A Perspective on DSN System Performance Analysis

Timothy Pham
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91214
Timothy.Pham@jpl.nasa.gov

Tracking Number: 56704

Abstract – This paper discusses the performance analysis effort being carried out in the NASA Deep Space Network. The activity involves root cause analysis of failures and assessment of key performance metrics. The root cause analysis helps pinpoint the true cause of observed problems so that proper correction can be effected. The assessment currently focuses on three aspects: (1) data delivery metrics such as Quantity, Quality, Continuity, and Latency; (2) link-performance metrics such as antenna pointing, system noise temperature, Doppler noise, frequency and time synchronization, wide-area-network loading, link-configuration setup time; and (3) reliability, maintainability, availability metrics. The analysis establishes whether the current system is meeting its specifications and if so, how much margin is available. The findings help identify the weak points in the system and direct attention of programmatic investment for performance improvement.

I. Introduction

Recently, the Deep Space Network (DSN) re-instituted a formal activity on the system performance analysis. Although this activity was in practice about two decades ago, it was neglected in the 1990's because of change in programmatic priorities. In November 2003 - July 2004, there were a confluence of mission critical events that DSN supported. Among these are the landings of two Mars Rovers Spirit and Opportunity, the orbit insertion of Mars Express, the Wild-2 comet encounter of Stardust, the launch of Deep Impact, the launch of Messenger, and the Saturn orbit insertion of Cassini mission. At the same time, there was a recognition that some equipment in DSN were sufficiently old, lacking necessary spares, and at risk of being able to provide adequate support. As a result, several new implementations were planned, with an operational date in late 2003. These deliveries and the criticality of follow-on mission events heightened the need for assuring the DSN readiness; thus, re-seeded the interest in system performance analysis.

Soon after a programmatic review in August 2003, a formal effort on performance analysis was re-established. The objectives are (1) to continually ascertain the system operational readiness, (2) to assess if the DSN is meeting its data delivery commitment to mission customers, and (3) to monitor system performance margin. Included in the analysis is the identification of the performance weak links, from which recommendations for improvement would be issued.

Figure 1 provides the context of performance analysis in the DSN product development/operational life cycle. The performance analysis activities take place in the Operation phase, once the ground system capabilities are committed to mission support. Both nominal operational data and failure-related data are input to the analysis. Other inputs include maintenance data and design/operability-related data. The results are fedback to Operations staff to correct for problems that can be fixed operationally, and to management and engineering staff to plan necessary software/hardware upgrade.

This paper presents the current progress in the DSN performance analysis. It describes the functions and processes within this activity. It focuses on data quantity, quality, continuity and latency (QQCL) metrics committed to mission users. The comparison between actual and committed performance is presented. Problematic areas that require attentions are also identified. Several interesting and necessary considerations required for proper

QQCL accounting are highlighted. Among these are the needs to account for expected outages (versus unexpected ones), and proper accounting of availability (versus only schedulable).

Besides typical telemetry, tracking, command (TTC) services, the analysis also addresses system performance in the areas of service management and science support. These areas require a different set of monitor metrics, compared to TTC. Instead of looking at successful frames/packets and link margin, the interest is on input and output errors and execution success rate for service management metrics. For science-support services, frequency and amplitude stability are a few key parameters of interest.

In addition to QQCL metrics, the analysis team also focuses on performance on another set of key parameters. The importance of these metrics, relevant to the mission operations, is explained. The observed current performance, available margin, and trends of relevant metrics are presented. Relevant lessons learned through this effort are also highlighted.

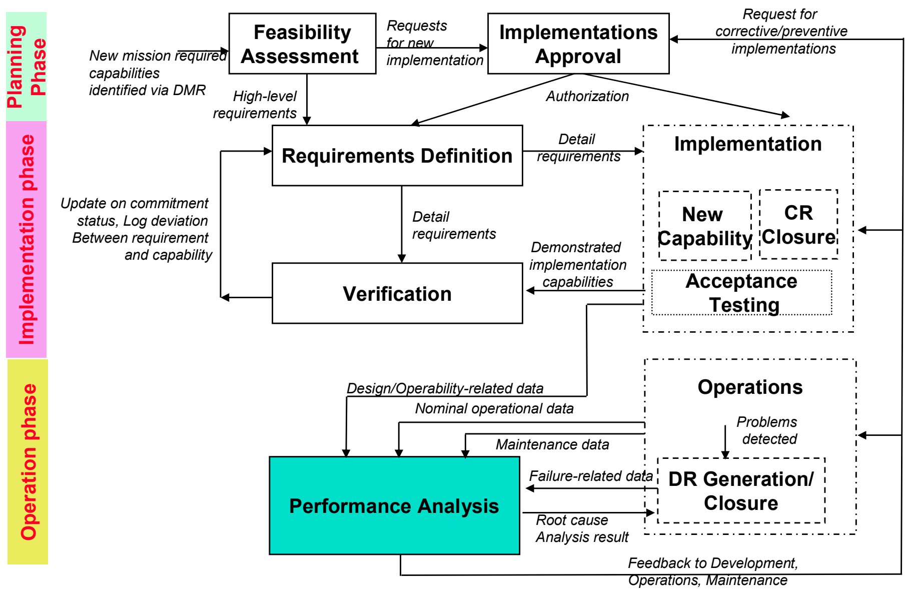

Figure 1. Context of Performance Analysis within DSN Development and Operation Cycle

II. Structure and Processing

Figure 2 shows the inputs, outputs and processing context of Performance Analysis. The inputs are coming from three different areas: operational data, behavior data, and reference model.

The operational data are actual performance of the system on the daily basis. They include:

(1) The *Discrepancy Reports* identifies any data outage that occurs during a tracking session. This rich data set identifies the system components at which error occurs, e.g. antenna servo controller, ranging processor, etc. It is identified by the mission set and ground tracking antennas. The data impact - whether it's complete lost, degraded or recoverable – is also categorized. The degraded category refers to performance degradation, such as large discrepancy in received signal-to-noise ratio. The recoverable data refer to temporary outages that are caused by failure in system components involved in data delivery but not data capturing.

(2) The *Link Monitor Log* contains a rich set of information related to system performance. It includes key parameters related to the link such as measured bit/symbol rates, signal-to-noise-ratio, system noise temperatures, spacecraft carrier frequency, antenna pointing correction under conical scanning or monopulse operation, etc. The log also contains alarms and warnings issued by the subsystems, which is important for failure diagnostics. The log further captures the transaction on system configuration; thus, allow derivation of parameters such as setup time and signal acquisition time.

