Collection and Analysis of Quality Data in a Distributed Simulation Test Environment

prepared by
Dean G Gonzalez and Jerry Black

1. Introduction

Since T&E using ADS is relatively new and continues to evolve, there are a paucity of tools available for general use that allow new T&E ADS personnel to evaluate the quality of their ADS system The quality focus of the tools presented is upon tools to assist in the detection of time errors, network-induced errors, and PDU-generation errors. Tools were developed that proved useful. These will be presented.

2. Background

The purpose of JADS testing is to assess the utility of ADS for T&E. This purpose is being accomplish through the execution of three distintly separate tests. The first of these is called the System Integration Test (SIT). The first phase of this test is the Linked Simulators Phase (LSP). It is this first phase of the first test that has recently been completed and which is the basis for the tools presented.

The SIT LSP tests simulate a shooter aircraft launching an air-to-air missile against a target aircraft. The shooter, target, and missile are represented by geographically separated simulators. The shooter is represented by the F/A-18 Weapon System Support Facility (WSSF) at China Lake, CA. The missile is the AIM-9 Sidewinder Simulation Laboratory (SIMLAB), also at China Lake, CA. The target is represented by the F-14 Weapons System Integration Center (WSIC) at Point Mugu, CA. Test control of this distributed test will be done from the Test Control and Analysis Center (TCAC) located at the JADS JTF in Albuquerque, NM.

The SIT LSP test replicates a "baseline" live fire test. The test evaluation method is to compare ADS test results with results from the identical "baseline" test (This is a simplified view of the JADS testing ... but it suffices for this discussion).

Data collection and storage, the data analysis system, and time synchronization topics are reviewed as preparatory material for the presentation of the tools that have been developed.

3. Data Collection and Storage

A simplified architecture of the ADS system for the LSP tests is shown below to assist with the discussion that follows:

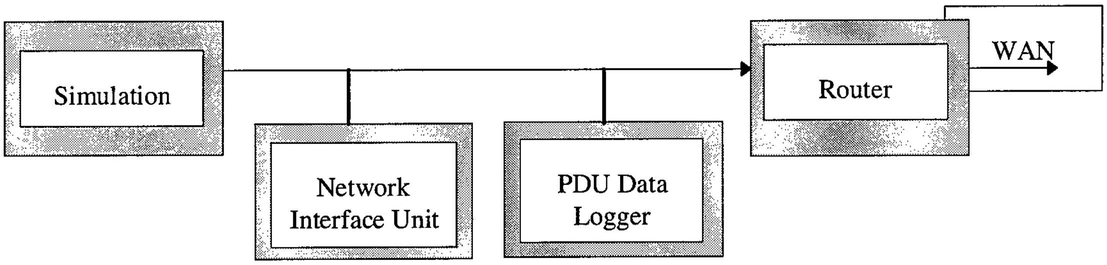

LSP Tests Network Architecture

3.1 Data Loggers

JADS utilized two primary data logging systems to record the ADS PDUs. Data files were also obtained from the participating simulators. These were termed the "high fidelity data". Part of the analysis assessed the effects of the ADS system upon the validity of the high fidelity data.

The STRICOM logger was used because it was low cost, i.e. it was actually no cost, since it was GFE. Also, it could be used on various UNIX platforms. Additionally, it provided accurate time stamping of logged PDU data.

The SNAP (Simulation Network Analyst Project) logger was obtained from Wright Patterson AFB. It was developed specifically for the capture and precision time-stamping of network messages. It was used to verify the time-stamping precision of the STRICOM logger as well as to collect network messages in general for network traffic analyses.

The high-fidelity data, obtained from each of the simulators, are the output of the simulators prior to input into the network interface units that transformed these data into PDUs.

3.2 Data Logging Setup

A Silicon Graphics workstation (SGI Indy) was used at each node with the STRICOM logger software to log the PDUs just after they were created by the Network Interface Units for the local simulation, and after they were received by the node from remote entities.

The nodes at which loggers were set up include the Weapon System Simulation Facility (WSSF) at China Lake, the Simulation Laboratory (SIMLAB) also at China Lake, and the Weapon Systems Integration Center (WSIC) at Pt. Mugu.

A STRICOM logger was also set up at the Test Control and Analysis Center (TCAC) at Kirtland.

Each Logger logged PDUs for all entities.

At the end of a day of testing, logger files and high fidelity data files were compressed (using the UNIX TAR and COMPRESS utilities), and were electronically sent (using the UNIX File Transfer Protocol (FTP) to the TCAC at Kirtland for storage and analysis.

3.3 Data Log Architecture

All source data files were stored in a common data directory in a file server at the TCAC.

The data directory contains a subdirectory for each mission day, as well as a subdirectory within each mission day for each logger location.

Each file was named with the month-day-year of the mission, the test number, and the location. The format was mmddyy_testnn_loc.xxx, where the xxx filename extension identified the type of data, i.e.

1. .lgr files were STRICOM logger data.
2. .dat files were Simulation (high fidelity) data files.
3. .snp files were snap data files.

4. Data Analysis System

The foundation of the data analysis system is the Metrica data management and analysis toolset. This system and its application to the JADS JTF data management and analysis is described in the following section.

4.1 Data Management System

The foundation for the Metrica data management and analysis system is a relational database. The requirement to structure collected data into database tables is an advantage as well as a disadvantage. The major advantage is the structure that is imposed on the datasets. The disadvantage is the need to develop this structure so that the data are stored in a useful and meaningful manner for all analysts.

One of the strong features of Metrica is its ability to contain array data as a field in its database tables. Thus, the JADS logger files database contains one record for each mission-day, trial-number, logger-location, logger-machine, and entity. Each record contains array fields for the logged time, PDU time, and the selected PDU data that include items such as latitude, longitude, altitude, etc. When it is desired to plot data for a selected record, the server does not need to access hundreds or thousands of records for each value of time and data. Since the data are stored as arrays, only one pointer is returned for each of the data to be plotted, and the server workload is minimal. The result is a very fast retrieval of data, even very large datasets, for plotting, tabulating, or other manipulation using Metrica-provided functions, or user-defined functions.

Metrica provides a versatile scripting language that facilitates the development of customized tools, including menus, etc.. A host of pre-defined plot types, such as 3D plots, histograms, etc. is also provided. Built-in array manipulation functions facilitate handling of large data sets. Also, Metrica provides a generous set of statistical functions.

The other major tool used for analysis is Microsoft Excel. This is a well known tool. It provides a wealth of analysis capability. The major reason it wasn't selected for all of the analyses is its limitations in handling large data sets.

For some analyses, C programs were developed due to the speed and versatility of the language. These were then called by the Metrica menus to accomplish the specific analyses.

4.2 Data Analysis Structure

Metrica uses a relational database as the foundation for its data management and analysis toolset. It is the nature of database systems that protection of the database is a paramount concern. This protection can cause nightmares and headaches for the database administrator, especially as the database is populated and software problems arise. The JADS analysis system design focuses upon preservation of the source data and upon insuring its integrity. The Metrica database itself, if destroyed, can be recreated with a minimum of effort using the source data and scripts that have been developed to recreate the table structures and import the data into Metrica.

Thus the database data and all of its structures can be destroyed with impunity. Database tables are easily recreated with scripts. A custom C program loads data from logger files into Metrica quickly. For example, loading of a complete day of LSP logger data takes less than five (5) minutes (The average day consists of about 20 trials, 4 loggers, 3 entities and constitutes about 10 megabytes of data).

5. Time Synchronization

The foundation for much of the JADS analyses, i.e. looking at data latencies between the various nodes, is accurate time synchronization over the wide area network's simulation machines and data logging devices. The following paragraphs review the time synchronization needs of JADS and how they were fulfilled.

5.1 Time Synchronization Requirement

The JADS requirement for time synchronization was to synchronize the clocks and time stamping to an accuracy of 1 millisecond. This was not easy. Other DIS T&E related programs which were reviewed, were satisfied with a time synchronization accuracy in the order of one (1) second. Establishing a one-second timing accuracy is relative easy. Taking this 3 orders of magnitude better is non-trivial.

