A Model to Assess the Mars Telecommunications Network Relay Robustness

André R. Girerd, Leila Meshkat, Charles D. Edwards, Jr., Charles H. Lee
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109, USA
(e-mail: andre.r.girerd@jpl.nasa.gov)

Abstract

The relatively long mission durations and compatible radio protocols of current and projected Mars orbiters have enabled the gradual development of a heterogeneous constellation providing proximity communication services for surface assets. The current and forecasted capability of this evolving network has reached the point that designers of future surface missions consider complete dependence on it. Such designers, along with those architecting network requirements, have a need to understand the robustness of projected communication service. A model has been created to identify the robustness of the Mars Network as a function of surface location and time. Due to the decade-plus time horizon considered, the network will evolve, with emerging productive nodes and nodes that cease or fail to contribute. The model is a flexible framework to holistically process node information into measures of capability robustness that can be visualized for maximum understanding. Outputs from JPL's Telecom Orbit Analysis Simulation Tool (TOAST) provide global telecom performance parameters for current and projected orbiters. Probabilistic estimates of orbiter fuel life are derived from orbit keeping burn rates, forecasted maneuver tasking, and anomaly resolution budgets. Orbiter reliability is estimated probabilistically. A flexible scheduling framework accommodates the projected mission queue as well as potential alterations.

1. Introduction

Once successfully in orbit around the red planet, current and projected Mars orbiters have demonstrated or are expected to attain relatively long operational durations. Evidence has shown that interplanetary spacecraft engineered to achieve a certain design life can often exceed that minimum, resulting in a functioning telecommunications node essentially limited by parts wearout, operational error, consumables, and funding. The combination of long spacecraft lifetimes and the aggregated international commitment over much of the past decade to utilize each Mars transfer opportunity has resulted in a positive net accumulation of functional assets over time. This gradual accretion of orbiting assets has resulted in a heterogeneous constellation known as the Mars Network.

The Mars Network acts as a communications intermediary between the Earth and individual missions on the surface (or near surface) of Mars. It allows landed missions to use short-range proximity link communications between Mars surface and Mars orbit, in addition to or instead of long-range direct-to-earth (DTE) communications. The existing elements of the Mars Network provide proximity communications in the UHF band, although with varying capabilities based on their technology and orbit. NASA's Mars program telecommunications strategy calls for each Mars orbiter to carry the Electra proximity payload, an advanced software radio designed specifically for Mars Network duty [1]. The Mars Reconnaissance Orbiter (MRO) is the first spacecraft carrying the Electra.

For many Mars surface missions the relay capability offered by the Mars Network compares favorably against DTE. Advantages include higher data rates at less power, reduced energy per bit, simplified or eliminated antenna pointing requirements, and reduced telecommunications subsystem mass and volume. Additionally, the network can provide contact opportunities for Mars locations at times when no DTE is possible. The current and forecasted capability of this evolving network has reached the point that designers of future surface missions consider complete dependence on it. Such designers, along with those architecting network requirements, have a need to understand the robustness of current and projected communication service.

A model has been created to identify the confidence level of Mars Network capability as a function of surface location and time. As currently implemented, the model is suitable for high-level Mars Program Office studies of a strategic nature.

2. Framework

Due to the decade-plus time horizon considered, the Mars Network will evolve, with emerging productive nodes and nodes that cease or fail to contribute. For the purposes of this discussion, a node is defined simply as any single Mars orbiter that provides proximity communications service. The model is a flexible framework designed to

- accept new node data as it becomes available
- manipulate existing node data under varying assumptions
- process node information into measures of robustness that can be visualized for maximum understanding

Each node provides a unique contribution to the Mars Network capability that varies as a function of its technology and orbit. Telecom performance information for five well-defined nodes is currently incorporated in the model. These nodes are the three existing Mars Network spacecraft (Mars Global Surveyor (MGS), Mars Odyssey (ODY), and Mars Express (MEX)), one pending arrival (MRO), and one cancelled candidate (Mars Telecommunications Orbiter (MTO)). In addition to its existing stockpile of node data, the model can accept new data for fresh nodes as they become defined, or merely leverage existing node performance data as reasonable approximations for future spacecraft that share similar proximity telecommunications capability.