(3) The *Support Data* are products needed for tracking. They cover tracking schedule, e.g., what mission to be tracked at what time on which antenna. They include tracking predictions, e.g., antenna pointing, received frequencies, spacecraft telecommunications configuration parameters. The schedule data sets the baseline from which data return and service availability is computed. The predictions help identify any potential misconfiguration when a problem occurs.

(4) The *Hardware Maintenance Records* is the log of works done on equipment, covering both preventive maintenance and repairs. This data set enables assessment of maintenance effort in providing the services. It identifies what equipment requires the most care, and thus, allows for possible replacement decisions. The data also enable characterization of mean time to repair failed components; thus, enable better reliability modeling of system behavior to maximize operation support.

The behavior data include test reports and anomalies identified at the time of new equipment delivery. They also include work-around procedures for functions that did not function as intended. These data points establish the degree of system operability. The system operability measure should correlate to the number of failures. A system that is hard to operate because of too many anomalies and work-around would likely be prone to errors. This is especially true when the operators are not the system designers.

The reference models are a set of baseline performance from which comparison against the actual operation performance is made. They include the performance models for various signal conditions, e.g. receiver loss as a function of signal to noise ratio and tracking bandwidth, or system noise temperature as a function of elevation. Included in this set is also the requirements established at the time of system development and validated at the commissioning phase.

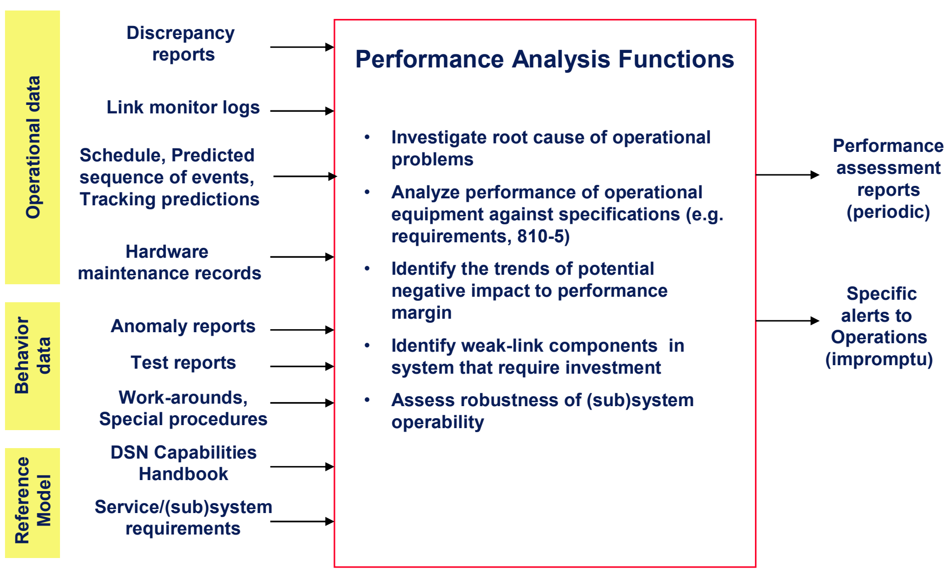

Figure 2. Context of performance analysis functions

Within the performance analysis, the following functions are carried out, as reflected in Figure 3:

(1) *Root cause analysis* – This analysis enables the true characterization of the problem, in particular for those that exhibit different symptoms. This process corrects for any incorrect attribution due to limited understanding at that time when problem surfaced. The root cause analysis also helps direct attention to the list of common problems; thus, heighten the chance for correction.

(2) *Assessment* – Here, the system operating performance is evaluated against the requirements and reference models. The data quantify whether the DSN is meeting its commitments to mission users, how much margin is there, and what future trends look like. The assessment helps to characterize the effort required to achieve the service commitments. It identifies the weak links that merit investment. The weak link

could be one piece of equipment near its design life cycle and requires lots of maintenance attention. It could also be a software module with operability problem that can benefit an upgrade.

The analysis results are feedback to the DSN engineering and management, so that appropriate action can be taken. Typically, the data are presented at a standard forum at periodic interval. Engineering change requests for improvement are submitted and funding solicited. However, when the detected potential problem has significant and immediate impacts, the finding would be directed to station maintenance personnel for immediate resolution.

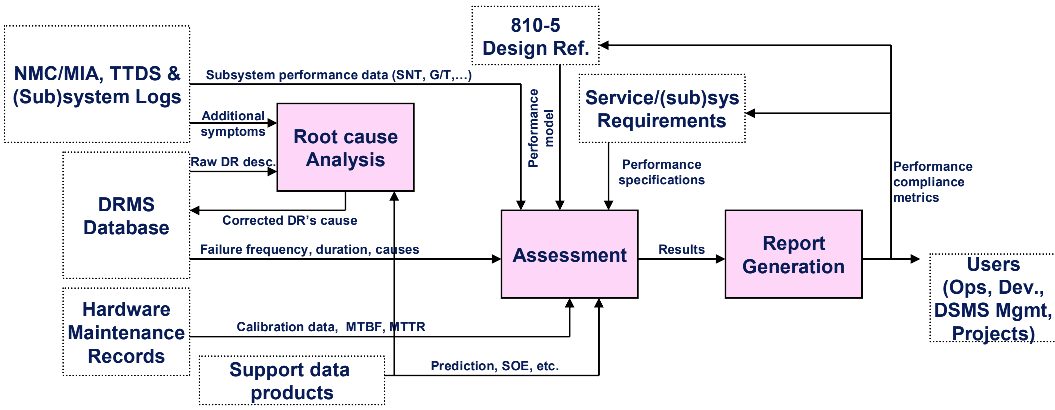

Figure 3. Key Processing Functions

Figure 4 shows a more detail view of processing within performance analysis. First, the data need to be captured and archived. Then, data extraction and mining follow. The assessment is focused in four areas. One focus is on data quantity, quality, continuity, latency (QQCL). These metrics are part of the DSN service commitments to enable missions meeting their science analysis objectives. Second focus is on link performance and margins. Reliability, maintainability and availability (RMA) aspect is also evaluated. And the last category is on design robustness. Trending analysis are conducted with projection for future performance.

At this point, most high-level assessments are done by the engineering team members. Automated data extraction capabilities have been developed to aid with the analysis.