Several major obstacles present themselves. In the first place, most computing equipment is not designed with accurate oscillators. Secondly, because we are using non-real-time operating systems, i.e. UNIX, it is difficult to quantify with certainty the accuracy of a timestamp. Just having a precision time source available with which to syncnronize a computer's clock is not sufficient. The logging machine is just another process, and in UNIX, all processes are sharing the computer and its various I/O and other resources. Thus, with a heavily loaded machine, a logger's request for system time may get hung up in the processing queue. And, then also, there's the problem of establishing and then verifying time synchronization over a wide area network with all of the propagation and network protocol devices' error sources.

5.2 Time Synchronization Software

Fortunately, GPS provides the accurate time source needed. And, fortunately, a considerable amount of work has been done to establish time synchronization in local and wide area networks. This work has been captured in a set of software under the name XNTP, developed largely by personnel at the University of Delaware. JADS used XNTP on all the logging machines to accomplish time synchronization. An SGI Indy at the TCAC was lashed up to a GPS receiver. It served as the Stratum 1 time server for all of the other machines on the local and wide area net. Data obtained to date indicate that the required time synchronization accuracy of one millisecond has been achieved.

Each of the simulators used precision time available at their site. Typically, this was a Cesium clock augmented with GPS time. Again, based on data obtained to date, it appears that the simulator and related PDU-formatting machines data was time-stamped to an accuracy of 1 millisecond.

5.3 Time Synchronization Tools

Of course, nothing can be taken for granted, especially time, when collecting test data. Things happen. Power glitches occur, computers halt and get restarted, etc.. Thus, it was incumbent upon the test team to continuously monitor time and its status at each of the logging machines.

XNTP records various sets of statistics that allow the assessment of the health of time. In particular, the XNTP clockstats and peerstats data were used by the JTF to monitor and assess time.

Clockstats provide the system time and GPS time at each GPS update. Peerstats provide the offset of system time to a peer's time. In the case of the Stratum 1 Time Server, peerstats provide the offset of the system time to the GPS time. For the other machines, peerstats provide the offset of the system time to the Stratum 1 time server.

The tools that have been developed by the JTF allow a point-and-click interface to import these clockstats and peerstats data into the Metrica analysis system and to obtain visual representation of the time statistics.

Two example plots are shown.

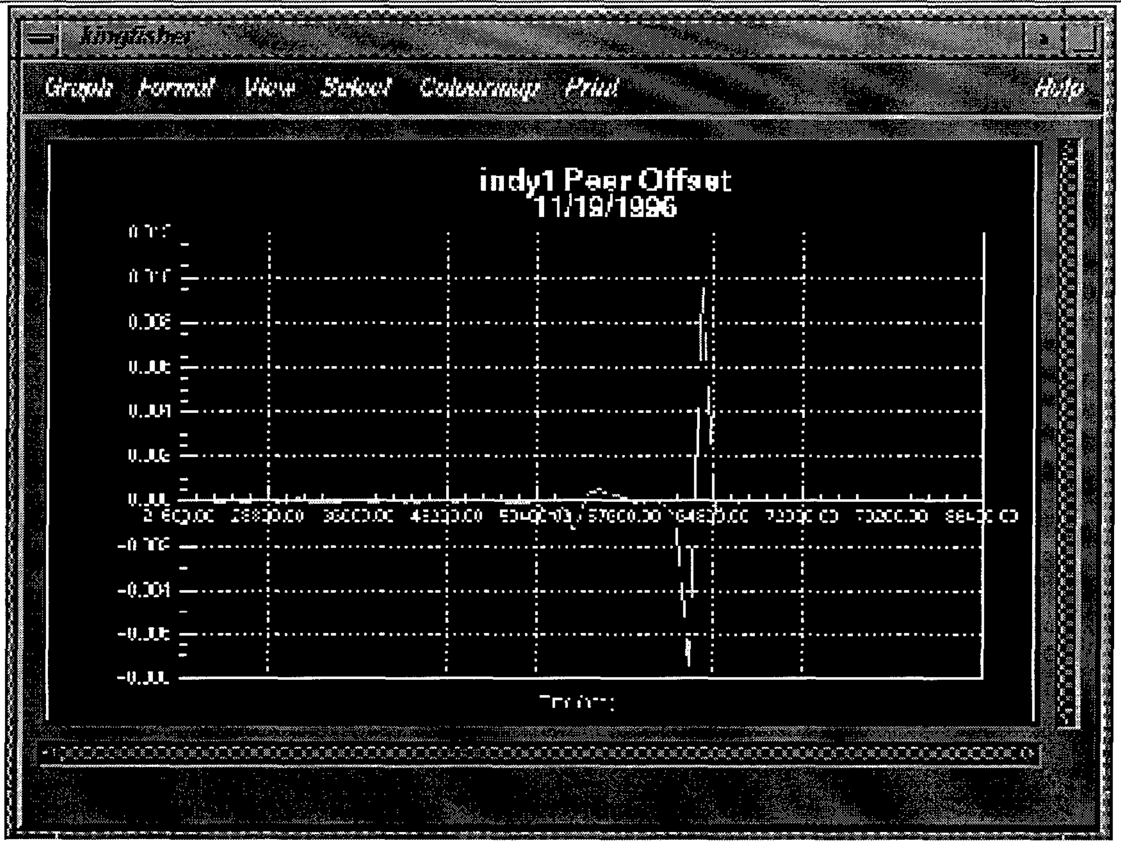

Time Synchronization Plot

This plot shows the peer offset of our Stratum 1 time server to the GPS time source. The last 18 hours of a day is represented along the x axis (each major unit is 2 hours). The vertical axis has 2 milliseconds for each major division. It can be seen that the time server was well within 1 millisecond for the bulk of the time, however, during the last two hours, the server lost it's time synchronization and was off by over 9 milliseconds for an hour or so.

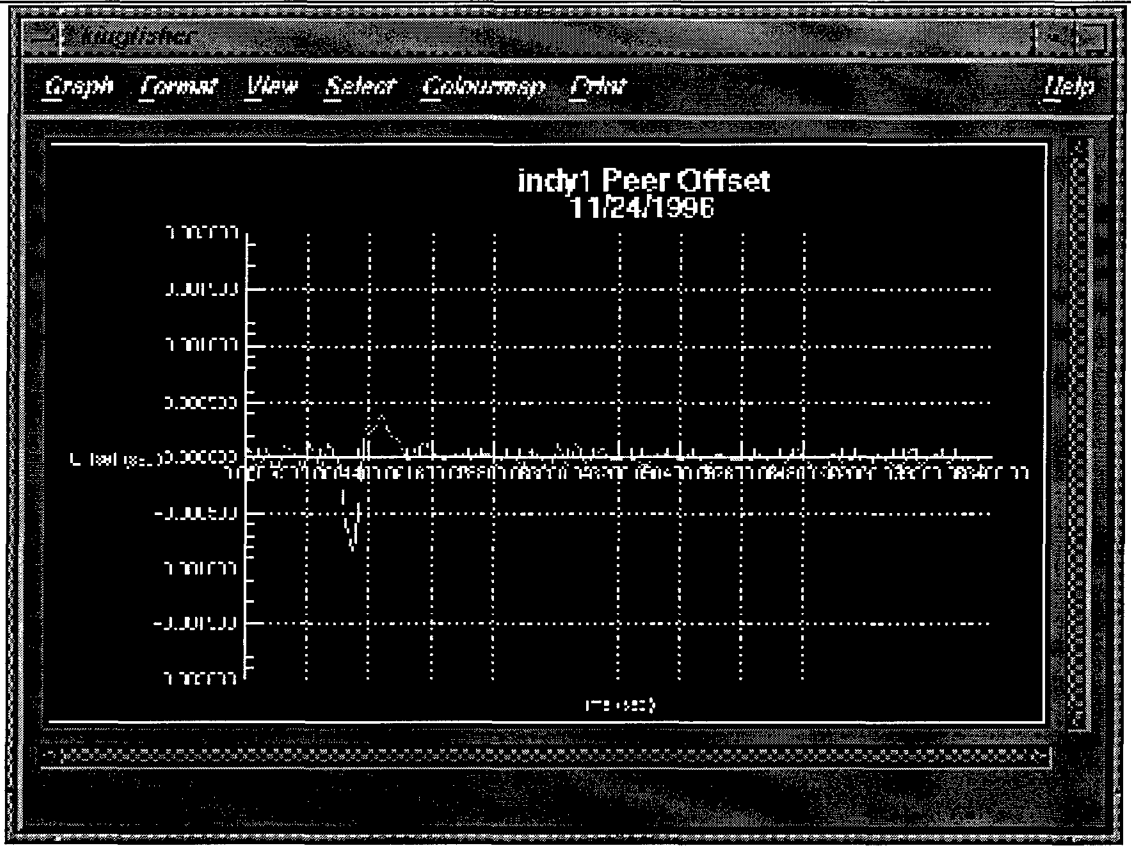

Another Time Synchronization Plot

This plot shows the peer offset data for a different day. The time increments have been changed along both axis. As can be seen, during this 24 hour period, the synchronization of the time server to the source GPS time was within a millisecond.