Even for a Mars Network scenario with an agreed-upon node constituency, the actual network capability is dependent on a variety of assumptions. These include launch and arrival dates for nodes not currently on-orbit, decommissioning dates for functioning nodes, hypothetical loss-of-mission events, and telecommunications-related choices such as adaptive vs. fixed rate. The framework allows such assumptions to be changed quickly for rapid what-if analyses.

The model emphasizes the analysis of the Mars Network as a system, and integrates constituent individual node data into system-level robustness metrics. After processing, the framework offers two ways to visualize the resulting information for maximum understanding. The first way lets the user change assumptions and view the system result in a series of two dimensional plots with negligible computational delay. This allows manual exploration of the multidimensional tradespace a slice at a time. The second way uses a comprehensive batch mode to offer more simultaneously viewable dimensions. The resulting multi-dimensional map provides an intuitive method to understand Mars Network robustness in a single rich visual, although with fixed assumptions and significantly increased computation time.

3. Components of the Analysis

To generate robustness information, the model integrates data from several sources. The three classes of necessary data, referred to as the components of the analysis, are proximity link performance, probabilistic estimates of orbiter fuel expenditure, and orbiter reliability. The first component concerns the capability each node can provide, while the remaining components address the likelihood of nodes being functional.

Proximity Link Performance

There are a number of important assumptions about the manner in which telecommunications performance is handled. For the purposes of the model, Mars Network capability analysis is focused exclusively upon proximity operations, since for the vast majority of conceivable cases the proximity link will be the communications bottleneck rather than the Mars-to-Earth link. Therefore, the model addresses UHF performance only since that frequency band remains the current and forseeable proximity standard. While it is possible that future surface assets might consider using X-band for proximity operations, only the cancelled MTO would have supported it. Nevertheless, the Electra proximity payload can accept enhancements to allow X-band functionality for future orbiters, should the demand arise. Another architectural fact of the current and projected Mars Network is that each orbiter's relay function is independent of any other orbiter. Not having cross-links allows individual bent-pipe relays to be analyzed in parallel which simplifies the combinatorics.

Outputs from JPL's Telecom Orbit Analysis Simulation Tool (TOAST) provide global proximity link performance parameters for current and projected orbiters around Mars. TOAST outputs are particularly useful for broad, high-level robustness studies because they come in the form of global maps instead of stand-alone link analyses for single planetary locations [2]. Several assumptions need to be explained to understand the use of TOAST for Mars Network predictions over long timeframes. The general approach focuses on the capability the Mars Network can supply rather than the demand of future-surface assets. This is due to high uncertainty in the surface mission lineup after 2010, and lower uncertainty in the relay performance characteristics of existing, forthcoming, and potential Mars orbiters. Additionally, the global coverage maps of TOAST obviate the need for exact surface location knowledge of any landed asset, and detailed information on orbiter relay hardware, orbit, and operational concept either already exists or can be readily modeled. As an example, Figure 1 shows the antenna gain pattern used for MRO. TOAST uses similarly detailed representative gain patterns for the other nodes of the Mars Network. Nevertheless, some assumptions have to be made for the surface end of the proximity link. Future landed asset UHF capabilities are split into two classes – large assets and small (energy-starved) assets. Both classes use the same helix antenna pattern and 8.5W transmitter power. The large class, having ample energy reserves, can transmit while in view of an orbiter, while the small class is limited to the best 10 minute opportunity. An example of the large class might be the Mars Science Laboratory, while a network of distributed seismology sensors or ESA's Beagle-2 Lander would fall into the small class.