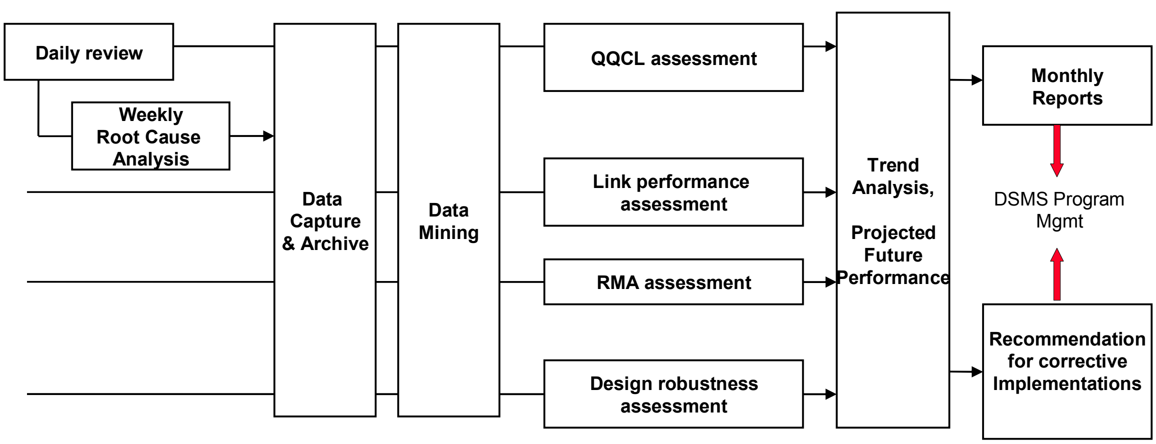

Figure 4. Performance Analysis Internal Processing

III. QQCL Assessment

The assessment on data quality, quantity, continuity and latency establishes the performance metrics of delivered data. This evaluation is attempted on key products, namely, telemetry, tracking and command data and radio science data.

Some of the goals set forth by the performance analysis team are:

- Institute a DSN-internal capability for determination of QQCL metrics. Up to now, the validation is relied on the mission users, e.g., the data management team. Up to recently, the only available internal DSN accounting metric is on data quantity. And such accounting is based on time (i.e., outage time versus scheduled time), rather than frame counts.
- Correlate QQ metrics between actual versus expected on the per pass basis. Events that cause expected outage such as spacecraft transitions from 1-way to 2-way configuration should be factored in the accounting.
- Establish an automated metric reporting mechanism in real time or near real time (e.g., minutes within the pass completion). It is expected that these data would be made accessible to the interested parties via web interface.

(1) *Quantity Assessment*

The benchmark commitment is 95% data delivery for nominal operations and 98% delivery for critical operations. The increase in data delivery percentage is due to additional staffing during critical period to quickly address any potential problem. It includes the option of using addition antenna as backup.

Sample of telemetry and tracking data return is shown in Figure 5*. The frame counting for two missions Cassini and Mars Reconnaissance Orbiter are given for January and February 2006. At this point, however, a correlation on quantity of expected versus actual is not yet available because of the delay in processing expected figure.

	DOY	Quantity, Total Frame Delivered	Out-Of-Sync Frames	Quality, %
Cassini	026-064	1165475	0	100
MRO		1994844	145	99.99

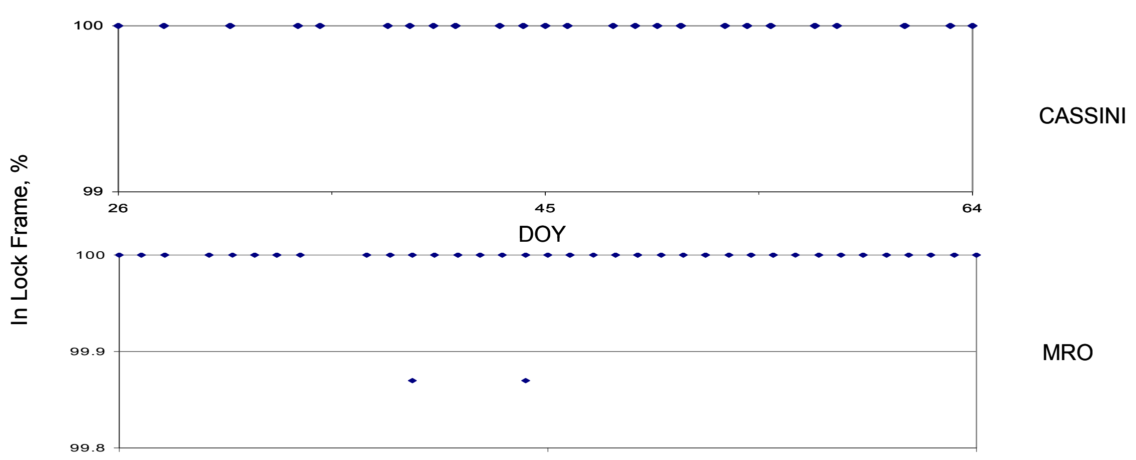

Figure 5. Telemetry frame accounting

* Dong Shin, "Telemetry and Tracking Data Assessment, January – February, 2006", JPL Internal Report

Applying the quantity metrics to the command data turns out to be problematic. The challenge is that the information on expected command link transmission units (CLTU) is not known. In command session, the mission operation team has lots of flexibility in determining the time within a pass to radiate the data, or to cancel the uplink session. These decisions are often made in near real-time – unless it is a critical operation. As a result, the knowledge on expected frames to be transmitted is not known by the DSN Operations.

To mitigate this effect, the performance analysis team keys into a metric that would define the problems that occurs during a command pass. The events of "CLTU session abort" are used for computation. When this event occurs, it signifies that there is a problem with the current sessions which no longer support commanding; thus, forcing the mission operations to re-establish the connection. The metrics of command abort is presented in Figure 6[†].