6. Analysis Tools

The analysis tools that have been developed for looking at PDU data have been grouped into a single menu system using the Metrica Technical Scripting Language. The menu provides for various functions, but this paper will only review the FILE menu items and the INDIVIDUAL ANALYSIS menu items.

6.1 File Menu

Under the FILE menu item provision is made for import of any selected individual trial or for import of a day's set of trial.

Provision is also made for exporting (in an Excel compatible format) any selected mission, trial, logger, and entity data set.

Also, provision is made for review of what exists in the database, so the analyst can decide whether or not importing of source data is required. The contents of the database are displayed by mission day, trial, logger, or entity, or any combination of these.

6.2 Individual Analyses Menu

The Individual Analysis Menu contains three submenus. These are the Tabulations menu, the Plots menu, and the Histograms menu. The discussion that follows will address each individually.

6.2.1 Tabulations Menu

Summaries of PDU arrival statistics and PDU latency statistics can be obtained in tabular form from the tabulations menu. Samples of each will follow.

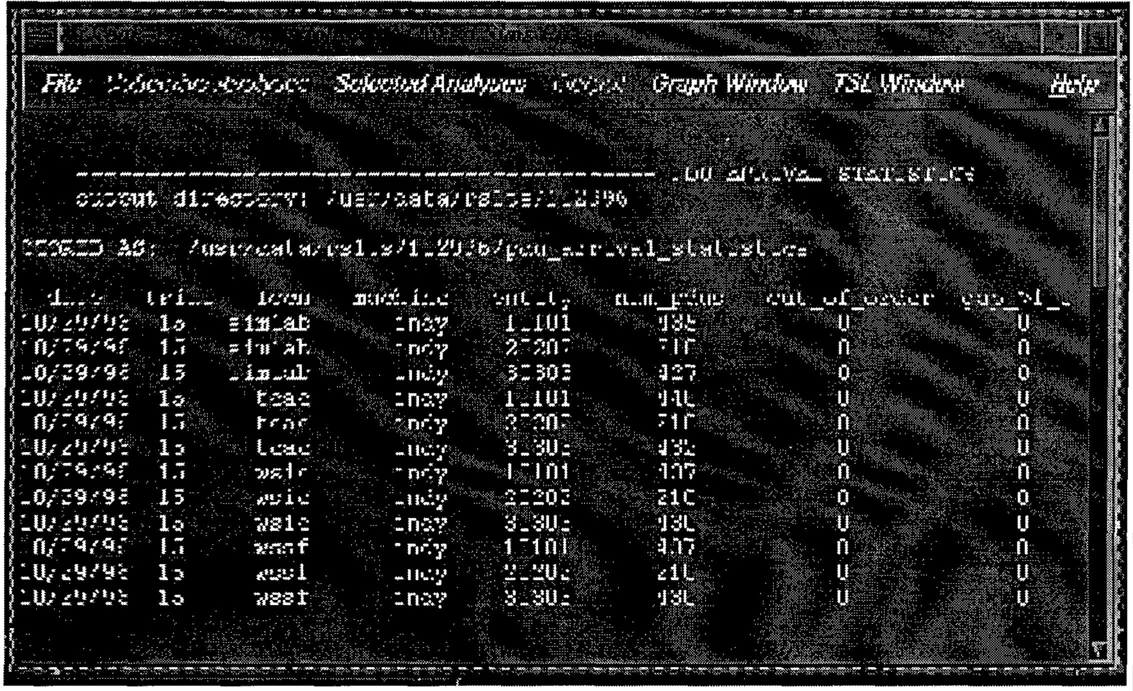

```
   ---------------------------------------- PDU arrival statistics
   output directory: /usr/data/rslts/102998

STORED AS:  /usr/data/rsl.s/102998/pdu_arrival_statistics

 date      trial   team   machine   entity   num_pdus   out_of_order   gap>1_
10/29/98   15    simlab    indy     1_101      435          0              0
10/29/98   15    simlab    indy     2_201      71?          0              0
10/29/98   15    simlab    indy     3_303      427          0              0
10/29/98   15      teac    indy     1_101      41?          0              0
10/29/98   15      teac    indy     2_201      71?          0              0
10/29/98   15      teac    indy     3_303      432          0              0
10/29/98   15      wsic    indy     1_101      437          0              0
10/29/98   15      wsic    indy     2_203      210          0              0
10/29/98   15      wsic    indy     3_303      43?          0              0
10/29/98   15      wsst    indy     1_101      437          0              0
10/29/98   15      wsst    indy     2_202      210          0              0
10/29/98   15      wsst    indy     3_303      43?          0              0
```

PDU Arrival Statistics Tabulation

Shown is a sample of pdu arrival statistics for a selected trial. The output is displayed directly from Metrica on the TSL window. As can be seen, the PDU arrival statistics include the number of PDUs logged, the number of PDUs that arrived out of order, i.e. the time of a received PDU is less than the time of the previously logged PDU, and the number of PDUs that had a time gap greater than 1 second (the 1 second is an arbitrary selection -- the counter is incremented when the time of a received PDU is greater than 1 second later than the time of the previously received PDU).

The tabulation provides a quick quality check of PDUs. In early JADS testing, there were many tests where PDUs were out of order, and there were numerous time gaps between PDUs greater than 1 second. Part of the reason for the time gaps was that the PDU data was from a live aircraft, and telemetry loss caused PDU loss. The reasons for the PDUs out of order were never ascertained. It is believed that the setup of the network interface units may not have been correct.