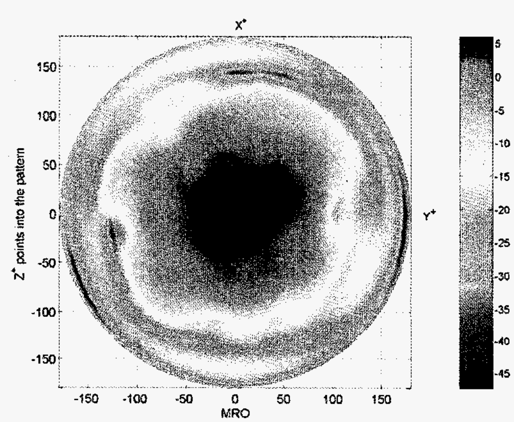

Figure 1: MRO UHF Antenna Gain

In addition to the selection of landed asset class, there are other assumptions governing TOAST output that can be varied within the model framework. An option exists to allow Electra-capable orbiters to use adaptive rates (in which the orbiter and lander continuously negotiate and adjust the data rate over the course of a communications session to take full advantage of time-varying channel conditions), while all other orbiters use their best single rate. In the case where future surface missions forego adaptive rate compatibility, all nodes can be forced to use best single rate. Multiple candidate orbits investigated by the MTO team are available to select from, offering a window into some basic choices for a non-polar orbiting dedicated telesat node.

TOAST provides many performance metrics for a given node, such as total UHF data volume per sol, average number of contacts, maximum communication gap, and maximum supportable data rate. Of these metrics, total daily UHF data volume is both the most meaningful and the most amenable for aggregation into a system-level performance metric for the Mars Network over an extended timescale. For example, deriving gap time metrics for any future system would be highly dependent on knowledge of relative orbit phasing between nodes, which is unlikely to be available for anything other than the current constellation and furthermore always subject to change. In contrast, the calculation of system data volume is nearly insensitive to any factor not already handled by the model. Figures 2 and 3 show global maps of UHF data volume for two nodes – MTO with a circular sun-synchronous orbit, and MRO in its polar orbit.. Notice that the polar orbiting MRO has maximum performance at the high latitudes, whereas MTO in a sun-synchronous orbit dominates the equatorial and mid-latitudes.

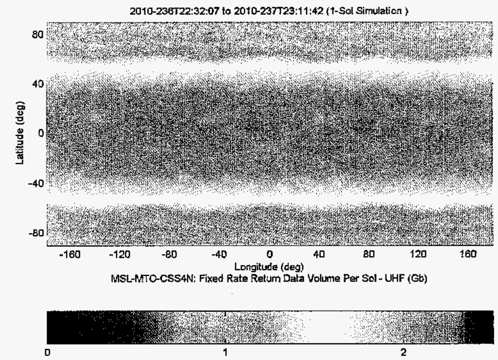

Figure 2: Global Map of UHF Data Volume of MTO in a Circular Sun-Synchronous Orbit

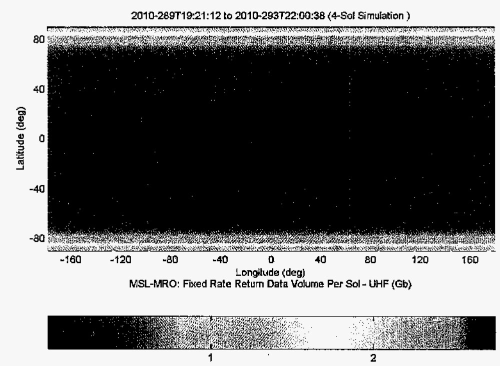

Figure 3: Global Map of UHF Data Volume of MRO in a Polar Orbit

Since latitudinal variations in performance are more significant and meaningful than longitudinal variations over extended time frames, TOAST performance metrics as a function of latitude and longitude are projected to a simpler function of latitude, for either the minimum, maximum, or mean performance across longitude. Figures 4 and 5 show total UHF data volume vs latitude for the same two nodes as Figures 2 and 3, respectively.

