

MARKET-BASED APPROACHES TO MANAGING SCIENCE RETURN FROM PLANETARY MISSIONS

Randii R. Wessen*, David Porter** & Robin Ianson**

* Jet Propulsion Laboratory, California Institute of Technology
Pasadena, California U.S.A. Fax: 818-393-5074 E-mail: randii.r.wessen@jpl.nasa.gov

** California Institute of Technology
Pasadena, California U.S.A. Fax: 818-793-8580 E-mail: dave@hss.caltech.edu

ABSTRACT: The return of science is the fundamental objective of any planetary mission. However, which constellation of science observations constitute the best return of "science" is hard to evaluate. Past approaches toward planning science observations have been based on co-location of the payload scientists who debate the merits of which investigation had the "stronger" science. This advocacy approach is time-consuming and does not provide appropriate incentives for science teams to reveal their tradeoffs.

An alternate approach, currently under evaluation by the Cassini Mission to Saturn, is one based on providing better incentives to the science teams. Incentives can produce better tradeoffs because the individuals who can make the best decisions about which science observations to propose, what resources are required to implement the observations, and which observations are most important are the science team's Principle Investigators (PI) themselves.

1.0 THE VOYAGER PLANNING PROCESS

To illustrate the difference between the new planning process and previous approaches, we will compare a market-based science planning process to the one used during the Voyager Mission. The Voyager Mission represents a good baseline from which to compare since both Cassini and Voyager are extremely long missions. The Voyager Planetary Missions had a total of six planetary encounters and lasted twelve years. Figure 1 shows the geometry of one of the encounters, namely the Voyager 2 encounter with Saturn (1981 Aug.). Cassini, on the other hand, will have only one encounter and that is with the planet Saturn. After seven years of cruise, it will arrive at Saturn and then spend the next four years exploring it. Figures 2 and 3 show Cassini's initial orbit about Saturn and the next 63 orbits about the planet, respectively.

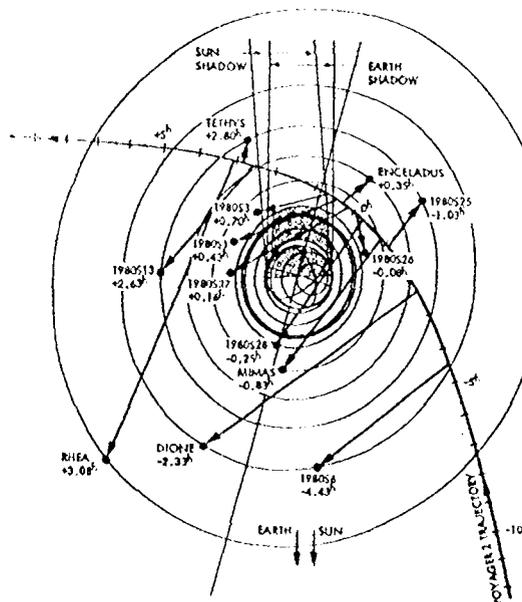


Figure 1. Voyager 2 encounter trajectory with Saturn. The encounter lasted four months with closest approach to the planet occurring on 1981 August 25.

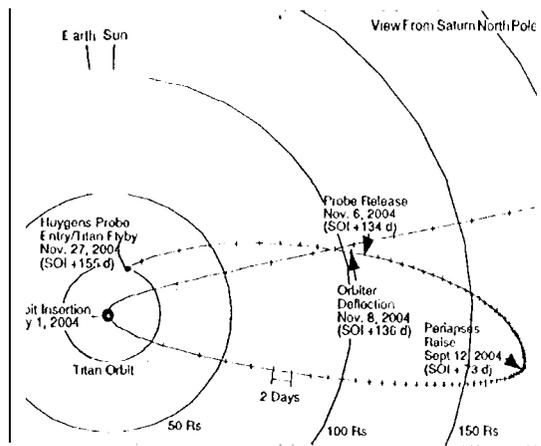


Figure 2. Cassini initial orbit about Saturn. Orbital period is approximately five months.