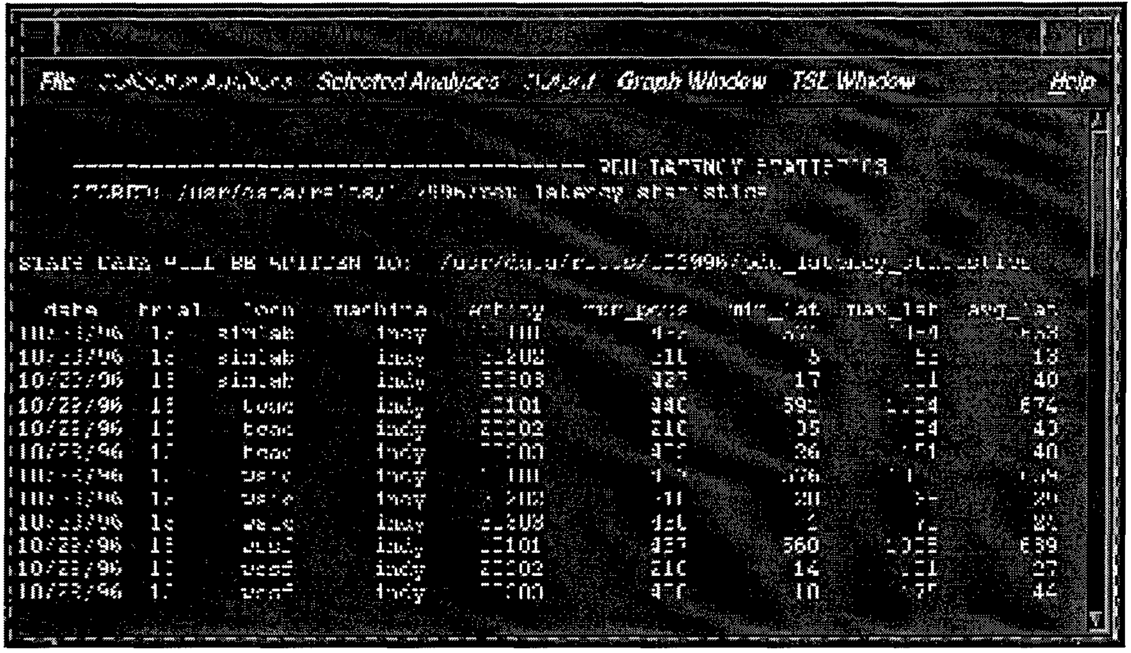

```
  -------------------------------------- PDU LATENCY STATISTICS
  [illegible] /usr/[illegible] latency statistics

STATS DATA WILL BE WRITTEN TO:  /usr/data/[illegible]/pdu_latency_[illegible]

  date      trial  logr     machine  entity   num_pdus  min_lat  max_lat  avg_lat
 [illegible] 1[illegible] simlab  indy  [illegible]  [illegible]  [illegible]  [illegible]  [illegible]
 10/[illegible]/96  1[illegible]  simlab  indy  [illegible]202  [illegible]  5  [illegible]  13
 10/[illegible]/96  1[illegible]  simlab  indy  [illegible]303  [illegible]  17  [illegible]  40
 10/2[illegible]/96  1[illegible]  [illegible]  indy  [illegible]101  [illegible]  59[illegible]  [illegible]  [illegible]
 10/2[illegible]/96  1[illegible]  [illegible]  indy  [illegible]202  [illegible]  [illegible]5  [illegible]  4[illegible]
 10/[illegible]/96  1[illegible]  [illegible]  indy  [illegible]  [illegible]  [illegible]6  [illegible]  40
 [illegible]  1[illegible]  [illegible]  indy  [illegible]  [illegible]  [illegible]  [illegible]  [illegible]
 [illegible]  1[illegible]  [illegible]  indy  [illegible]  [illegible]  [illegible]  [illegible]  [illegible]
 10/[illegible]/96  1[illegible]  [illegible]  indy  [illegible]  [illegible]  [illegible]  [illegible]  [illegible]
 10/2[illegible]/96  1[illegible]  [illegible]  indy  [illegible]101  [illegible]  560  [illegible]  [illegible]39
 10/2[illegible]/96  1[illegible]  [illegible]  indy  [illegible]202  [illegible]  14  [illegible]  [illegible]
 10/[illegible]/96  1[illegible]  [illegible]  indy  [illegible]  [illegible]  10  [illegible]  4[illegible]
```

PDU Latency Statistics Tabulation

Shown is a sample of the PDU Latency statistics tabulation. The tabulation shows the number of PDUs, and the minimum, maximum, and average latency of PDUs for the selected data.

Ideal data for the selected dataset would show a latency of 20 to 30 milliseconds for an entity logged at its node, and 40 to 70 milliseconds for an entity logged at a remote node*. As can be seen, there were numerous apparent latency problems with the selected data.

* The wssf logger was located at the shooter (entity 10101) simulation site.
The simlab logger was located at the missile (entity 20202 simulation site.
The wsic logger was located at the target (entity 30303) simulation site.

6.2.2 Plots Menu

Under the PLOT menu, are the GENERAL PLOTS, SPECIALIZED PLOTS, and ENGAGEMENT PLOTS menu items. All plots are automatically labeled with the mission date, trial number, logger, and entity as applicable. Also, the x and y axes are labeled with the selected parameter names. Examples of each of these plot items follow.

6.2.2.1 General Plots

The GENERAL PLOTS menu item allows the user to select any of the database PDU data for the x axis and also for the y axis.

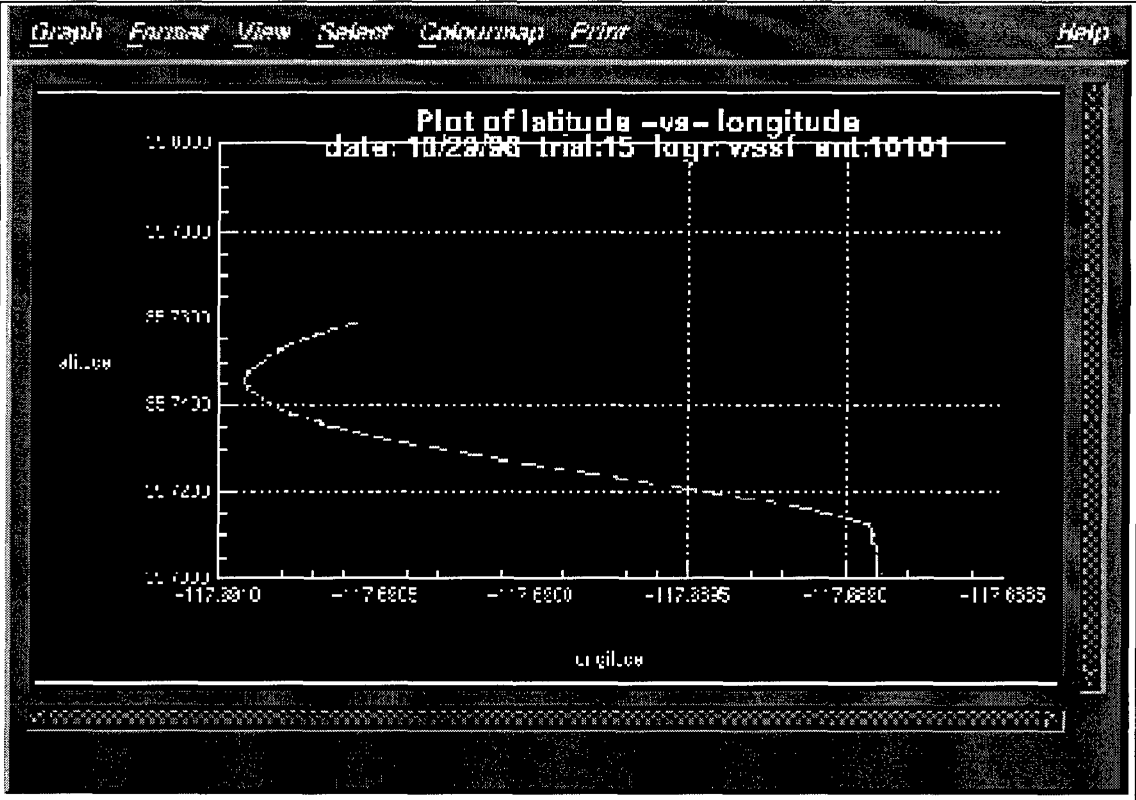
Graph
Format
View
Select
Colourmap
Print
Help
Plot of latitude -vs- longitude
date: 10/28/98 trial:15 logr: vssf ent:10101
-117.3335

Parameter vs Parameter Plot

Shown is a sample of a parameter versus parameter plot where the PDU values of latitude have been plotted against longitude for a selected date, trial, and entity as recorded at a specific logger.

6.2.2.2 Specialized Plots

Under SPECIALIZED PLOTS, the user may select PDU rate plots, Log Rate plots, Latency plots, or Histogram Plots.

The PDU rate plot displays the rate at which PDUs were generated as perceived by the logging device, i.e. PDU time difference between successive logged PDUs.

The Log Rate plot displays the rate at which PDUs were logged by the logging device, i.e. log time differences between successive PDUs.

The Latency plot displays the elapsed time from PDU creation until it is logged by the logging device, i.e. log time minus PDU time.

Histogram plots display the frequency of occurrence for selected intervals of PDU rate or Latency.

Note: **There are two times that are used in this and the following discussions.**

PDU time is the time contained inside the PDU and represents the time when the PDU data were created.

Log time is the time when the PDU was received and logged at a given logger.

Note: **For the LSP tests, the dead reckoning algorithms were disabled so that entities generated PDUs at fixed intervals.**

An example of each of these plots follows.

