The Impact of Traffic Prioritization on Deep Space Network Mission Traffic

Esther Jennings, John Seguí, Jay Gao, Loren Clare, Douglas Abraham

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
{Esther.H.Jennings, John.S.Segui, Jay.L.Gao, Loren.P.Clare, Douglas.S.Abraham}@jpl.nasa.gov

Abstract—A select number of missions supported by NASA's Deep Space Network (DSN) are demanding very high data rates. For example, the Kepler Mission was launched March 7, 2009 and at that time required the highest data rate of any NASA mission, with maximum rates of 4.33 Mb/s being provided via Ka band downlinks. The James Webb Space Telescope will require a maximum 28 Mb/s science downlink data rate also using Ka band links; as of this writing the launch is scheduled for a June 2014 launch. The Lunar Reconnaissance Orbiter, launched June 18, 2009, has demonstrated data rates at 100 Mb/s at lunar-Earth distances using NASA's Near Earth Network (NEN) and K-band. As further advances are made in high data rate space telecommunications, particularly with emerging optical systems, it is expected that large surges in demand on the supporting ground systems will ensue. A performance analysis of the impact of high variance in demand has been conducted using our Multi-mission Advanced Communications Hybrid Environment for Test and Evaluation (MACHETE) simulation tool. A comparison is made regarding the incorporation of Quality of Service (QoS) mechanisms and the resulting ground-to-ground Wide Area Network (WAN) bandwidth necessary to meet latency requirements across different user missions. It is shown that substantial reduction in WAN bandwidth may be realized through QoS techniques when low data rate users with low-latency needs are mixed with high data rate users having delay-tolerant traffic.[1] [2]

TABLE OF CONTENTS

1. INTRODUCTION

NASA's Deep Space Network (DSN) consists of three Deep Space Communications Complexes (DSCCs) positioned strategically across the Earth to enable continuous coverage of the wide range of space exploration missions. The DSCCs are located at Madrid, Spain; Canberra, Australia; and Goldstone in California, United States. Each GSCC consists of numerous antennas and associated signal processing resources for receiving and transmitting communications with spacecraft ranging from Earth orbit to the far reaches of the solar system.

The DSN provides services for both uplink (command) and return (telemetry) data communications, as well tracking services. It also provides the data transfer services to transport the data terrestrially between the DSCC and the end user Mission Operations Center. These data transfer services rely on the Wide Area Network (WAN) capabilities provided generally to all NASA enterprises, programs and centers by the NASA Integrated Services Network (NISN). The DSN also provides science services, using the unique capabilities of the DSCC resources (radio science, VLBI, radar science). In this paper, we focus on DSN return data services, focusing on the transfer services, which do not include radio science, VLBI and radar science.

Return data is transmitted from a spacecraft and collected at a DSCC, where it is then forwarded to the Deep Space Operations Center (DSOC) that provides further data distribution to the individual Mission Operations Center (MOC).

The DSN supports S-, X- and Ka-bands and provides frame, packet, and file services. The DSN is based on a service provision model using international standards.

Missions that are supported by the DSN are highly diversified. These variances are more extreme than in other space ground networks. For example, mission lifetimes may be very long, and distances can be extreme (e.g., Voyager 2 is 14B km from Earth, 12.8 hr one-way light time, 160 b/s return data rate). Pass durations may be minutes to hours.

---

[1] 978-1-4244-7351-9/11/$26.00 ©2011 IEEE.
[2] IEEEAC paper #1714, Version 5, Updated January 11, 2011

In particular, there is a wide variation in the offered data rates among the different missions that the DSN supports. The Joint Dark Energy Mission (JDEM) is anticipated to use a 150 Mb/s data rate, for example. In addition, return data from a given mission will comprise substantially different requirements on Quality of Service: a small percentage of the total downlink data will be time-critical, consisting of engineering health & status or quick-look information, and the remainder will be bulk science data that has a much more relaxed latency requirement. Also associated with these QoS differentiations are reliability, such that a status message is better to drop and be replaced with a more current value than to be retransmitted automatically by the underlying communications transport system. In this paper, we will focus on latency issues, with the understanding that mechanisms to implement differing reliability requirements (such as Automatic Repeat reQuest ARQ) may also be involved.