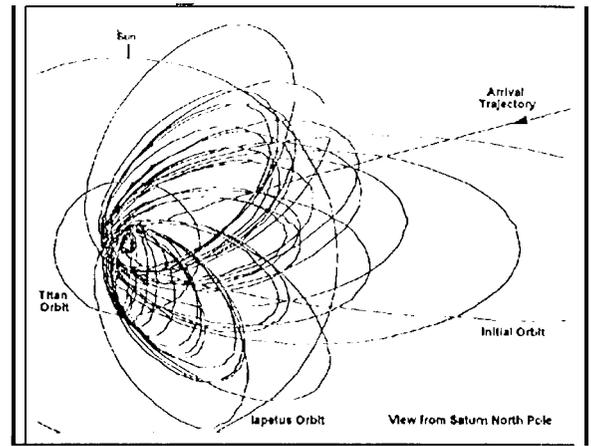


Figure 3. Cassini's four-year tour about Saturn. Orbital periods are approximately two months long and circle the planet 63-times.

The Voyager Planning Process¹ evolved during the 12 years of its use. The final process, used for both the Voyager 2 Uranus and Neptune encounters, began with Science Workshops. The results from the workshops were distributed to the Discipline Working Groups, who developed observation strategies. These strategies were used to answer questions identified by the previously mentioned workshops.

The Voyager Flight Science Office (FSO) had the difficult task of developing the timelines from the observations recommended by the working groups. The results, known as Scoping timelines, were finally presented to the PIs, who then evaluated and proposed changes (see Fig. 4).

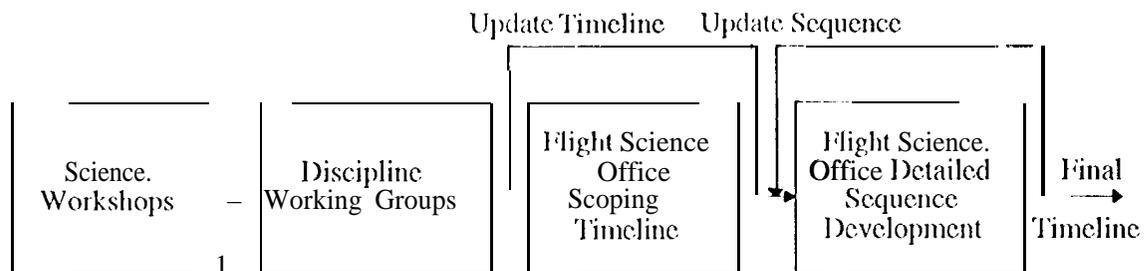


Figure 4. The Voyager Planning Process. Updates were required to both the Scoping timelines and the detailed sequences to incorporate PI changes.

1.1 VOYAGER DISCIPLINE WORKING GROUPS

The Voyager approach began with Science Workshops. These workshops brought together scientists who had expertise in the study of the particular planet. Their goal was to define the current state of knowledge of the target, and produce a list of major planetary objectives.

The objectives were divided into four disciplines, the totality of which defined the planetary system. The four Voyager disciplines were Atmospheres, Magnetosphere, Rings and Satellites. The science experts were then grouped according to their particular specialty. The resulting groups made up the Discipline Working Groups. In each group one individual was assigned as

chairperson. None of the Voyager PIs were included in the Discipline Working Groups to remove the possibility of a bias towards one particular investigation.

The chairperson of each working group presented their major science objectives to the Voyager PIs. This was then followed by a list of suggested observations, their durations and approximate times of execution in order to acquire the previously mentioned science objectives. The observations themselves were prioritized from 1 = highest value science, 2 = high value science and 3 = moderate value science.

1.2 VOYAGER SCIENCE SCOPING

Once each Discipline Working Group submitted their prioritized list of observations, the FSO integrated them into a single, conflict-free timeline. This timeline contained the "most important" science observations identified by the working groups. However, since the FSO had very limited information about tradeoffs between the many observations, it could only address the most obvious issues,

The approach taken to produce the timeline was to choose a priority 1 observation from each discipline. The FSO first selected the observation that required the most spacecraft resources (e.g., computer memory, integration time, propellant, etc.), as these observations were the most difficult to incorporate into the timeline.

Selecting difficult observations first can lead to reduced science return. Observations that were flexible were incorporated only after the difficult ones. This ensured that difficult observations made it into the timeline, but not necessarily the most important ones for the particular discipline. This approach can provide incentives to make observations resource-intensive and rigid,

Once the first round of observations were included into the timeline, a second round of priority 1 observations were selected. Observations were taken from all four disciplines, one at a time, to ensure that the timeline had "balanced" science. In this context, balanced science means that all disciplines have approximately the same number of observations selected from their lists.