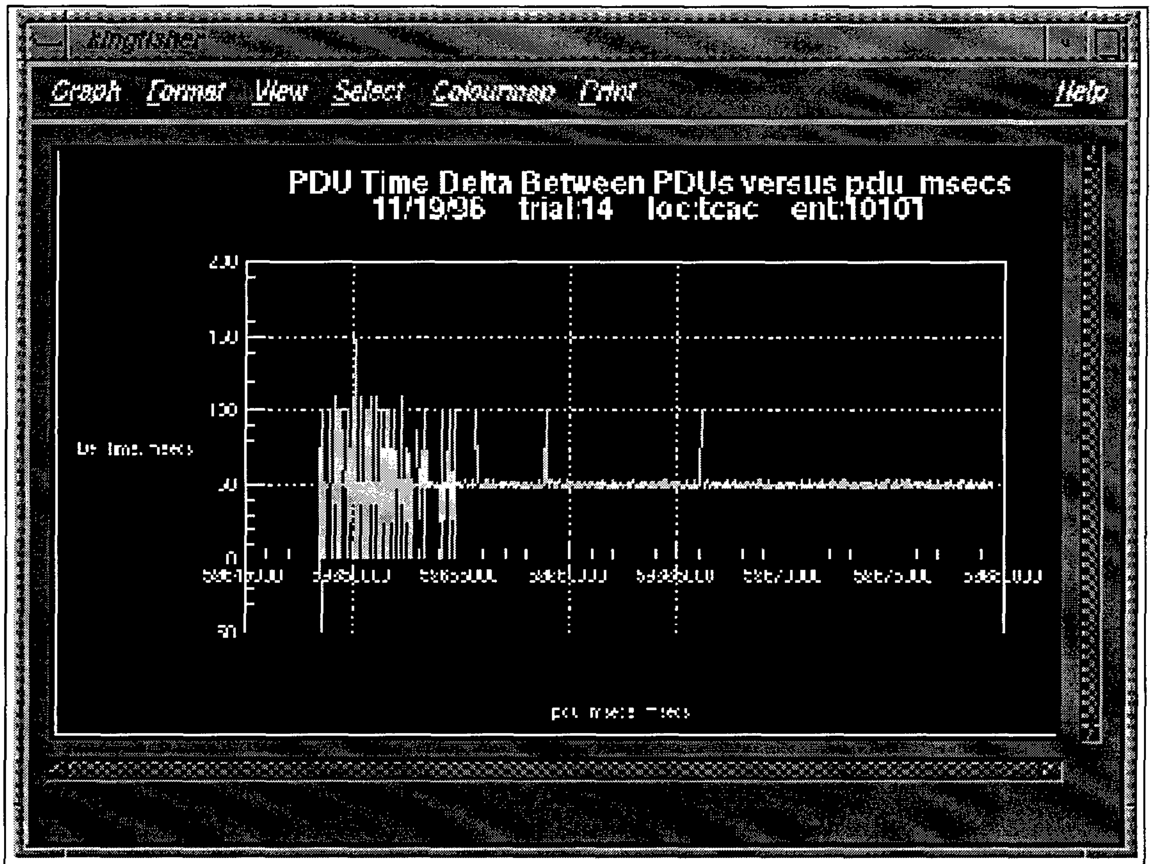
Kingfisher
Graph
Format
View
Select
Colourmap
Print
Help
PDU Time Delta Between PDUs versus pdu_msecs
11/19/96 trial:14 loc:tcac ent:10101
100
0

PDU Rate Plot

Shown here is a PDU rate plot. This plot shows the rate at which PDUs were generated, i.e. (PDU time difference between successive logged PDUs for a given entity), as perceived at the logger. The plot is available as a plot again PDU number (where the first PDU logged for an entity is number 1, the second for that entity is 2, etc.), logged time, or PDU time.

For the particular plot shown, the nominal PDU rate is 50 milliseconds (actually the rate is the reciprocal of this number, i.e. (1.0/.050) or 20 samples per second. It can be seen that the perceived PDU generation rate is not ideal, i.e. for the first five to seven seconds (the scale is 5 seconds per major division), the rate vacillated between 0 and 150 milliseconds between successive PDUs. This is an indication that the simulation is producing PDUs at an uneven rate and/or that PDUs are being lost, or arrive late or out of order at the logger. A review of individual PDUs in this case showed that the simulation was generating PDUs at 20 samples per second as expected, but that the PDU time was not always updated, i.e. two successive PDUs would have the same time, and then the next PDU would be 100 milliseconds later. Thus, the problem in this case is that the simulation network interface unit that created the PDUs did not put the appropriate PDU time in the PDU.

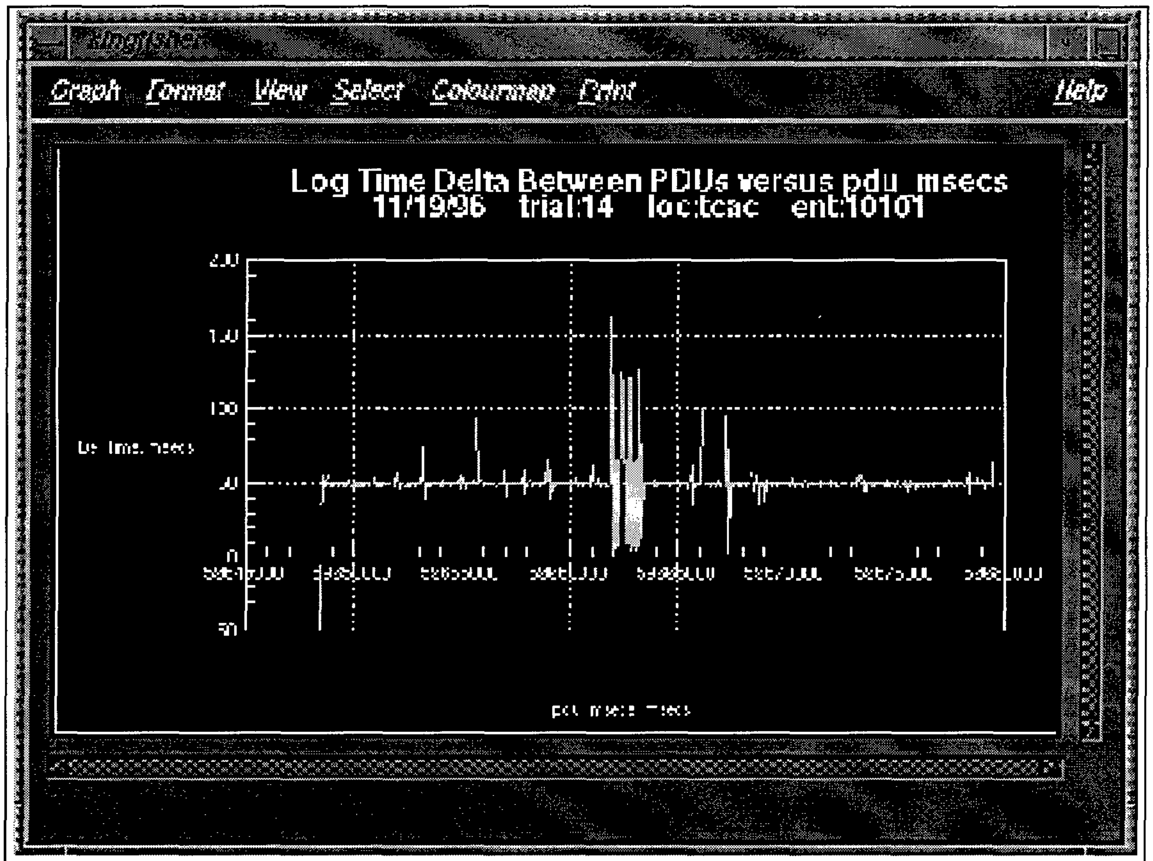
Graph
Format
View
Select
Colourmap
Print
Help
Log Time Delta Between PDUs versus pdu_msecs
11/19/96 trial:14 loc:tcac ent:10101
0

Log Rate Plot

Shown is a log rate plot. This plot displays the log time differences between successive PDUs for a given entity. It shows the frequency with which that entity's PDUs are logged at the logger site. It also may be plotted against PDU Number, Log Time, or PDU time.

The combination of this plot and the previous PDU rate plot may provide some indication of source of data rate problem sources. For the selected data, the Log Rate plot shows no log rate problems in the same time interval where the previous plot showed pdu rate problems. Thus, the PDUs were being logged at the expected time intervals of 50 milliseconds, but the time inside the PDU was not at the expected 50 millisecond intervals. This indicates problems with the PDU itself as discussed in the previous slide.