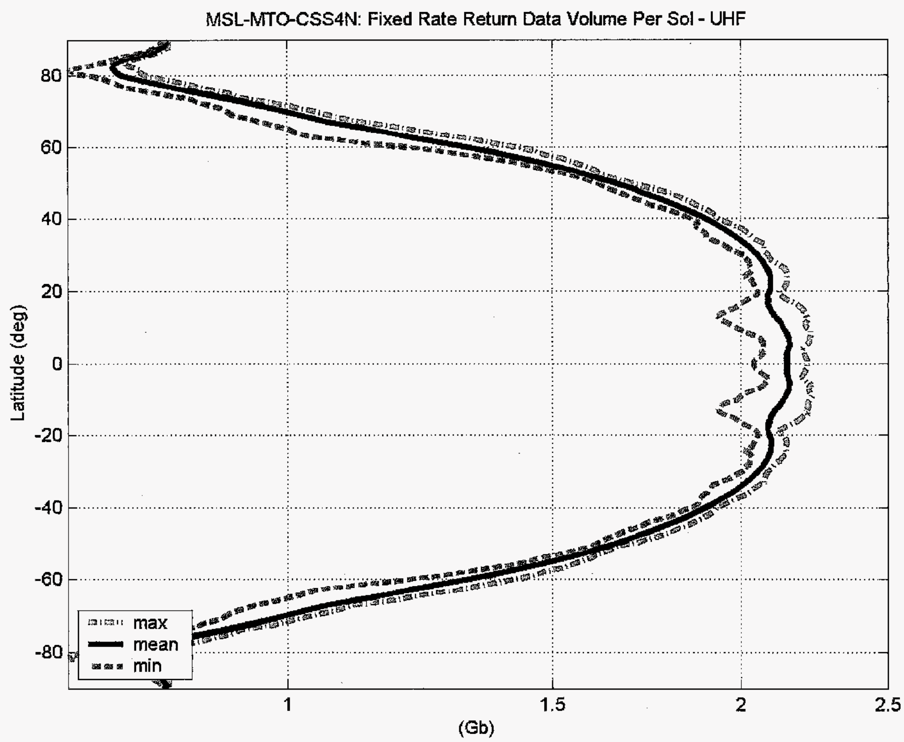

Figure 4: Latitudinal Distribution of UHF Data Volume for MTO in a Circular Sun-Sychronous Orbit