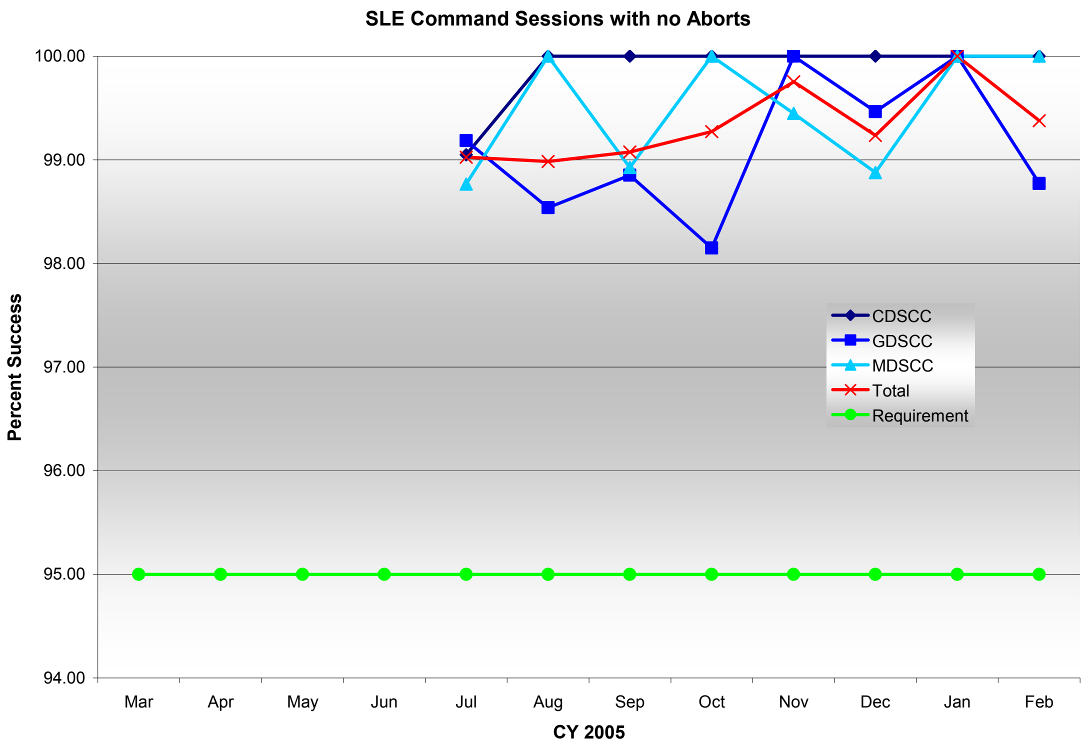

Figure 6. Figure of merit for command quantity/quality metrics

(2) *Quality Assessment*

The benchmark commitment for telemetry is a maximum frame rejection rate of 10^{-6} for concatenated convolution and Reed Solomon code. For turbo coding, the maximum error rate is 10^{-4} for long frames (8912 bits) and 10^{-5} for short frame (1784 bits).

Figure 5, Cassini support is meeting the specification. For MRO, the return is 99.99%, translating to a frame rejection rate of 10^{-4} which is meeting the specification for long turbo code frames.

(3) *Continuity Assessment*

The DSN commitment for telemetry data continuity is maximum 8 gaps per 10,000 frames. Gaps are defined as consecutive undecoded frames. A sample of actual statistics for Mars Reconnaissance Orbiter mission produced by the Data Management Team is shown in Figure 7.[‡]

A study on Odyssey gap statistics by C. Eggemeyer at JPL[§] revealed some interesting patterns. The gaps appeared to have bimodal signature, of short and long duration. Both types incurred the same amount (1%) of

---

[†] T. Cornish, "Uplink Performance Assessment, January – February, 2006", JPL Internal Document

[‡] Data are made available from Multi-mission Data Management Team, http://ddo/ddo/dmt/

outage. Short gaps had an averaged width of about 300+ packets, or less. Within this category, 20% of gaps occurred over hand-over period between two tracking stations; the other 80% was within one station. In contrast, long duration gaps have an averaged span 10 times greater, about 4000+ packets. The bulk of gaps (above 95%) occurred over the hand-over.

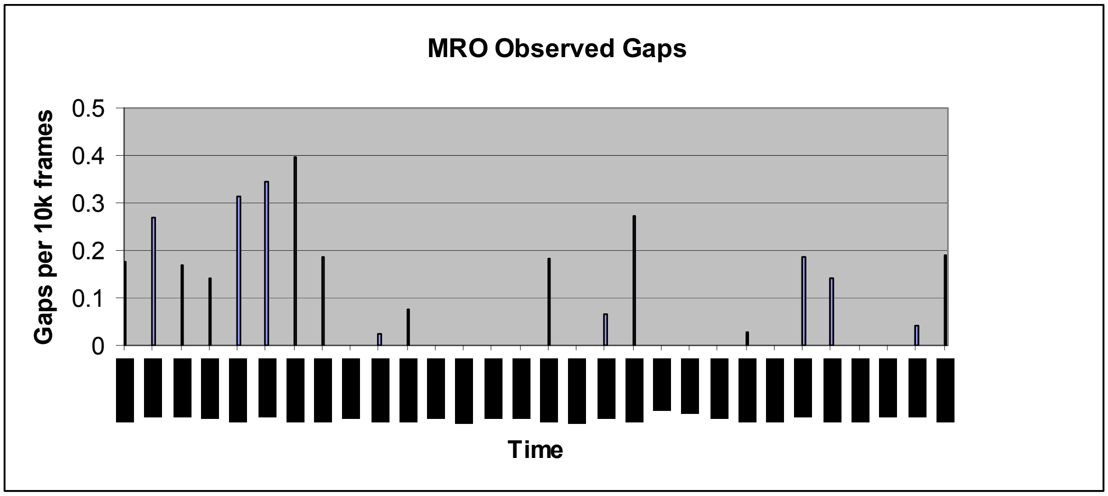

Figure 7. Observed gap statistics on MGS mission

(4) *Latency Assessment*

The DSN commitment to delivery latency is as follows:

- *Telemetry:* 10 seconds for real time engineering data; less than 24-hr for scientific data
- *Tracking:* within 30 minutes in special support configuration; within 24 hrs for normal operation.
- *Command:* within 10 ms of command transmit time. The short duration is a result of either real-time commanding or timetag commanding.

Figure 8 shows the averaged and maximum latency for telemetry science data delivery for MRO mission over January and February 2006. The latency in this case is defined as the difference between Earth-received time recorded at the DSCC tracking complex and arrival time at the time at the Network Operation Center at JPL. An average latency of about 1 second indicates that the data are delivered timely. The worst delay seen is about 1 minute, which is very small compared to the promised 18 hrs in the case of MRO commitment.

The latency observed for real-time telemetry data is shown in Figure 9. This data is for SOHO mission, collected over January-February, 2006 time frame. The worst delay is less than 5 seconds, with an average of 0.5 second. There is a wide margin in performance relative to the 10-second commitment.