The Scoping timeline was complete when as many of the observations, recommended by the working groups, were incorporated into the timeline. The remaining observations all required one or more resources that had already been allocated.

1.3 VOYAGER DETAILED SEQUENCE DEVELOPMENT

Once the Scoping timeline was produced, it went through a relatively large number of review cycles. During each cycle, the FSO presented the "new" timeline or sequence to the PIs, along with a detailed summary of the spacecraft resources required to implement the sequence.

The presentation to the PIs was the first time the investigators had a chance to evaluate which observations were incorporated into the sequence and how they were implemented. Though the Discipline Working Groups were experts in their field, they did not know all of the observations that were important to the PIs themselves. This lack of information forced many time-consuming review cycles. It also illustrated that the PIs do bring information into the sequence development process and that they should be brought in as early as possible to reduce the number of review cycles.

With each review cycle, recommendations were made to remove certain observations and replace them with others. Some observations may have required longer integration times which would force the FSO to try to "find" time in the sequence to accommodate them. Other

recommendations might require the science office to find unallocated spacecraft resources from which observation implementation problems could be solved.

In some cases, observations would compete for the same spacecraft resources in such a way that no compromise existed. One observation would have to be selected over the other. In these rare cases, the PIs appealed to the science office. With each appeal, the PI brought their case to higher levels of science office management, until finally reaching the Project Scientist.

It was the Project Scientist's responsibility to ensure that the "greatest" amount of balanced science was returned from the mission. Appeals would require conflicting investigation teams to prepare time-consuming arguments to present to the Project Scientist. On the basis of these presentations, conflicts were resolved.

1.4 PROS AND CONS OF THE VOYAGER PLANNING PROCESS

The multiple planetary rendezvous allowed the science community, and the associated science planning process, a chance to mature with each successive encounter. PIs learned what observations were of high value to the other investigators and which were not. This, in turn, taught the PIs which observations were available for trades. These trades between investigations could be outright exchanges of time on the timeline, spare spacecraft command words, or a change in the rate at which the spacecraft collected data from each instrument.

The process was well defined, robust and produced timelines that were able to provide the science community a wealth of information that is today still revealing secrets about Saturn. Unfortunately, the process was time-consuming, labor-intensive, and did not provide the FSO with the detailed information needed to make the best observation tradeoffs.

2.0 MARKET-BASED APPROACHES TO SCIENCE MANAGEMENT

The Voyager Planning Process described above has as its roots the idea that each science discipline should get its "fair share" of observation resources. Since the individuals who are scheduling observations do not know the resource tradeoffs for each investigation, they are left with the imprecise task of trying to schedule observations based on prioritized lists and a notion of "balanced" science. The questions they are confronted with are:

1. How many resources does the observation really require (e.g., Can I reduce the observation's duration or does it need all of the stated time)?
2. Are several priority 2 observations equal to one priority 1 observation?
3. Is it safe to assume that similar priority items have similar science values and tradeoffs?
4. If I try to schedule the most difficult high priority items first, what is the loss in science value that is imposed on other observations?
5. What level of resource reserves should be held to assist in the rescheduling of resources among users?

The individuals in the best position to answer these questions are the PIs. The FSO, in order to obtain a balanced timeline, must make these decisions and as a result, significant resources are used to argue for and make changes to the timeline. The question is whether the planning process can be revised so that it provides incentives to the science teams to supply the appropriate information so that the final timeline reflects their relative science tradeoffs.

This type of allocation problem is not unique, and is faced by almost all organizations that must allocate shared resources among users. Examples include use of supercomputer time, railroad tracks, computer networks, telescopes and laboratories, club facilities, etc. There are many different ways in which organizations deal with the allocation of shared resources. Our focus will

be on decentralized or market-based approaches. The main feature of market-based approaches is that the decision-making is left to those individuals that are in the best position to make the tradeoff (see Ledyard (1993)² for a description of decentralized mechanism design). In the commercial sector, this allocation problem is solved by charging for the use of congested services. For examples of pricing schemes used to reduce congestion in shared resource facilities see Westland (1992)³, Senkow (1992)⁴, and Sankaran (1989)⁵.