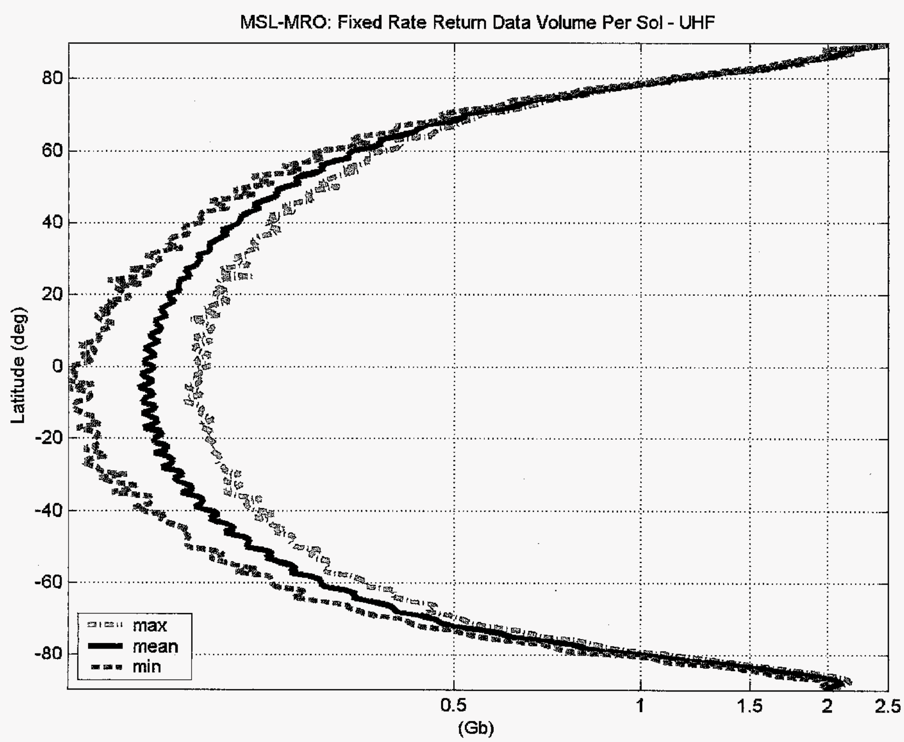

Figure 5: Latitudinal Distribution of UHF Data Volume for MRO in a Polar Orbit

Probabilistic Estimates of Orbiter Fuel Expenditure

Running out of fuel means the end of an otherwise functional node. For the known and a few hypothetical nodes, the model builds a probabilistic estimate of this eventuality's day of reckoning by examining three contributory factors that dominate fuel expenditure. The first factor is the steady state fuel burn rate, handled in a deterministic manner. This factor varies from spacecraft to spacecraft, and is influenced by orbit altitude, vehicle mass, and operations profile. It is dominated by momentum wheel desaturations and any necessary stationkeeping. The second factor is the fuel consumption allocated for planned targeting maneuvers. Not to be confused with science-oriented targeting maneuvers, this factor addresses a programmatic desire to adjust the orbit phasing of a relay node to provide critical event coverage (e.g. Entry descent and landing or Mars orbit insertion) for newly arriving Mars probes every 26 month Mars transfer opportunity. Although the actual expenditure is a function of the degree of phase shift required (expensive plane changes are not considered) and the amount of lead time allowed [3], both likely unknowns for future missions, the periodic nature of such maneuvers allows them to be treated in a deterministic though conservative manner. The third factor is the fuel usage for anomalies, handled in a probabilistic manner. This is composed of two parts, the distribution of fuel usage per anomaly, and the distribution of anomaly frequency. Calibrated against the knowledge of remaining fuel for current orbiters and the predicted remaining fuel (after reaching Mars operations status) for future orbiters, the three factors allow an estimate of the probability of having sufficient remaining fuel versus time. Figure 6 shows the probability distribution function of the expected fuel depletion date for a representative, though unspecified node. The presence of the steady state and periodic planned targeting maneuvers causes a noticeable shift in the curve's shape away from symmetry. Figure 7 shows the associated cumulative distribution function as the probability of running out of fuel, and its inverse the fuel confidence, or probability of having sufficient remaining fuel.

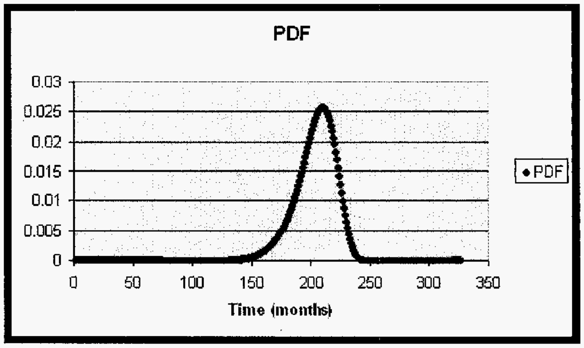

Figure 2: Representative Probability Distribution of Expected Fuel Depletion Date

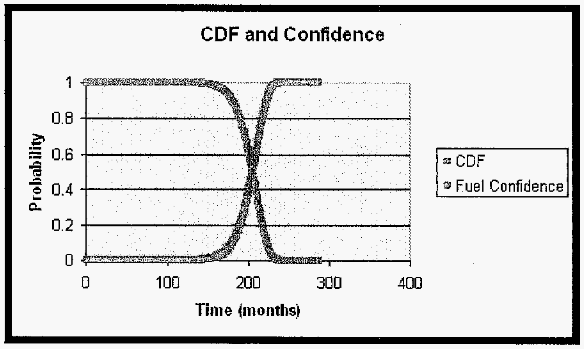

Figure 3: Cumulative Distribution Function for Fuel Depletion and Fuel Confidence

Orbiter Reliability

In order to understand the robustness of the Mars Network as a system, it is necessary to have data on the reliability of each constituent spacecraft node over time. To this purpose, the model uses time-based probabilities of node functionality. A paper by Meshkat et. al [4] describes the probabilistic risk assessment techniques used to determine node reliability in more detail.