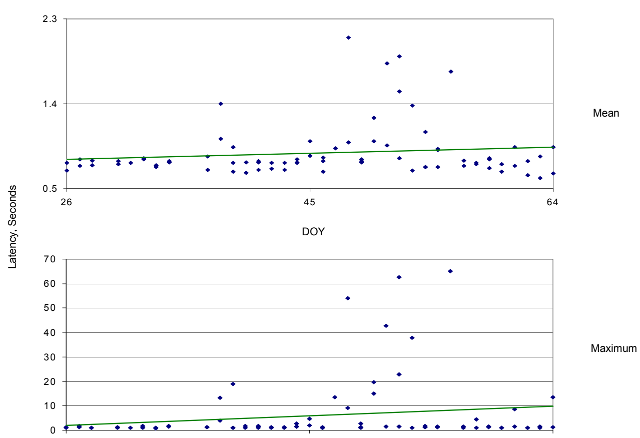

Figure 8. Latency of telemetry science data, for MRO

§ C. Eggemeyer, JPL Internal Report on Gap Statistics,

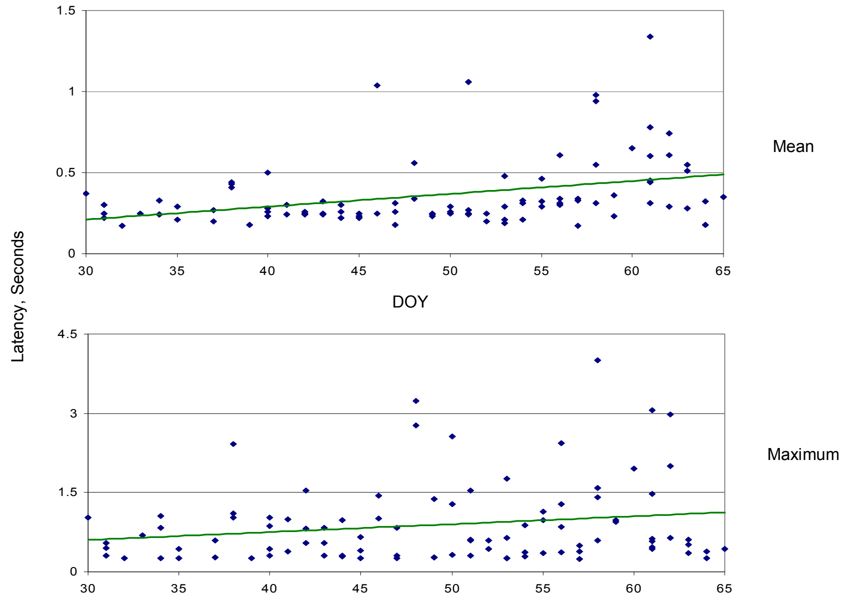

Figure 9. Latency of real-time telemetry data, for SOHO

IV. Link Performance Assessment

On link-related performance, some key parameters of interest are:

(1) Antenna pointing
(2) System noise temperature
(3) Doppler accuracy
(4) Frequency and amplitude stability
(5) Wide-area network (WAN) bandwidth loading
(6) Synchronization of frequency and timing references
(7) Link configuration setup time

A. Antenna pointing

As the DSN is moving toward higher operating frequency Ka-band (32 GHz), the antenna beam width becomes smaller and the antenna gain accentuated. Thus, antenna pointing becomes more critical to realize the gain in performance. The requirement on Ka-band pointing at the 34-m antenna is targeted at 4 mdeg blind pointing and 2 mdeg monopulse pointing. The blind pointing is needed to enable accurate tracking of natural planetary body such as quasar radio sources that serve as a calibrator for Delta differential one way ranging. To a lesser extent, it is also needed to bring the antenna close to the spacecraft target so that monopulse tracking acquisition can be successful. Once acquired, the monopulse close loop tracking algorithm would enable the antenna to stay on source. Reference TBD provides more detail discussion.

The need for X-band pointing is much less because of broader beam width. For the 34-m, an X-band pointing can be relaxed to 8 mdeg, if necessary, and still result in less than 0.2 dB pointing loss. For the 70-m antenna, an X-band pointing of 4-mdeg is still needed to maintain 0.2 dB loss.

Figure 10 provides a sample of actual pointing performance of the 34-m and 70-m antenna at Canberra[**]. The data are derived from pointing correction during X-band conical scan (conscan). On the average, the 34-m pointing error is about 7 mdeg, and 70-m antenna 3 mdeg. These two values keep the pointing loss within 0.2 dB level for X-

[**]Mark Harris, CDSCC Pointing Report for January – February, 2006, JPL internal document

band mission. For Ka-band on the 34-m antenna, a special pointing model tailored specifically to the two Ka-band missions – Cassini and MRO – enables the performance of blind pointing to be less than the required 4 mdeg.††

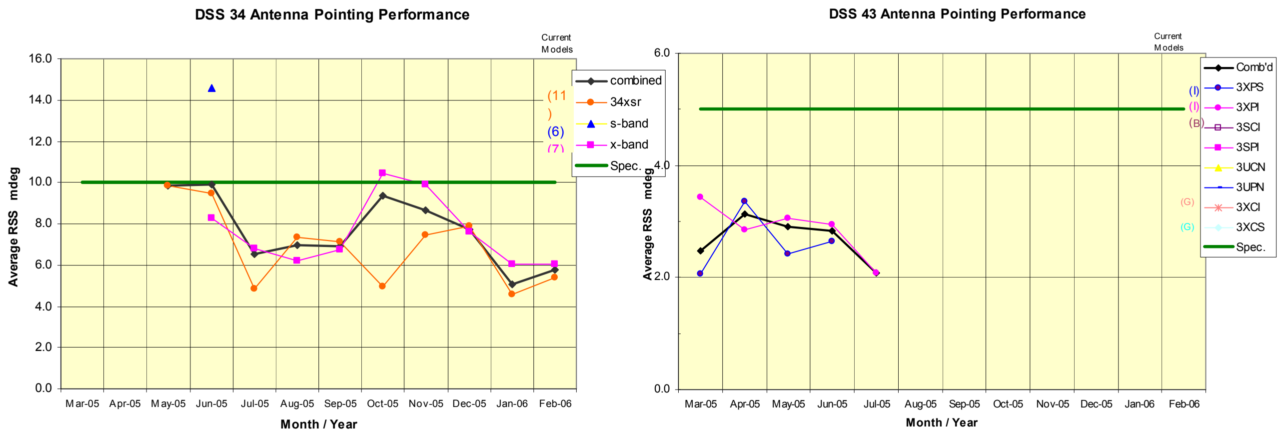

Figure 10. Pointing performance of the 34-m and 70-m antenna at Canberra Tracking Complex

B. System Noise Temperature

Data on system noise temperature are extracted from the link monitor logs. The values vary as a function of tracking elevation. To simplify the comparison against the model, only SNT of high elevation (beyond 80-degree) are extracted since the dependency on elevation becomes minimum at that point. Since the noise temperature is also dependent on the microwave configuration, e.g. diplexed (to allow concurrent uplink and downlink) versus listen only (downlink only), maser versus HEMT (high efficient mobility transistor), proper grouping of data is also necessary. A sample of actual noise temperature, averaged for a given configuration, for individual antennas is shown in Figure 11.

