (CAID) of the conflict pair. Suppresses are correlated with the sector/facility at which they were initiated: Controlling sectors of the conflict pair are not evaluated.

Conflict Alert – Immediate Alerts

This value represents the number of Immediate Alerts in which at least one aircraft was controlled by this sector/facility. POWER calculates this value by counting all applicable Immediate Alert Summary output messages (LOG_O_IAS) in the DART log files.

Assigned Altitude Changes

This value represents the number of assigned altitude changes for this sector/facility. Assigned altitudes for all controlled aircraft are recorded in DART track data files. For any given interval, POWER stores a temporary array of these values. POWER then searches the array for changes in altitude and tabulates the total.

Interim Assigned Altitude Changes

This value represents the number of interim altitude changes entered into ATC system for this sector/facility. Interim altitude changes are identified by a "T" displayed in character position B4 in

Field B of the data block tag of a controller's radarscope. This information is recorded in Full Data Block (FDB) messages in the DART log reports. POWER collects and sorts Field B4 information for the pertinent sector(s) within a given processing interval. Each unique interim altitude is extracted and tabulated.

Controller Entries

These values represent frequency counts of Computer Readout Device (CRD) entries for this sector/facility, sorted by R-side, D-side, and A-side. Controller CRD entries are recorded as input messages in DART log reports. It is important to note that this value represents the number of entry messages and does not necessarily reflect the exact number of actual keystrokes that might be required to generate the message. POWER classifies messages as R-, D-, or A-side according to the device. Therefore, these entries do not necessarily identify the duties of the individuals making the entries.

Entry Errors

These values represent frequency counts of the number of entry errors appearing on the CRD for this sector/facility, sorted by R-side, D-side, and A-side. All errors displayed on the CRD are recorded as output messages in DART log reports (LOG_O_ERROR and LOG_O_REJECT). POWER classifies errors as R-, D-, or A-side errors by the device that displayed the error message.

Distance Reference Indicator Requests

This value represents the number of distance reference indicator (DRI) request entries for this sector/facility. POWER computes this value from the content of DRI output messages recorded in DART log reports (LOG_O_DRIDO). The message content indicates whether the entry was a request or a delete.

Distance Reference Indicator Deletes

This value represents the number of distance reference indicator (DRI) delete entries for this sector/facility. POWER computes this value from the content of DRI output messages recorded in DART log reports (LOG_O_DRIDO). The message content indicates whether the entry was a request or a delete.

Route Display Entries

This value represents the number of route display entries for this sector/facility. POWER computes this value from the content of output accept messages recorded in DART log reports.

Pointout Entries

This value represents the number of pointout entries for this sector/facility. POWER computes this value from the content of output accept messages recorded in DART log reports.

Pointout Entries (breakdown)

These values represent the number of pointout entries for this sector/facility, sorted by R-side, D-side, and A-side. POWER classifies pointouts as R-, D-, or A-side entries by the device used to make the entry.

Data Block Offset Entries

This value represents the number of data block offsets for this sector/facility. POWER computes this value from the content of output accept messages recorded in DART log reports.

Track Reroute Entries

This value represents the number of track reroute entries for this sector/facility. POWER computes this value from the content of output accept messages recorded in DART log reports.

Start Track Entries

This value represents the number of start track entries for this sector/facility. POWER computes this value from the content of output accept messages recorded in DART log reports.

Hold Entries

This value represents the number of hold entries made by this sector/facility. POWER computes this value from the content of output accept messages recorded in DART log reports.

Strip Request Entries

This value represents the number of strip requests made by this sector/facility. POWER computes this value from the content of output accept messages recorded in DART log reports.