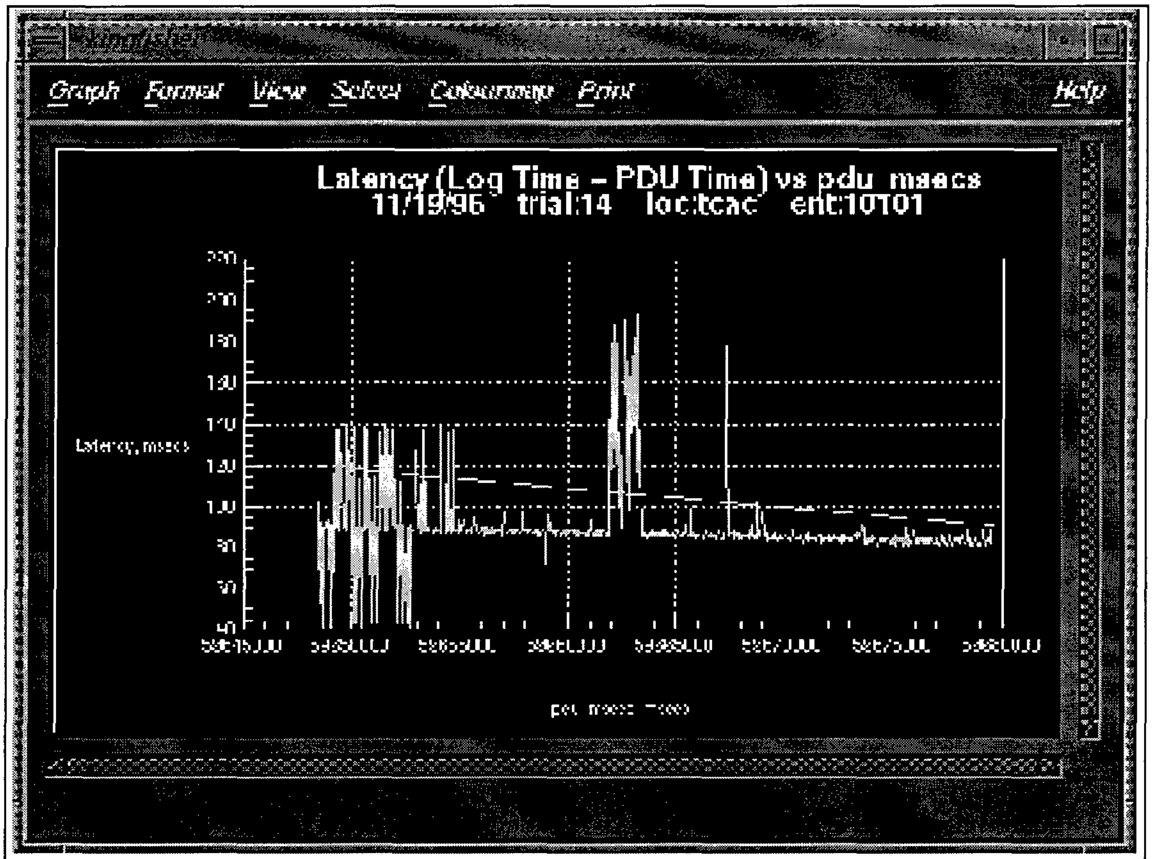
Graph
Format
View
Select
Colourmap
Print
Help
Latency (Log Time - PDU Time) vs pdu msecs
11/19/96 trial:14 loc:tcac ent:10101
Latency, msecs

Latency Plot

Shown is a sample pdu latency plot. It displays log time minus pdu time versus log time. The particular plot shows that the latency from the time the selected simulation PDU was created until it was logged, in general, decreased from about 90 milliseconds at the start of the trial to about 80 milliseconds at the end of the trial, 30 seconds later. This is attributable, in this case, to the simulation's network interface unit, and it's delays in producing the PDUs. The point-to-point variations in latency are attributable to the problem discussed earlier, i.e. the PDUs were not being time-stamped correctly.

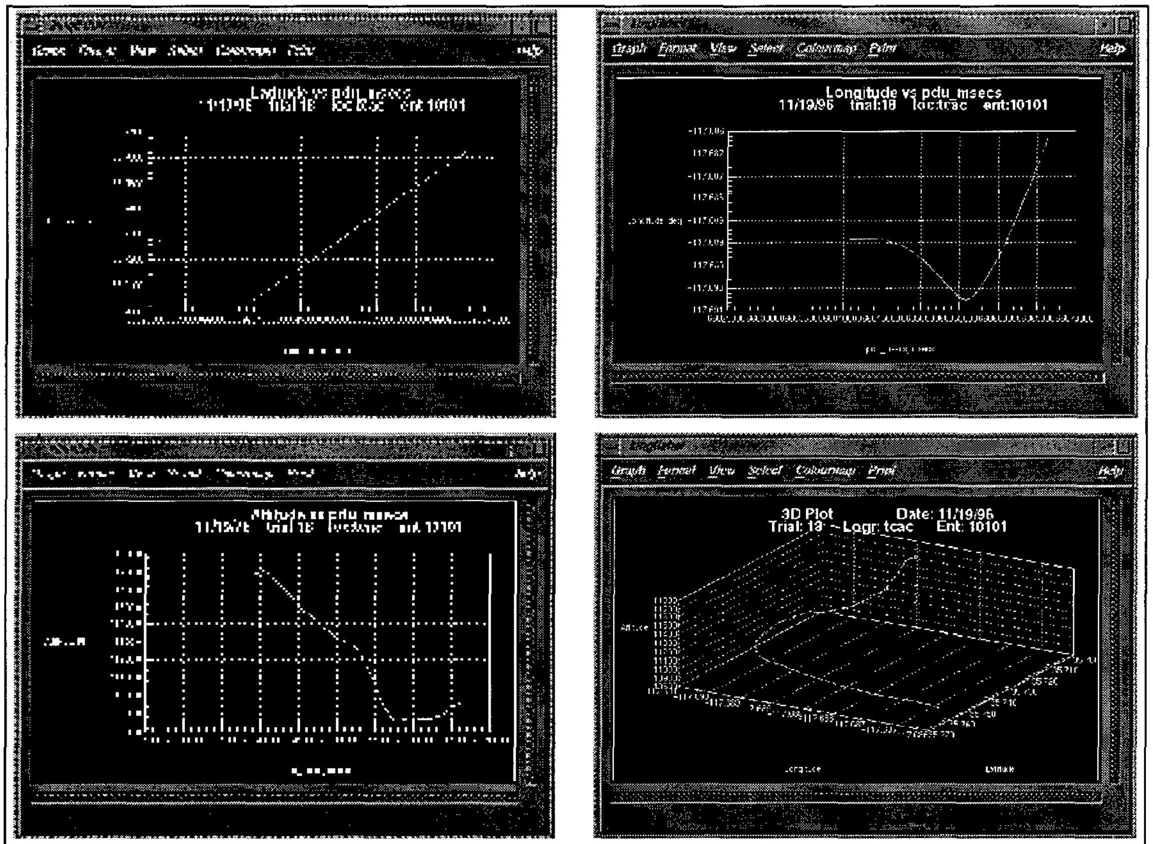
Latitude vs pdu_msecs
11/19/96 trial:18 loc:tcac ent:10101
Graph Format View Select Colourmap Print Help
Longitude vs pdu_msecs
11/19/96 trial:18 loc:tcac ent:10101
Altitude vs pdu_msecs
11/19/96 trial:18 loc:tcac ent:10101
3D Plot Date: 11/19/96
Trial: 18 Logr: tcac Ent: 10101

XYZ Plots

Shown in these plots are the result of selecting the XYZ plot option of the specialized plots menu. Plots of latitude, longitude, and altitude as well as the 3 dimensional view are automatically generated and displayed for the analyst.

The XYZ Plots provides a plot of Latitude, Longitude, and Altitude versus either PDU Number, PDU time, or Logged time. Also, the item provides a 3D view plot of the selected entity.