When direct pricing is not a viable option, some organizations turn to fixed budgets. That is, individuals are provided with a budget that can be used for a variety of services. For example, professors at the University of Chicago are given annual budgets that they can use for secretarial support, travel, computer hardware and software, etc. The use of fixed budgets relieves the manager from making tradeoff decisions that are best understood by the individual.⁶

As another example of this type of decentralized approach, the University of San Diego Supercomputer has used point budgets that allow users to prioritize their jobs. The more points assigned to a job the higher the likelihood it will be completed quickly. Since the point budgets are fixed, users must make decisions about when they run their jobs and how important the job is relative to all of their jobs (see Grether et al. (1995) for more details about these types of scheduling systems).⁷

2.1 A MARKET-BASED MECHANISM

The design question that we seek to address is whether we can construct a decentralized science planning system for the Cassini Project that can outperform the Voyager method. This is an ambiguous question since we must first define what we mean by outperforming and then we must be able to demonstrate that such a mechanism change results in better performance. Since we have no way of making inter-investigation science comparisons, it seems senseless to talk about maximum science value. If there is such a metric available, then it should be used directly as part of the planning process.

For our purposes, we are interested in two measures of relative performance. First, if one uses a market-based mechanism, is every investigation/discipline no worse off in the science it recovers, and at least one investigation does better than if a Voyager Planning Process is used. Second, within an investigation, we want to know what the relative science loss and gains are from each investigation. The method by which we will make these measurements is the use of experimental methods in economics (see the next section for a description of this methodology). What we will do next is describe the market-based mechanism we plan to test.

The mechanism begins with fixed scheduling point budgets being allocated to each science discipline and PI. The Discipline Working Groups then use scheduling points to rank observations instead of course priority classes. For lack of a better term, the allocation of scheduling points to observations will be called a "bid". The Cassini Science Office (CSO) then creates a conflict-free timeline from the bids by maximizing the number of scheduling points. Thus, conflicts for the List of resources by competing observations is resolved by the number of points submitted with the observation. Discipline Working Groups are also allowed to give allocations of the scheduling points provided to them, to specific investigations or to the CSO, to influence the timeline after the observations are first incorporated into it. In addition, after seeing the current timeline, bids can be revised. Specifically, bids on unincorporated observations can be INCREASED and new observations can be tendered.

Once the preferences of the Discipline Working Groups have been registered and a preliminary timeline has been formed, the Ds and the CSO can fine-tune the timeline by using their scheduling points to ensure observations stay in the timeline, or make new bids for alternative observations. The process is open and all bids are available for all to see. Every so often the timeline is updated with an algorithm that maximizes the number of scheduling points in

obtaining a conflict-free timeline. Thus, the process provides feedback to the teams to redesign their observations to fit into the timeline. The scheduling points are used to signal the relative worth of the observations. Hence, incentives are provided to not over-claim resources since this will require a large portion of scheduling points from a fixed budget. The process stops when changes to the schedule stop or are small. Figure 5 provides an overview of this process.