Spacecraft reliability information was generated for the three existing nodes of the Mars Network and four future candidates. These candidates are MRO, MTO, MTO2, and a generic single-string Mars Scout orbiter. For each node a system schematic was constructed capturing key information about functional relationships and system architecture, such as the degree of redundancy and cross-strapping for critical subsystem components. For existing nodes, the known current spacecraft state was used for initialization and future reliability estimated using expert opinion about failure rates and the most updated consumable information for life-limiting components. The technique for future nodes was similar but incorporated the additional uncertainty of successfully passing through upcoming critical event gates such as launch, cruise, and Mars orbit insertion. As with the unique system schematics, the nature and sequence of the critical events were tailored to the expected mission profiles of the future nodes with considerable effort spent to be as representative as possible. For example, MRO included an aerobraking event while MTO included some potentially risky experimental activities.

The probabilistic reliability calculations were focused on events and phase-dependent system configurations leading up to and including relay operation. Anything unrelated to the critical path or unnecessary for Mars Network duty was left unmodeled. An example of a subsystem willingly omitted is the science payload of MRO. Because the failure behavior of the systems considered is related to the sequence in which events occur, dynamic fault trees were used. In addition, because of the multi-phase nature of future nodes, with phase-dependent system configurations and inter-phase dependencies, a combinatorial and Markovian approach was used to solve the classic phased-mission system.

4. Combinatorial Modeling Method

To generate system-level metrics for the Mars Network, the component data streams of proximity link performance, probability of sufficient fuel, and orbiter reliability have to be combined for each contributing node in an appropriately time-based manner. The probabilistic availability of each node, that is, the likelihood of the node functioning as intended, is merely the combined probability of the node's reliability and its probability of sufficient fuel. Unlaunched or decommissioned nodes have zero availability. The composition of the Mars Network will change with time in both the number and type of available nodes. A given state of the network, therefore, is a unique set of functioning nodes and unavailable nodes. Figure 8 shows part of a state tree for a six-node-maximum network, revealing 32 states. The probability of each network state is found from the combined probability of the required node conditions.

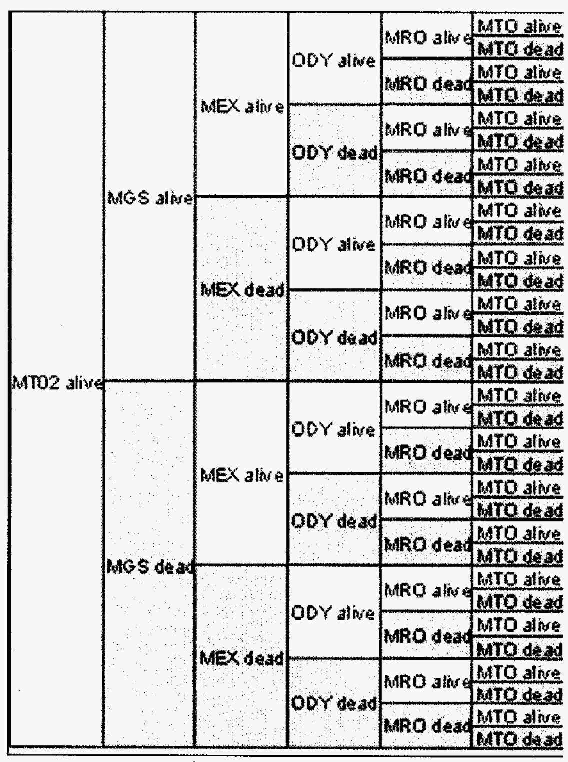

Figure 4: Network State Tree

The proximity relay capability of any network state depends on the performance offered by its constituent available nodes. This network state capability varies by latitude, since its node contributions do. For a given latitude, then, the question arises how to combine the individual node capabilities. With respect to data volume, the node capabilities are additive for the large asset class. This method is considered sufficient for the high-level strategic intent of the model, although it may overestimate the network state capability slightly by ignoring transmission exclusions for simultaneous view periods. For the energy-starved small asset class, the network state data volume capability is determined solely from the node that offers the highest performance ten minute pass.