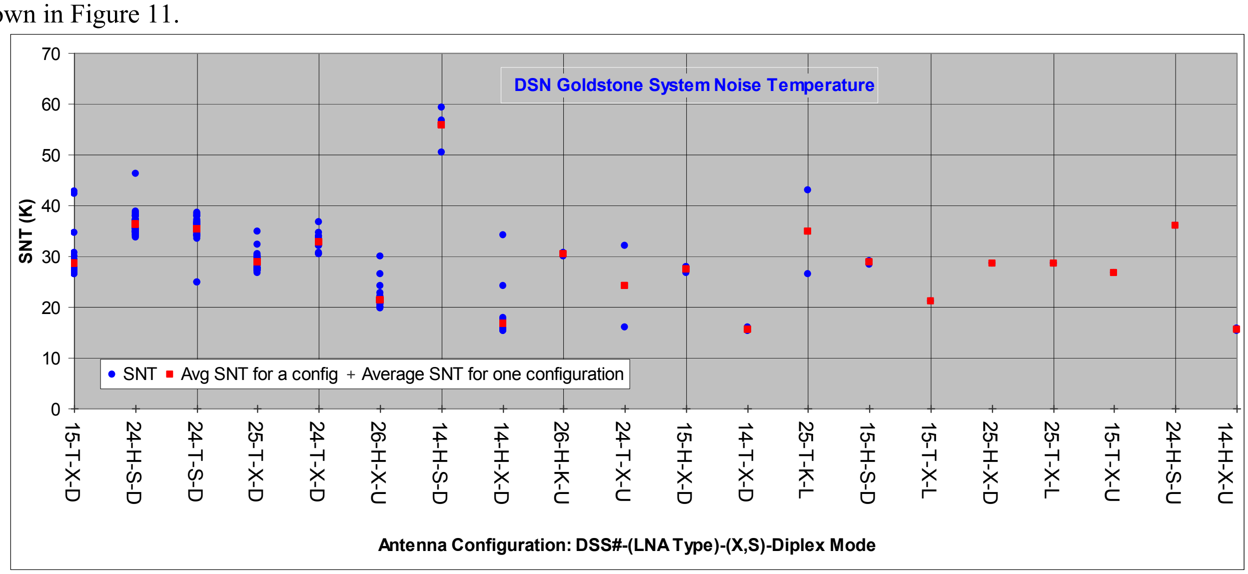

Figure 11. Observed SNT at Canberra DSCC

C. Doppler Noise

Significant change in Doppler noise affects navigation accuracy. A significant deviation from the norm or from performance requirement implies equipment failures. The cause could be in the input frequency references or the equipment is being saturated which results in incorrect estimation of the signal frequency and phase. A sample of

†† DSS-34 Ka-band Delivery Review, JPL internal document.

observed Doppler noise on the received carrier is shown in Figure 12. Model performance, as a function of signal power, is also included.

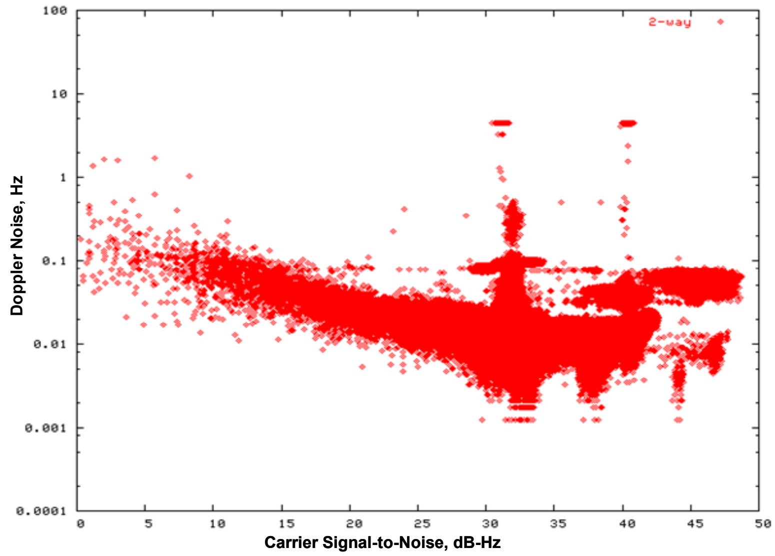

Figure 12. Doppler noise

D. Frequency and amplitude stability

Beside the telemetry, tracking and command functions, the DSN ground system is also used as a scientific instrument in radio science experiments. Gravitation wave search requires system with very stable long-term frequency stability, in the order of hours that correlates to the roundtrip light time of the experiment. In planetary ring or atmospheric occultation study, physical characteristics of the ring or atmosphere are deduced from the amount of attenuation the signal gets exposed as it travels behind the object of study. To do so, it requires a ground system of high amplitude stability.

Figure 13 shows the actual vs. modeled frequency stability. The model includes the effects of ground system, space media, and spacecraft equipment. The space media effects include solar scintillation, Earth troposphere and ionosphere.

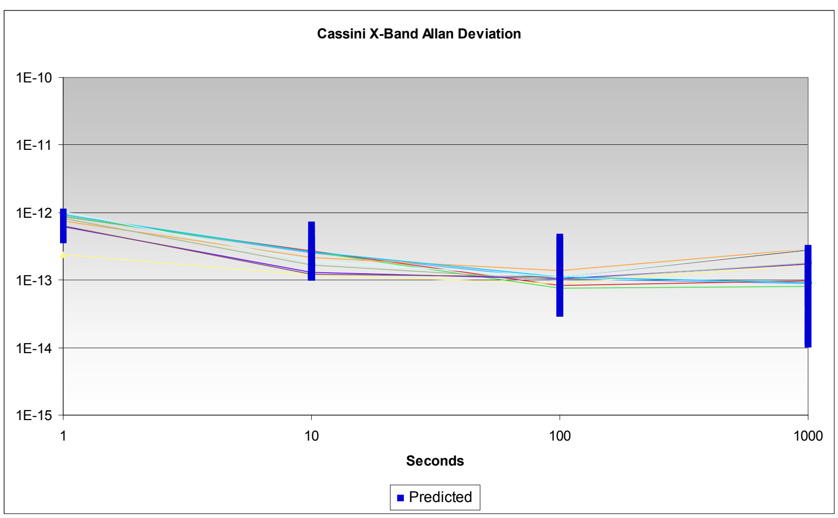

Figure 13. Frequency stability between observed and model

Figure 14 shows the amplitude stability observed at one of the spacecraft track at one of the 34-m Beam wave guide antennas. The data observed stability at X-band is close to the 0.2 dB specification. However, it should be noted that the observed data include the effect of space media and spacecraft, in addition to the ground system. Thus, it is reasonable to expect that the observation is more than what is required for the ground system.