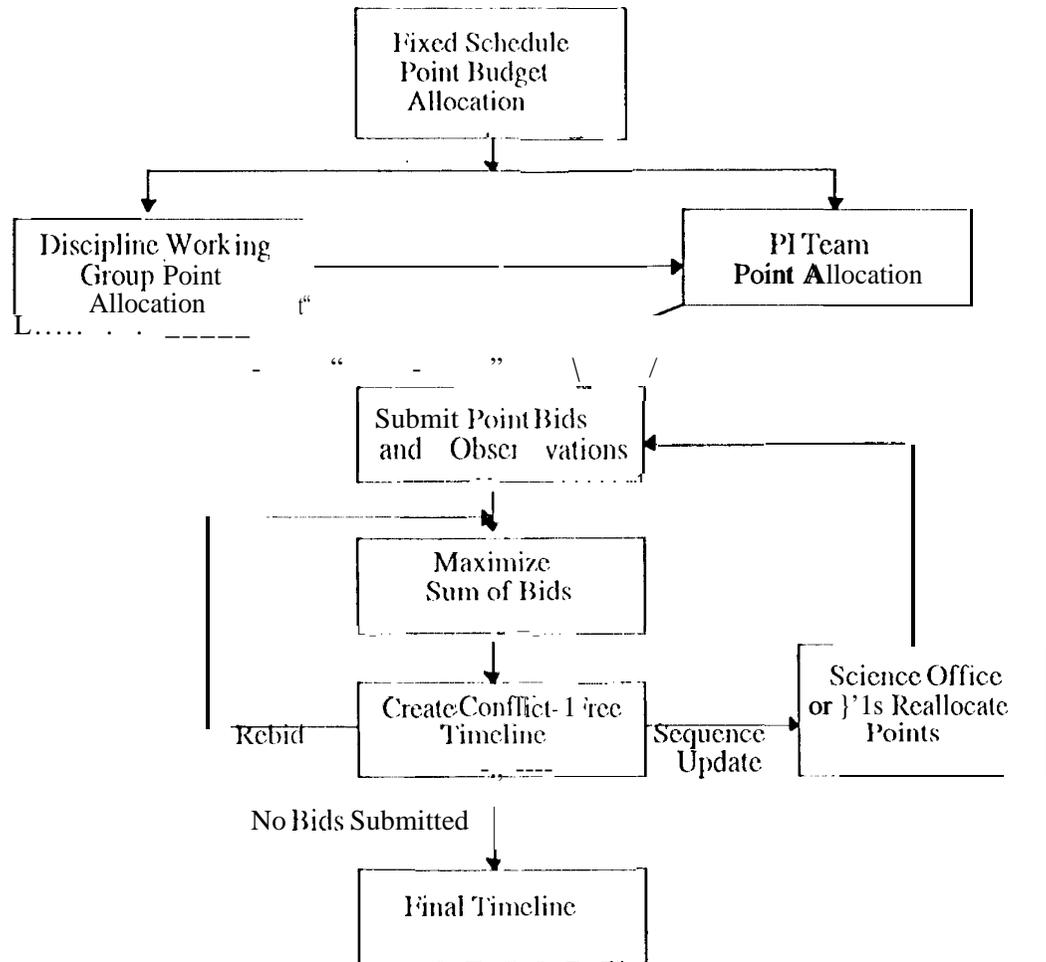


Figure 5. A market-based approach for generating a conflict-free science timeline based on inputs from the PI teams. Notice that either the Discipline Working Groups or the PIs may allocate the points.

Notice that this mechanism replaces the subjective scheduling decisions of the CSO and the rescheduling/adjudication process to the timeline with a single decision of allocating scheduling points to relevant participants. While this is not an easy decision process, it seems much less demanding, time consuming and less costly than the Voyager model. Given the limitations of space, in this paper, we will not provide a discussion of methods to make the initial allocation of scheduling points.

2.2 TESTING THE EFFICACY OF THE MECHANISM

While the intuition behind the design of the mechanism presented above might be apparent, how the cognitive processes of individuals interact in such an intricate mechanism is not as clear. What rules make the process more transparent and what type of information feedback works best are open issues. To test the ability of this mechanism to get real people to make tradeoff decisions, we will use experiments.

An experiment is basically a small scale, prototype of the process described above in which real individuals make decisions within the mechanism. These experimental subjects are motivated by cash payments. Specifically, subjects are recruited with the understanding that they will make money based on the decisions made in the experiment (mechanism). Subjects are provided with a description of how their specific allocations are transformed into individual monetary payments. Since the experimenter controls the underlying values, we can both replicate the experiments and can also make measurements and comparisons across mechanisms by knowing the underlying preferences. For those interested in the details of experimental economics methods see Smith (1982), Plott (1994)⁹ and Kagel & Roth (1995)¹⁰.

2.3 AN EXPERIMENTAL DESIGN

In order to test the difference in the performance and behavior of individuals in a Voyager Planning Process and the market-driven process described above, we plan to conduct the following experiment at design. An environment which involves scheduling science observations is constructed. The experiment has seven subjects representing seven PI teams. Each subject is then given a table showing the "science payoff" if their observations are included in the final timeline. See Table 1 for an example payoff sheet for a subject.

The table represents tradeoffs for the Imaging Science System (ISS) for high resolution surface mosaics. It describes the science value tradeoffs for this investigation, and this observation in particular, over three revolutions (e.g., rev. 1, rev. 2 and rev. 3) past Titan. The first two columns show the start and end time of the observation relative to Titan closest approach. The next three columns show the science return to the investigation by obtaining the corresponding start and end time for each revolution. So, if this investigation obtained time in the timeline from -02:20 to -01:00 on rev 2, they would get a science value of 70. This will be translated into dollars by a fixed proportional amount that will be kept by the subject. The table shows the tradeoffs for each combination of revolutions and various start and end times.