When all traffic is handled identically by the communications service provider, irrespective of the different latency requirements of distinct traffic classes, then sufficient resources must be provisioned to meet the requirement of the most stringent (lowest latency guarantee) class across the total offered load. On the other hand, the service provider may utilize QoS differentiation, in which service class is recognized by some means, and priority handling is provided, generally with higher priority given to the lower latency traffic type. By using QoS differentiation, a substantial reduction in total resources may be realized while satisfying the vector of QoS latency requirements for the traffic classes.

In this paper, we characterize the reduction in resources that may be achieved through the use of QoS differentiation in the context of the DSN. This reduction equates to lower bandwidth required for the WAN elements supporting return data transfer from the DSCCs to the DSOC.

The next section provides a brief overview of the inputs, model, and simulation tool used for the analysis. Section 3 presents the characterization of WAN bandwidth reduction achievable through the employment of QoS differentiation methods by the DSN service provider. Section 4 briefly describes further extensions of the work regarding joint consideration of the data distribution to the individual MOCs. Conclusions are presented in Section 5.

2. MODEL AND SIMULATION TOOL

2.1 Generation of the Input Data

Before describing the DSN return data transfer model and associated simulation tool, we first want to understand the source of the future mission traffic being applied to them. The missions of interest, their associated space-ground link characteristics, and their tracking requirements are derived from the Mission Set Analysis Tool (MSAT) – a tool which analyzes the communications properties of potential future DSN-supported missions appearing in NASA's Space Communications Mission Model (SCMM) [1]. The Orbital Trajectory Inference Engine (OTIE) is then used to develop visibility files for each mission relative to the DSN ground tracking stations. The link budget and tracking requirements from MSAT and the visibility files from OTIE are then fed into the DSN Simulator – a tool which, among other things, generates a simulated schedule for each mission's ground station contacts over the time frame of interest [2]. It is this simulated schedule, along with the associated mission data rate information, that serves as the input to the modeling and simulation that are the focus of this paper.

2.2 Description of the Model and Simulation Tool

The DSN return data transfer model is defined as follows. Numerous space missions are supported simultaneously, and are scheduled to use DSN resources. At the appropriately scheduled time, a spacecraft's transmitted signal will be received at a (typically one) ground tracking station, called a Deep Space Station, which is the antenna and control equipment and RF equipment. Several DSSs are located relatively closely together, and their signals are transferred to a Signal Processing Center (SPC), which performs digital processing. Thus a DSCC consists of the set of DSSs and the associated SPC for one of the three locations (Canberra, Madrid, or Goldstone). For the purposes of this paper, it is assumed that the SPC processing is performed to a common level (e.g., frame layer) for all spacecraft return data flows. Furthermore, the return data may undergo buffering at the SPC to whatever extent is allowed by the defined QoS latency requirement. That is, the mechanisms are in place for such store-and-forward buffering, and sufficient storage resources are available. It is noted that these assumptions may pose challenges for commonplace Internet (TCP/IP-based) equipment, due to large-scale buffers and latencies that may exceed conventional parameter settings. (Use of Delay/Disruption Tolerant Networking [3] provides one solution for this.)

Of importance to our problem is the aggregate of return data that is to be transferred from each DSCC to the DSOC. Therefore, details of the individual spacecraft, GSSs and each SPC are modeled as separate traffic sources entering a single DSCC entity that combines all associated return data that is offered to the DSCC site. That is, the superposition of all return traffic that is received at a given DSCC is used in the model, for all three DSCCs. For purposes of sizing the DSOC-to-MOC links, individual spacecraft mission data is modeled as de-multiplexed at the DSOC.

Figure 1 below depicts the model of the return data transfer from the DSCCs to the various MOCs via the DSOC. In this figure, an example of DSCC to DSOC path is shown in pink; an example of DSOC to MOC path is shown in blue. In our abstract model, we treat these paths as single links. The pink circles represent the DSCC data aggregation point and DSOC telemetry ingress point. Section 3 describes bandwidth sizing of the DSCC to DSOC paths (pink path). A QoS differentiation analysis is derived here. Section 4 extends the end-to-end path to include DSOC to MOC paths (blue path).