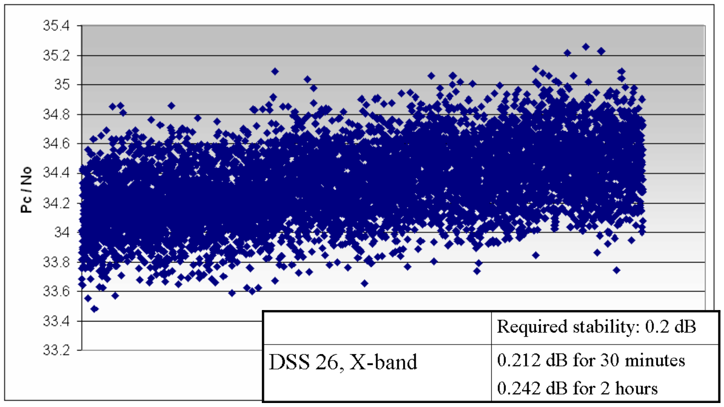

Figure 14. Amplitude stability

E. Wide Area Network Loading

In today configuration, the DSN facilities are interconnected by a set of dedicated connections. The connection is made of a few T1/E1 circuits. The bandwidth is limited, about 4 Mbps for overseas tracking complexes and 5.5 Mbps for Goldstone tracking complex. It is important to monitor the network loading to gain an understanding on how heavy the traffic is and how much margin remains. This information is also important for the validation of the network traffic model. Together, the data and the model allows for proper planning on expansion for future need.

A sample of the network monitoring is provided in Figure TBD. The data are grouped under various traffic types; each has a different latency requirement. As shown, the average consumption is 40 % of total capacity (2.1 Mbps consumption for Goldstone). The percentage of time when the load exceeds 80% of capacity is ~30% for Goldstone[‡‡]. Such utilization is less for Canberra and Madrid Complexes.

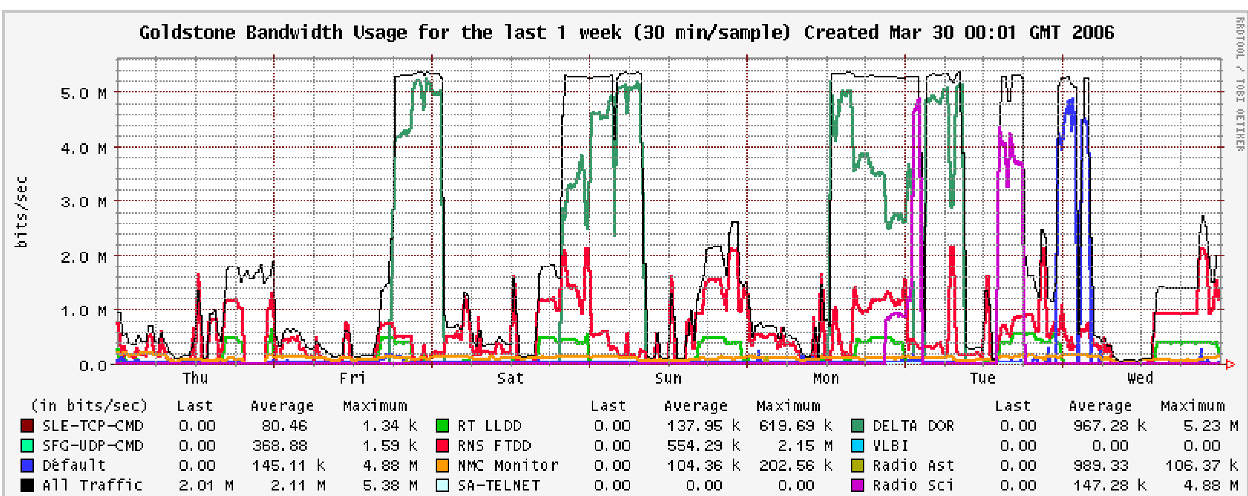

Figure 15. Network Utilization

F. Time and Frequency Reference Synchronization

It is critical to maintain synchronization of the frequency and timing references at each DSCC to a standard reference of NIST. This is to ensure the accuracy of data time tags as well as navigation application. Figure 16 shows a sample of metrics being monitored and assessed.

[‡‡] J. Liao, Assessment of WAN saturation, JPL internal communications, March 2006.

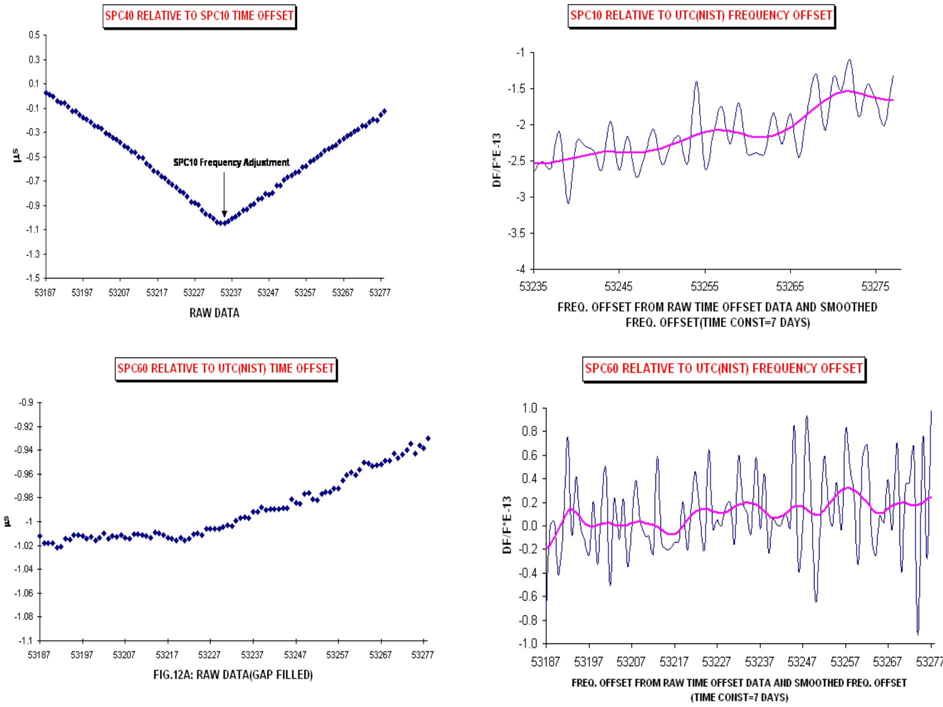

Figure 16. Frequency and timing synchronization offset

G. Link Setup Time

Current DSN scheduling allows for a fixed time to setup certain services. This setup involves configuring equipment into the link, going through some internal validation, moving antenna to point, warming up transmitter if needed. For telemetry service, 30 minutes is given. With concurrent command and tracking services which involves the uplink, the setup time is increased by 15 minutes.