6.2.2.3 Engagement Plots

Two options are provided in the Engagement Plots menu. One is the missile planview plot. The other is the 3D view (of the engagement) plot. An example of each is included.

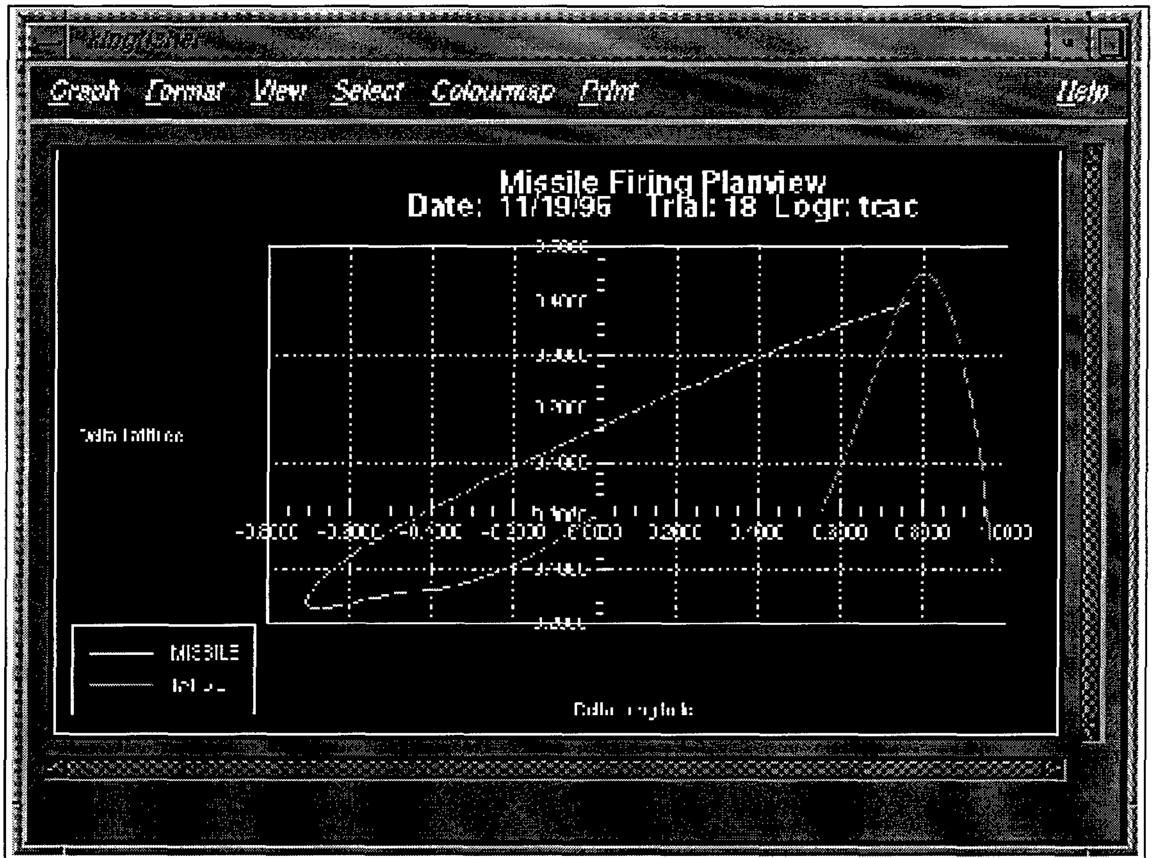

Missile Planview Plot

Shown here is a specially concocted missile planview plot. Actual data could not be used because they are classified. The plotting routine limits the data to that occurring during the missile flyout only, i.e. the target is not displayed prior to missile launch, or after missile flyout termination. The missile always starts at the origin of the plot. The x coordinate in this case is longitude, and the y coordinate is latitude (minus the missile latitude and longitude respectively at missile launch).

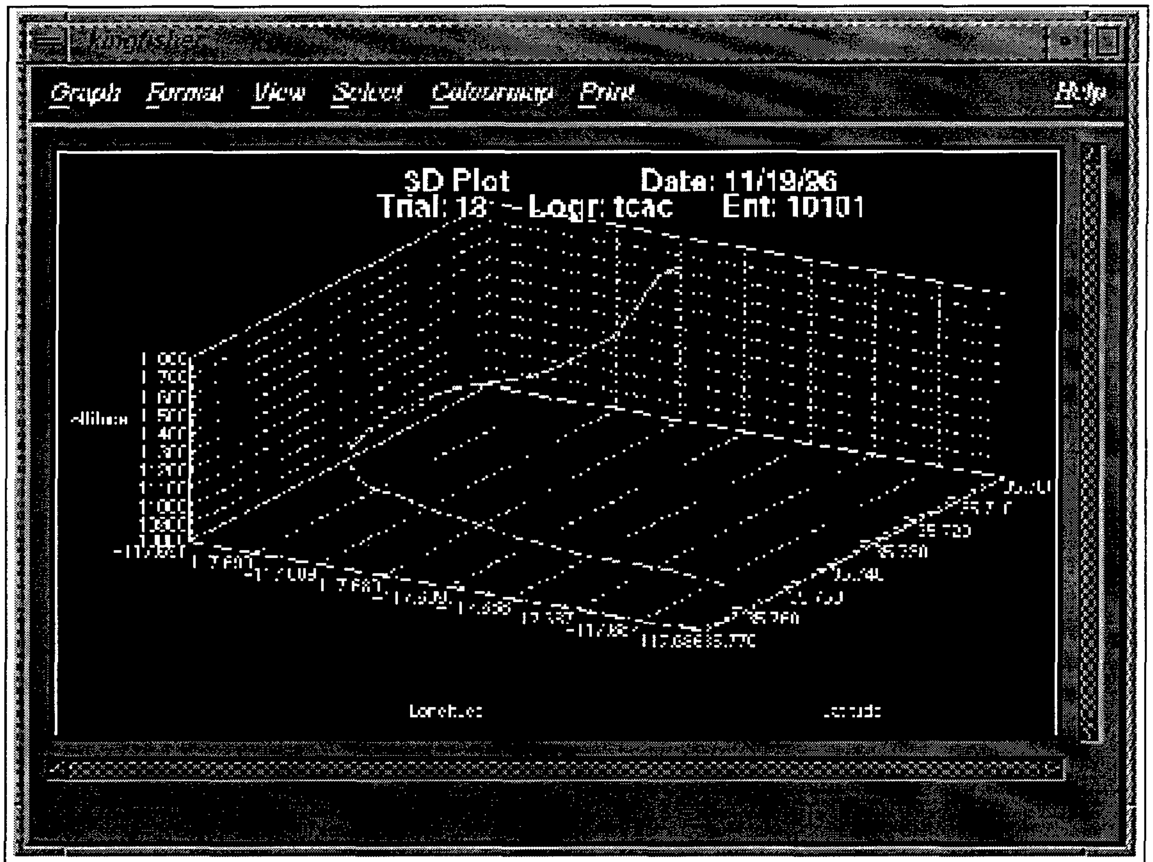

3D View Plot

Shown is the result of selecting the 3D view under the engagement plots option of the menu.

6.2.3 Histograms Menu

The PDU rate and PDU latency plots shown earlier provide a display of PDU rate and latency for each sample. A histogram of these data can be obtained under the histograms options of the specialized plots menu. Samples of these plots follow.