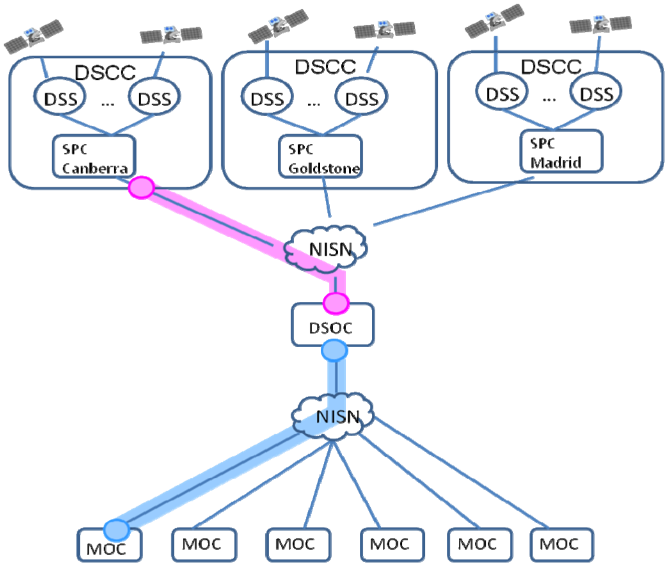

Figure 1 – Data Transfer Model

Our analysis is based on use of the space networking analysis tool MACHETE [4]. In general, this tool provides broad modeling capabilities for characterizing orbital and planetary motion kinematics, link engineering, and communications protocol behaviors. For the purposes of the investigation described in this paper, relatively coarse-grain abstraction of the problem may be used, allowing straightforward application of basic queueing models of store-and-forward operation at each DSCC. The core capability within this used is the QualNet® discrete-event simulation environment.

3. QoS DIFFERENTIATION GAIN

In our simulation scenarios, we generate high rate data traffic representing anticipated missions such as the JDEM mission with data rate as high as 150 Mbps and pass length can be as high as a couple of hours. Other representative mission data at lower rates are also generated for the simulation. The simulation scenarios cover a 10-day period in 2018. A schedule of communication is produced by Space Communications Mission Model Team's DSN mission set scheduling tool. For each mission, the schedule lists the start and end times of each pass and the communication data rate. This information is used as a basis for traffic generation in our simulations.

A representative 10-day total traffic load for one DSCC is presented in Figure 2, in b/s vs. time in days. The bursty nature of the traffic is apparent due to support for the occasional very high data rate passes. In this example, the sample coefficient of variation for the data set is found to be 2.98, indicating the high dispersion in the distribution.

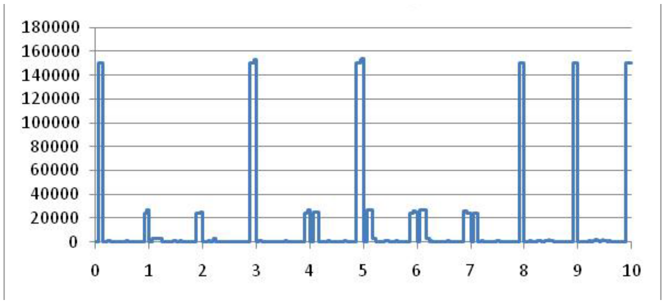

Figure 2 – Sample Traffic for DSCC

The purpose of this study is to characterize bandwidth savings when using QoS versus not using QoS for the DSCC to DSOC paths. From the scheduled communications, we assume that 10% of the data has a time-critical delivery requirement and the other 90% are delay-tolerant.

The network topology is represented as three DSCC paths to the DSOC. It is not our intent to model the NISN-provided DSCC-DSOC path in detail, thus, we currently map each path to a link in the simulation to determine the WAN bandwidth needed to meet various latency requirements of either prioritized or not prioritized data traffic. In addition to the latency requirement, it is also required that all data be delivered. In other words, the bandwidth must meet *both* the latency and no data loss requirements.

Since the majority of data is delay-tolerant, we need to allocate enough store-and-forward capacity at each DSCC to hold the high volume of data so that a lack of buffer space does not become the driver of bandwidth. The protocols used in the simulation are IP-based for the terrestrial WAN. In this study, the simulation tool models only the delay due to transmission, queuing and propagation, not protocol processing, since this is assumed to be comparatively negligible in this case.

Two scenarios were simulated: one with QoS prioritized data and one without QoS. For each experiment, we vary the latency requirement of the high priority data. Note that, when QoS is not used, we cannot distinguish between time-critical data and delay-tolerant data, so all data must meet the time-critical requirement. We ran simulations where the time-critical latency started at 10 seconds and becomes incrementally relaxed to 5,000 seconds (1 hour, 23 minutes and 20 seconds). The delay-tolerant latency requirement is assumed at 8 hour. Our simulation shows that as the latency requirement is relaxed, the needed bandwidth also decreases.