The monitoring of actual setup time supports assessment on how well the DSN is executing the track preparation. It can help identify the bottleneck in the process. It also offers opportunity to increase the efficiency of network operation with appropriate reduction of setup time. Obviously, provisions need to be made to accommodate the exceptions where extra time is needed to allow operators work out the problem without being penalized for not meeting the requirement.

Figure 17 provides a sample of the setup performance for the month of January and February, 2006. This data set shows that, on the average, Goldstone Operations personnel got the link ready within 71% of the allocated time. 90% of the time, they finish the job within allocated time; leaving 10% being the exception.

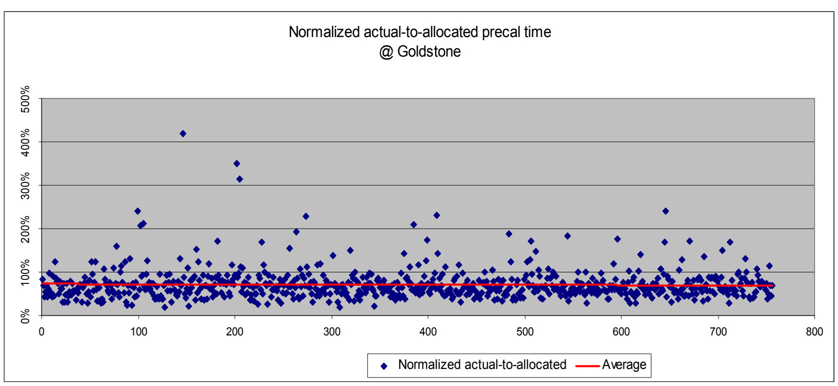

Figure 17. Pre-track setup time, Jan. - Feb. 2006

V. Reliability, Maintainability, Availability Assessment

The aspects of reliability, maintainability and availability (RMA) constitute a big part of effort from performance analysis team. Failures are examined for underlying causes so that corrective actions can be taken. Collaboration of various experts help brings about resolution to operational problems. Common or severe weak points in the system are identified. The resulting "top-ten" items are considered by program management for improvement.

Figure 18 shows a sample of non-availability breakdown to subsystem levels. Figure 19 shows a trending comparison of uplink equipments. This particular subsystem has progressively improved over the last few years, as reflected in the drop in the number of failure discrepancy reports.

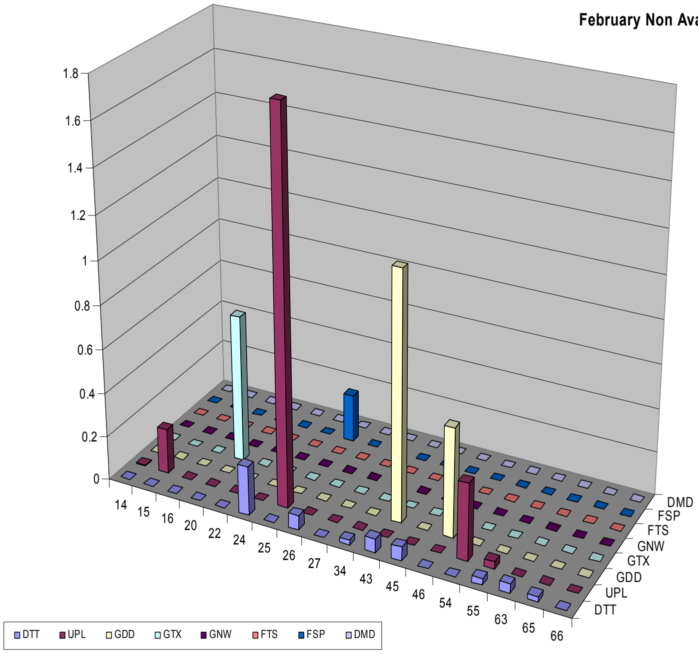

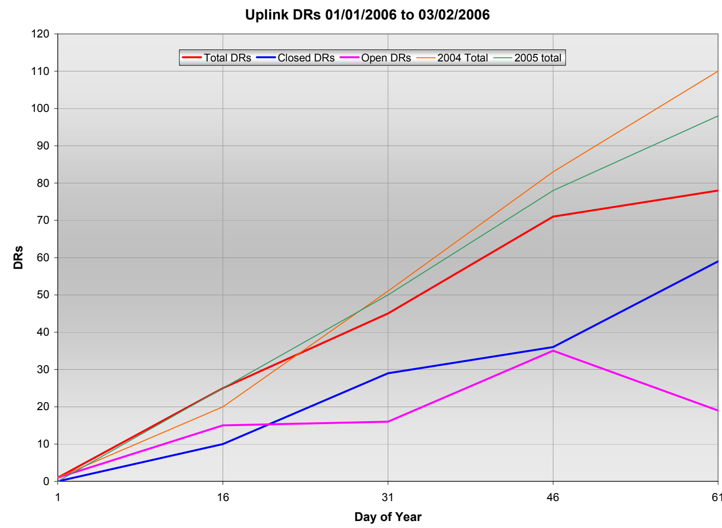

Figure 18. (Non-) availability statistics breakdown to subsystem level

Figure 19. Failure trending comparison for uplink equipment

VI. Conclusion

To recap, this paper presents the context of ongoing performance analysis activity being taking place in the Deep Space Network. The activity includes root cause analysis of failures which enables proper, effective solution; and performance assessment which establish if the requirements are being met operationally and the margin available. The assessment establishes a baseline performance on QQCL metrics as well as specific parameters that would impact services being provided. These include antenna pointing which becomes more important for operation at higher Ka-band frequency, system noise temperature, Doppler accuracy, bandwidth loading, offset in frequency and time references among the three tracking complexes, and the setup time for tracking passes.

Acknowledgments

The system performance assessment products presented therein are derived from the joint efforts by various engineering teams within the Deep Space Network. Special recognition is given to the following individuals for their leaderships in the areas of expertise: Jason Liao on network monitoring and data extraction that enable metrics calculation; Tim Cornish on uplink equipment; Andrew Odea on downlink telemetry processing; Dong Shin on tracking and telemetry QQCL processing; Mick Connally and Alina Bedrossian on science data processing; Arthur Freiley, John Cucchissi, Paul Cramer, James Bowen on antenna and microwave equipment; Mark Harris of Raytheon Australia and Tim Gregor of ITT for antenna pointing calibration; Ron Norman on service management processing; Robert Tjoelker on referenced frequency and timing; Russell Byrne and Michael Rafferty on ground communications aspects; John Luvalle of ITT for the general statistics on failure reports; Luis Huertas of ITT and Hossein Hosseini on power components.

.