Since the exact future state of the network is unknown over the course of time, there will be a multitude of potential state capability levels each with an associated probability. Therefore, for each latitude and point in time, the expected network capability will be distributed with respect to probability. Individual probabilities of network states that meet or exceed a capability floor are combined into a total probability of meeting that capability. A similar sorting approach is used to determine expected capability given a certain probability threshold.

5. Output

Perhaps the simplest measure of Mars Network robustness is the expected number of orbiters over time. Figure 9 shows the expected network population for a hypothetical scenario with MTO-class orbiters arriving in 2010 and 2015. For this graphic and all subsequent ones throughout this paper, simplified orbiter availability data was used as a substitute for the still-evolving expert-driven reliability and fuel consumption models. Therefore, the data displayed in the following figures should not be taken as representative of actual expectations of Mars Network robustness, but as examples of the types of output the model can produce. In Figure 9, one can see the effect of the arrival of MRO, MTO1, and MTO2, as well as the probabilistic decline of MGS and MEX.

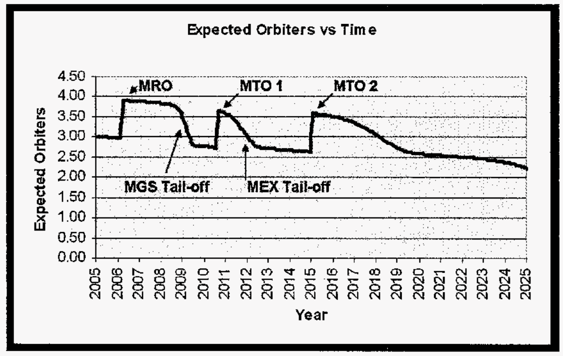

Figure 5: Expected Orbiters vs. Time

Robustness measures of Mars Network addressing performance handle the interplay between confidence and capability. Figure 10 illustrates the confidence versus time for a given 3.2 Gb/Sol desired UHF data volume capability for a large asset at 22 degrees south Mars latitude with the same hypothetical scenario used in Figure 9. Until the arrival of the first MTO-class orbiter, the desired capability is impossible. After MTO1, the confidence slowly degrades with time until boosted by the appearance of MTO2. Figure 11 shows an alternative way of viewing performance-based robustness for the same scenario. Here the expected capability is displayed given a confidence threshold of 90%. A more conservative threshold of 95% would result in reduced expected capability.

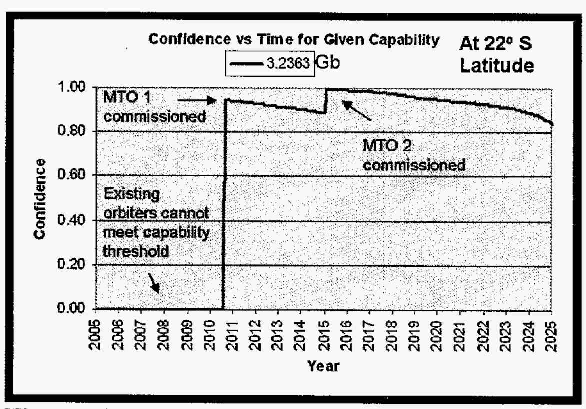

Figure 6: Confidence vs. Time for a Given Capability

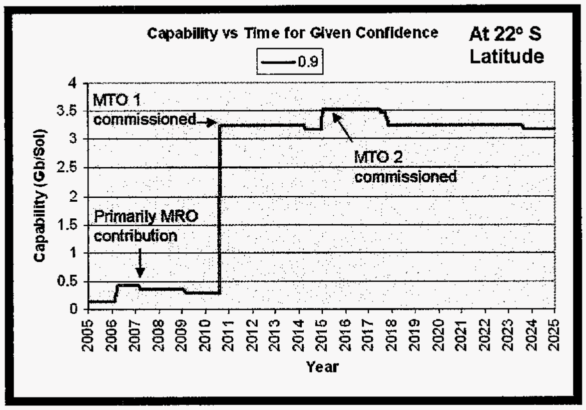

Figure 7: Capability vs. Time for a Given Confidence

Processing capability and confidence levels for every combination of time and latitude band allows multidimensional displays of the type provided in Figure 12. Here, the physical axes represent data volume capability, latitude, and time, while color indicates confidence. Confidence increases in shades from blue to orange to red. Continuing with the same scenario as before, one can see that polar capability dramatically increases with the arrival of MRO in 2006, but that equatorial and mid-latitude coverage remains in a relative trough until the arrival of MTO1 in 2010. The arrival of a second MTO in the 2015 timeframe increases potential capability dramatically but only extends the high-confidence region marginally above that provided during MTO1's early service.