To determine the required bandwidth meeting both latency and no data loss requirements, we start with an initial bandwidth "guess", run the simulation, and check the produced statistics to see whether both requirements are met. We then use a binary search approach to select link bandwidths where the new estimate is dependent on the simulation outcome of previous iterations. This is iterated over a predefined maximum number of iterations. Then we select the minimum bandwidth among all the iterations meeting both requirements as our answer. Figure 3 shows this iterative process.

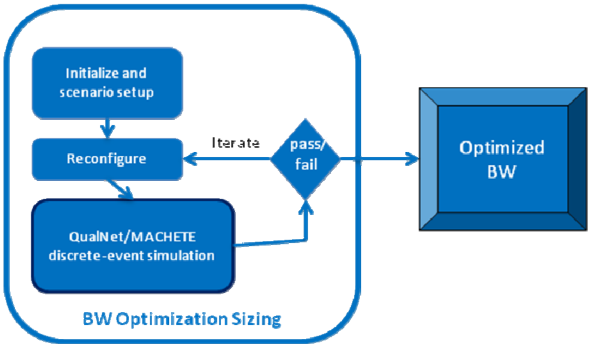

Figure 3 – Iterative Bandwidth Sizing Process

Each experiment for a specific latency requirement value is a set of QualNet/MACHETE simulations that iteratively finds the optimal bandwidths.

Figure 4 shows bandwidth savings on the Goldstone DSCC to DSOC path. The x-axis lists the different latency requirements used for the time-critical data (in seconds). The y-axis is the bandwidth needed to meet both latency and no data loss requirements.

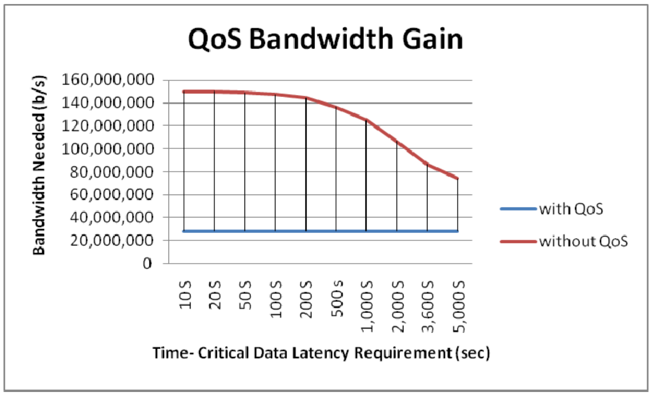

Figure 4 – Simulation Result on Required Bandwidth from DSS to DSOC

From the simulation, we observe a factor of 3 to 5 of bandwidth savings when QoS is deployed. The highest data rate in the mission traffic set is 150 Mbps, so approximately 15 Mbps of data are time-critical and may be considered a trivial lower bound for the bandwidth. Currently, our simulation shows that a bandwidth of 28 Mbps is sufficient to meet latency requirement with no data loss when QoS mechanisms are employed. The delay for time-critical data is less than 5 seconds and for delay-tolerant data is less than 6 hours. This indicates that we may be able to reduce the bandwidth further while still meeting the latency requirement. However, the driver to the 28 Mbps result was due to storage constraints; data loss was observed if bandwidth was reduced below 28 Mbps. The trade-offs among storage and latency requirements need to be further investigated.

4. MOC DATA DISTRIBUTION INTEGRATION

In the previous section, we described experiments focusing on the paths from DSCC to DSOC. We can extend this study to include the latencies incurred on the paths from the DSOC to various MOCs. This extension is shown as the blue path in Figure 1.

One of the key issues in this extended scenario is that the MOCs have different distances from the DSOC, therefore different WAN service costs. Bandwidth optimization that combines both the DSCC-DSOC links as well as the DSOC-MOC links must account for fairness issues, since the end-to-end latencies will vary due to the different DSOC-MOC distances and differences among user mission traffic to individual MOCs.