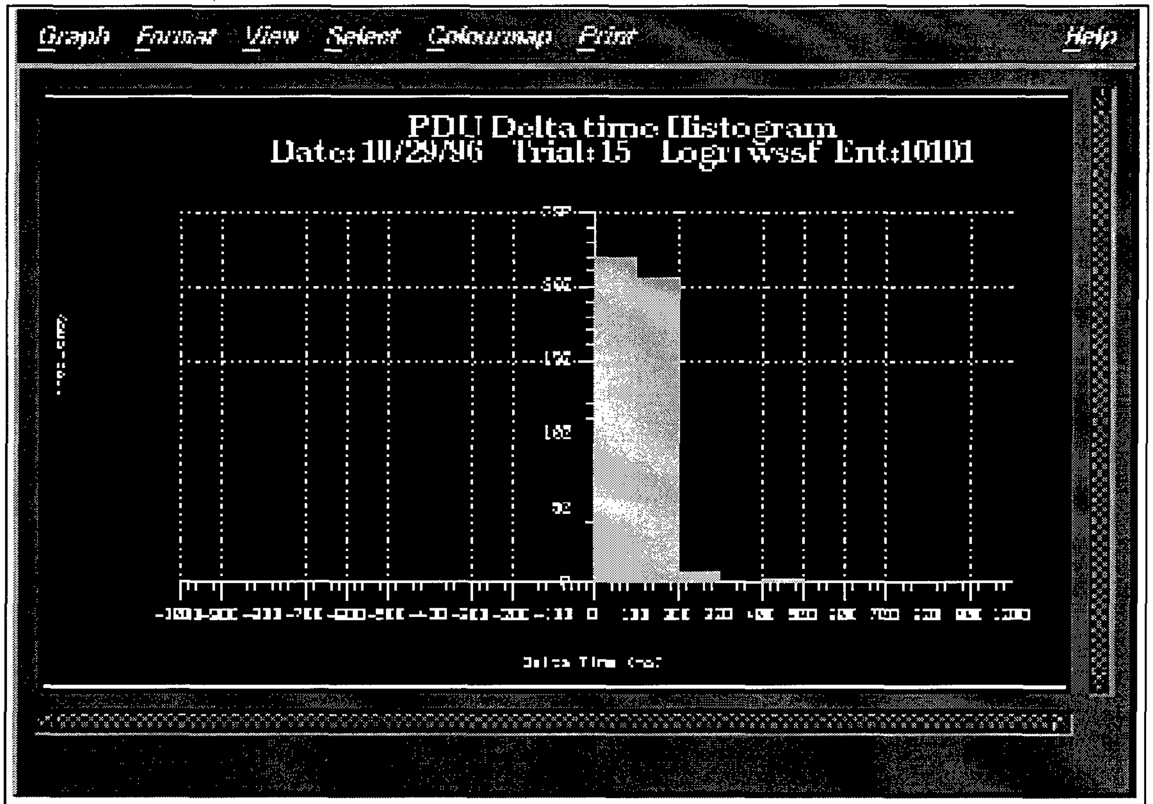

PDU Rate Histogram

Shown is a sample PDU Rate histogram. As can be seen, the PDU rate was not ideal for the selected trial, entity. For this particular simulation entity, the PDU rate was set at 20 samples per second (dead reckoning was disabled). Thus, without problems, all of the PDUs would have occurred within 100 milliseconds of each other. The histogram shows that about half of the PDUs were somewhere between 100 and 200 milliseconds apart, and that there were about 10 PDUs that were over 200 milliseconds apart, and 2 or 3 that were about 500 milliseconds apart. This behavior is likely caused either by the simulation PDU generation device, by the network dropping PDUs, or by network delays of the PDUs.

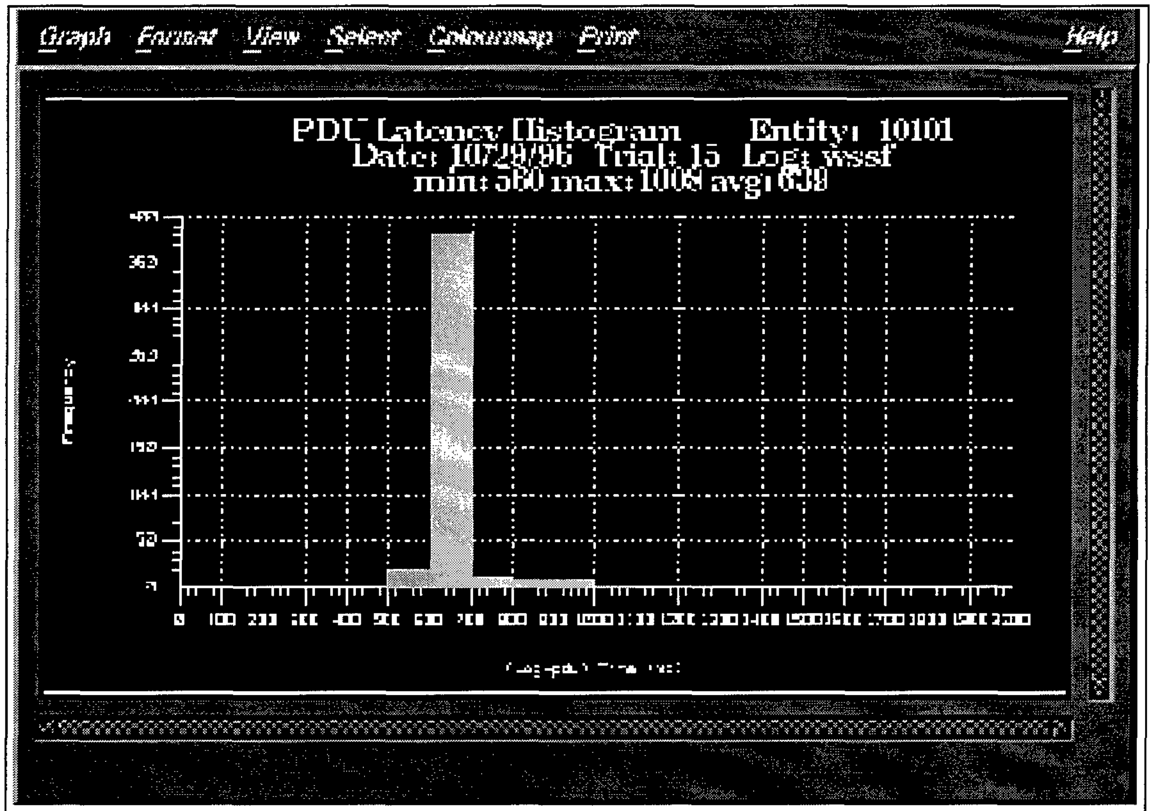

PDU Latency Histogram Entity: 10101
Date: 10/28/96 Trial: 15 Log: wssf
min: 580 max: 1008 avg: 638

PDU Latency Histogram

Shown is a sample PDU latency histogram. For the specific data selected, it can be seen that the latency (time delay from the time the PDU was created until it was logged varied between 500 and 1000 milliseconds. What is selected is the logger at the same node as the displayed entity. The expected latency is in the order to 20 to 50 milliseconds. Delays of the magnitude shown indicate problems in any of several places. In this particular case, the most likely source of the problem is the Network Interface Unit using bad time, or somehow introducing a half-second delay prior to time-stamping the PDU.

7. Summary

It has been our experience that great gobs of data are generated by ADS systems, and that great gobs (a gob is different and distinct from a heap) of results can be generated by toolsets such as the ones presented today. The challenge for an analysis team is in the selection of the appropriate data and the display of the appropriate data, so as to not become inundated.

The tools developed provide a very flexible point and click interface to the source data. They can "handle" great gobs of data, i.e. the various plots and tabulations shown are generated within seconds. In this manner, the analyst can quickly preview data looking for patterns, etc., and then press on to develop more detailed analyses of the pertinent data.

The tools shown have utility for any ADS-based test program in that they provide the tester with a means of analyzing data quickly and easily.

The tools are portable, since they are based on a UNIX operating system.

The tools are in a state of development. A snapshot of the current state of these tools has been presented. A more comprehensive and refined toolset will be developed before this program is over. These tools could be a legacy to assist future test programs utilizing ADS.

INTERNET DOCUMENT INFORMATION FORM

A . Report Title: Collection and Analysis of Quality Data in a Distributed Simulation Test Environment

B. DATE Report Downloaded From the Internet 1/28/99

C. Report's Point of Contact: (Name, Organization, Address, Office Symbol, & Ph #): Joint Advanced Distributed Simulation
Joint Test Force
ATTN: Ann Krause (505) 846-1291
11104 Menaul NE
Albuquerque, NM 87112-2454

D. Currently Applicable Classification Level: Unclassified

E. Distribution Statement A: Approved for Public Release

F. The foregoing information was compiled and provided by:
DTIC-OCA, Initials: VM_ **Preparation Date:** 1/28/99__

The foregoing information should exactly correspond to the Title, Report Number, and the Date on the accompanying report document. If there are mismatches, or other questions, contact the above OCA Representative for resolution.