Start hh:mm	End hh:mm	rev. 1	rev. 2	rev. 3	rev. 1&2	rev. 1&3	rev. 2&3	rev. 1,2&3
-02:35	-01:00	80	75	70	95	90	85	100
-02:20	-01:00	75	70	65	90	85	80	95
-02:35	-01:15	70	65	60	80		75	90
-02:05	-01:00	65	60	55	80	75	70	85
-02:20	-01:15	60	55	50	75	70	65	80
-02:35	-01:30	55	50	45	70	65	60	75

Table 1. Science values for ISS high resolution surface mosaics.

Given the specified payoff tables provided to subjects, we can compare the performance of each scheduling regime. We will also examine the robustness of the market-based system under different initial scheduling point allocations.

3.0 CONCLUSION

This paper has described a research plan to design and test a new method for planning and negotiating science observations. The current method of a hierarchical process of science working groups and challenges to the timeline suffers from several ills. First, the process is time-consuming and many resources are used to make and remake the schedule. Second, the use of simple priority designations and a notion of "balanced" science is not enough information for schedulers to make important science tradeoffs. The information needed to make these tradeoffs is side with the 1's. Third, the incentives provided by the current process can result in poor outcomes because the information given to schedulers is not complete and the system is open to manipulation.

A more direct way to obtain science tradeoff information in which participants are given an incentive to provide accurate information is through a market-based approach. The market-based approach that we consider is one in which participants are given fixed budgets of scheduling points that are allocated by the project. The points are used to provide an intensity of preference for the observations being scheduled. In this way, schedulers no longer have to infer the science value of observations. The schedulers just try to maximize the number of scheduling points that result in a conflict-free timeline. Incentives are placed on the participants because they have a fixed budget from which to make their tradeoff decisions. Another important feature of the proposed market-based process is that there will be feedback so that individuals can rebid based on the current timeline. This feature has been shown to be important in obtaining high-valued outcomes (see Ledyard, Noussair and Porter [1996]).

Since the proposed process is new and has not been tried in the context of planning science timelines, the processes will be tested and designed using experimental methods in economics. This method allows for the direct testing of the performance of resource allocation schemes. This methodology will be used to provide scientific evidence on the performance of mechanisms that are used to allocate and develop science timelines.

3.0 REFERENCES

1. Miner, E. D., Stembridge, C. H., Doms, P. H., "Selecting and Implementing Scientific Objectives", *Journal of the British Interplanetary Society*, Vol. 38, 1985, pp. 439-443.
2. Ledyard, J., "The Design of Coordination Mechanisms and Organizational Computing," *Journal of Organizational Computing*, 3(1) 1993.
3. Westland, J., "Congestion and Network Externalities in the Short-Run Pricing of Information Services," *Management Science*, Vol. 38 Number 7 1992.
4. Senkow, J., "Pricing Mechanism for Resource Allocation in Decentralized Organizations: Experimental Evidence," PhD Thesis, University of Minnesota, 1992.
5. Sankaran, J., "Bidding Systems for Allocating and Pricing Indivisible Items in Certain Non-Market Contexts," PhD Thesis, University of Chicago, Graduate School of Business, 1989.
6. Ledyard, J. O., Porter, D., Rangel, A., "Using Computerized Exchange Systems to Solve an Allocation Problem in Project Management," *Journal of Organizational Computing*, Vol. 4(3), 1994, pp. 271-296.

7. Grether, D., Ishikida, T., and Porter, D., "Alternative Mechanism for the Funding and Allocation of Supercomputing Resources," *Economic Research Series*, Jet Propulsion Laboratory, 1995.
8. Smith, V., "Microeconomic Systems as an Experimental Science," *American Economic Review*, Vol. 72, 1982.
9. Plott, c., "Market Architectures, institutional Landscapes and Testbed Experiments," *Economic Theory*, Vol. 4 Number 1 1994.
10. Kagel, J. and Roth, A. , *Handbook of Experimental Economics*, MIT Press, 1995.

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of '1'ethnology, under a contract with the National Aeronautics and Space Administration.