Our approach to the problem is to first optimize the DSCC to DSOC bandwidth with certain initial assumptions on the DSOC to MOC links. Then using the result of the bandwidth optimization sizing from DSCC to DSOC, apply the same iterative bandwidth optimization sizing to the DSOC to MOC links. While extending the end-to-end paths to include MOCs, there is additional store-and-forward at the DSOC, and the latency budget is levied across both hops. Alternatively, one can begin by fixing the DSOC to MOC bandwidth to expected obtainable service capacity for each MOC, and perform the bandwidth optimization on just the DSCC to DSOC side.

In addressing the fairness issue, the variation in the ratio of high to low priority traffic classes and DSOC to MOC distance should be factored into the latency allocation. Therefore, additional constraints can be introduced by weighting the latency allocation on the DSOC to MOC segment based on its distance and high priority traffic loading.

We ran experiments using the approach to first determine the bandwidths needed for DSCC to DSOC and then determined the bandwidths needed from DSOC to MOCs. The same iterative method was used to adjust bandwidth of each link for the next run according to the result of the previous run. We used representative mission data traffic generated by SCMM as input to the network simulation. The data traffic covers a 31-day period in 2018. Figure 5 shows the bandwidth required from the DSOC to each MOC, while using QoS versus not using QoS.

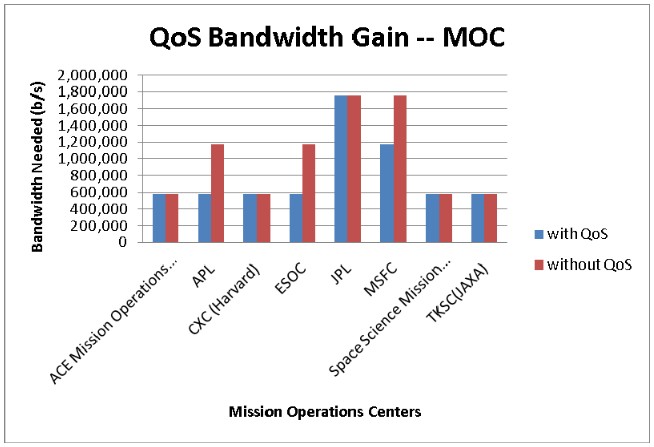

Figure 5 – Simulation Result on Required Bandwidth from DSOC to MOC

An initial set of results showed that we may achieve a factor of 2 bandwidth savings while using QoS for some of the links between the DSOC and MOCs. The experiment results also confirmed the intuition that the use of QoS will only improve on bandwidth required. It should be noted here that due to the large trade space of the problem, simulations runs were restricted to coarser granularity based on a finite set of bandwidth values. Therefore one may see instances where the simulation predicted identical bandwidth requirements with and without QoS, although we expect finer resolution analysis may reveal use of QoS would require lower bandwidth. For example the bandwidth requirement predicted for ACE Mission Ops is 586Mbps, with or without QoS. While the high priority traffic latency without QoS is 7.56sec, which meets the 10sec latency requirement, the maximum latency for high priority traffic with QoS is only 1sec, indicating the potential for additional bandwidth saving. However, the amount of bandwidth saving is below the granularity of the simulation so the simulation stops at this value.

5. CONCLUSIONS

The advantages of incorporating QoS mechanisms have been quantified in terms of the reduction of WAN bandwidth required. The NASA Deep Space Network (DSN) was used as the illustrative case from which the results were drawn. Realistic models of anticipated DSN mission traffic were used in deriving the results in terms of data rates and variability arising from scheduling highly subscribed DSN resources and spacecraft view opportunities. Parametric variation of user QoS latency values over different traffic classes served to illustrate the range of possible performance.

Our study shows that substantial bandwidth savings is achievable by introducing a basic prioritization mechanism into the data transfer service. The benefit of the QoS mechanism mainly arises from its ability to exploit the trade between storage and bandwidth in order to maximize transmission capacity on the high priority traffic during periods of traffic bursts.

It is observed that without a QoS capability, the system is unable to differentiate between data that could be deferred for later transmission versus data that requires immediate forwarding, therefore driving bandwidth requirement to the peak rate in order to meet the most stringent latency. This is particularly detrimental to cost containment for a network with a bursty traffic profile because the cost of bandwidth is much higher compared to storage.

6. ACKNOWLEDGEMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology and the National Aeronautics and Space Administration.