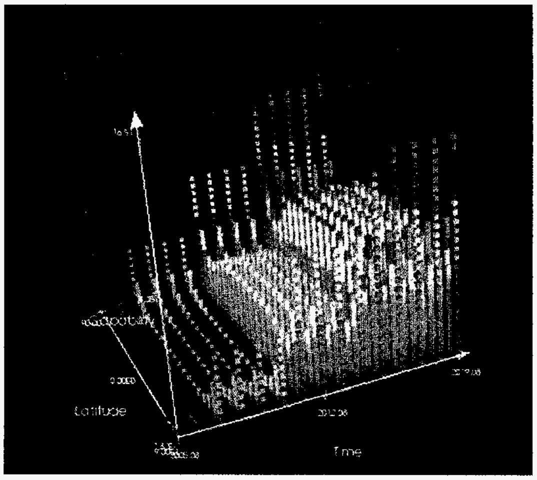

Figure 8: Mars Network Multidimensional Plot

6. Applications

The model allows the following what-if questions to be quickly addressed:

- What is the network significance of a missed launch opportunity?
- How does a failure at any point in any mission's life affect latitude coverage?
- In what location and timeframe is the network most or least sensitive to any failures?
- How does one Mars roadmap compare to another?
- To what extent could competed Mars Scout missions augment Network robustness?
- What are the implications for programmatic and international collaboration?

Future applications include post critical event planning, with network robustness updates given after retirement of known risks. In this manner, the framework could serve as an engine for dynamic strategic planning. Additionally, in conjunction with cost and operational utility information, the model could assist analyses of node productivity, addressing the issue of when it makes sense to decommission a functional node.

7. Summary

The emergence and continuing evolution of a capable Mars Network offering significant advantages to future surface assets highlights the need to understand the robustness of its current and projected relay service. A model has been developed to identify Mars Network robustness for all surface locations as a function of time. Components of the model include global telecom performance data for contributing spacecraft, probabilistic estimates of orbiter fuel expenditures, and probabilistic risk assessments for orbiter reliability. The model uses a framework allowing rapid comparisons between potential mission queues and operational assumptions. The resulting data outputs are suitable for high-level Mars Program Office studies of a strategic nature.

8. References

[1] C.D. Edwards, T.C. Jedrey, E. Schwartzbaum, A. Devereaux, R. DePaula, M. Dapore, T. Fischer, "The Electra Proximity Payload for Mars Relay Telecommunications and Navigation," 54th International Astronautical Congress, Bremen, Germany, 29 Sep. – 3 Oct., 2003.

[2] C.H. Lee, K-M. Cheung, C.D. Edwards, S. Kerridge, G. Noreen, A. Vaisnys, "Orbit Design Based on Global Maps of Telecom Metrics," IEEE Aerospace Conference, Big Sky, Montana, March 2005.

[3] C.D. Edwards, D. Bell, T. Ely, R. Hastrup, T. Jedrey, "NASA's Telecommunications Strategy for Mars Exploration," 19th AIAA International Communications Satellite Systems Conference and Exhibit, Toulouse, France, 17-20 April, 2001.

[4] L. Meshkat, A. Girerd, and C.D. Edwards, "An Integrated Approach for the Probabilistic Risk Assessment of the Mars Relay Network," to appear in RAMS Conference Proceedings, Newport Beach, CA January 2006.

9. Acknowledgements

The authors would like to thank the contributions of Pasquale Esposito, Jeff A. Lewis, Robert Mase, Sergio Guarro, and Karin Feldman.

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

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.