REFERENCES

[1] B. E. MacNeal, D. S. Abraham, R. C. Hastrup, J. P. Wu, R. J. Machuzak, D. P. Heckman, R. J. Cesarone, R. P. Tikidjian, K. Tran, “Mission Set Analysis Tool for Assessing Future Demands on NASA’s Deep Space Network”, IEEE Aerospace Conference paper 1052, 2009.

[2] Doug Abraham, “Overview: Mission Model & Architecture Analysis,” Presentation to the SCaN Program System Engineering Team (Pasadena: JPL, 11 May 2010), pp. 15-21.

[3] RFC 4838, Delay-Tolerant Networking Architecture, V. Cerf et. al., April 2007.

[4] Esther H. Jennings, John S. Segui, and Simon Woo, “MACHETE: Environment for Space Networking Evaluation,” AIAA International Conference on Space Operations (SpaceOps), Huntsville, AL, April 2010.

BIOGRAPHY

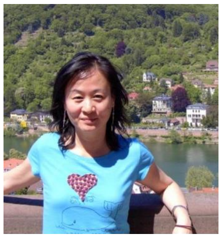

Esther Jennings *is a research staff with the Communication Systems and Research Section of the Jet Propulsion Laboratory. She received a Ph.D. in Computer Science from Luleå University of Technology, Sweden. She was a postdoctoral fellow at the Industrial Engineering Department at Technion, Israel Institute of Technology. Prior to joining JPL, she was an assistant professor at the Computer Science Department of California State Polytechnic University, Pomona. Her research interests are in distributed graph algorithms, reliable multicast protocols, energy-efficient algorithms for wireless network and algorithm simulation.*

John S. Seguí *is a member of the Communication Networks Group at the Jet Propulsion Laboratory, managed for NASA by the California Institution of Technology, Pasadena, California. He works on network architectures and protocol evaluation for the Interplanetary Network Directorate (IND), Space Communication and Navigation (SCaN) organization*

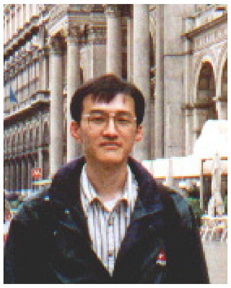

Jay Gao *joined the Jet Propulsion Laboratory in 2001 and is currently a senior research staff in the Communications Networks Group in the Telecommunication Research and Architecture section. His research is primarily focused on space-based wireless communications and networking, with emphasis on quality-of-service, demand access, and multilink technologies for the envisioned Interplanetary Network (IPN) and optimization and protocols for deep space Ka-band communications. He also supports requirements definition and interface design activities for the Department of Defense's Transformational Communications MilSatcom project and system engineering effort for NASA's Exploration System and Mission Directorate (ESMD), supporting the Constellation Program for return of human to the Moon and Mars. Other research interests include sensorweb, distributed communication/sensor systems, energy efficient routing and self-organization algorithm for cooperative signal processing and sensor networks. He received his B.S., M.S., and Ph.D. degree in Electrical Engineering from UCLA in 1993, 1995, and 2000, respectively.*

Loren Clare *is the supervisor for the Communications Networks Group at the Jet Propulsion Laboratory. He obtained the Ph.D. in System Science from the University of California, Los Angeles in 1983. His interests include wireless communications protocols, self-organizing systems, network systems design, modeling and analysis, and distributed control systems. Prior to joining JPL in May 2000, he was a senior research scientist at the Rockwell Science Center, where he acquired experience in distributed sensor networks, satellite networking, and communications protocols for realtime networks supporting industrial automation.*

Doug Abraham *is a Senior Systems Engineer within the Architecture, Strategic Planning, and Systems Engineering Office of JPL's Interplanetary Network Directorate. In this capacity, he oversees efforts to forecast future mission customer requirements and trends, assesses their implications for Deep Space Network evolution, and assists in the development of the roadmaps and plans needed to guide this evolution. Doug also supports NASA HQ-led study activities pertaining to future space communications architectures.*

Prior to his current assignment, Doug worked on the Galileo, Ulysses, and Cassini flight projects. He has also worked on several pre-project formulation activities, including Pluto Fast Flyby and the "Fire and Ice" collaboration with Russia.

Doug began his career as a graduate student intern in the International Space Station Program Office (1988). He graduated Magna Cum Laude from Texas A&M University in Physics (1986) and earned an M.S. in Technology and Science Policy, with specialization in technology assessment and electrical engineering, from Georgia Tech (1990).