NASA TECHNICAL NOTE

NASA TN D-7783

NASA TN D-7783

APOLLO EXPERIENCE REPORT -

A USE OF NETWORK SIMULATION TECHNIQUES IN THE DESIGN OF THE APOLLO LUNAR SURFACE EXPERIMENTS PACKAGE SUPPORT SYSTEM

by Richard A. Gustafson and Jeffrey N. Wilkes

Lyndon B. Johnson Space Center
Houston, Texas 77058

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION • WASHINGTON, D. C. • SEPTEMBER 1974

1. Report No.
NASA TND-7783

2. Government Accession No.

3. Recipient's Catalog No.

4. Title and Subtitle
APOLLO EXPERIENCE REPORT — A USE OF NETWORK SIMULATION TECHNIQUES IN THE DESIGN OF THE APOLLO LUNAR SURFACE EXPERIMENTS PACKAGE SUPPORT SYSTEM

5. Report Date
September 1974

6. Performing Organization Code

7. Author(s)
Richard A. Gustafson and Jeffrey N. Wilkes

8. Performing Organization Report No.
JSC S-404

9. Performing Organization Name and Address
Lyndon B. Johnson Space Center
Houston, Texas 77058

10. Work Unit No.
921-10-10-00-72

11. Contract or Grant No.

12. Sponsoring Agency Name and Address
National Aeronautics and Space Administration
Washington, D.C. 20546

13. Type of Report and Period Covered
Technical Note

14. Sponsoring Agency Code

15. Supplementary Notes

16. Abstract
A case study of data-communications network modeling and simulation is presented. The applicability of simulation techniques in early system design phases is demonstrated, and the ease with which model parameters can be changed and comprehensive statistics gathered is shown. The discussion of the model design and application also yields an insight into the design and implementation of the Apollo lunar surface experiments package ground-support system.

17. Key Words (Suggested by Author(s))
- Systems Compatibility
- Lunar Exploration
- Computer Simulation
- Data Links

18. Distribution Statement
STAR Subject Category: 31.

19. Security Classif. (of this report)
Unclassified

20. Security Classif. (of this page)
Unclassified

21. No. of Pages
12

22. Price
$3.00

APOLLO EXPERIENCE REPORT EDITORIAL COMMITTEE

APOLLO EXPERIENCE REPORT
A USE OF NETWORK SIMULATION TECHNIQUES IN THE DESIGN OF THE APOLLO LUNAR SURFACE EXPERIMENTS PACKAGE SUPPORT SYSTEM

By Richard A. Gustafson and Jeffrey N. Wilkes
Lyndon B. Johnson Space Center

SUMMARY

The Apollo lunar surface experiments package ground-support system and the Manned Space Flight Network are described functionally. A suspected design problem concerning the experiments package/network interface is discussed. It was learned that analysis of the data-communications interface problem was most feasible by the application of modeling techniques. The data-handling conventions of the Manned Space Flight Network communications system are discussed, and the embedding of these conventions in a general-purpose simulation-system network model is described. The model was run using anticipated network traffic loads for a lunar-landing mission and extensive data-traffic statistics. The results of the network simulation confirmed the suspected data-handling problems and suggested alternative solutions. The validity of the simulation results and of the implemented solution was confirmed by postmission analysis of the logged data traffic.

INTRODUCTION

In the summer of 1968, the NASA Lyndon B. Johnson Space Center (JSC) (formerly the Manned Spacecraft Center (MSC)) was charged with the real-time data acquisition, command, and control of the Apollo lunar surface experiments packages that were to be emplaced on the lunar surface during the lunar explorations. The support requirements necessitated by the Apollo lunar surface experiments package (ALSEP) program were considerably different from those previously encountered in the support of manned missions. Some of the unique ALSEP support requirements are stated in the following paragraphs.

Sustained, continuous support was required for long mission phases. Each ALSEP required 45 days of continuous support after deployment on the lunar surface, 2.5 days of continuous support after each terminator crossing (lunar sunrise and sunset), and 2 hours of support per Earth day. Because each ALSEP had a projected lifetime of 1 to 2 years, multiple mission support was also required. These long-duration, long-term support requirements had to be contrasted to the relatively short-duration manned missions (8 to 9 days for Apollo missions) for which the MSC Mission Control Center (MCC) was intended.

A high degree of support-system reliability was not required for ALSEP support, and data-system redundancy was not required. Restoration of failed systems had to be accomplished only within 2 hours. The rather modest system reliability and availability requirements for the ALSEP support were relaxed considerably from the severe requirements placed on manned-mission support systems; for example, the real-time computer complex (RTCC) required the attainment of a reliability (probability of zero failures) goal of 0.9995 for the critical phases of Apollo missions.

Support of the ALSEP did not require large support systems comparable to the scale of the MCC Apollo systems. Processor sizing estimates indicated that a medium-scale computer system would meet communications and data-processing requirements. The MCC Apollo system involved the use of a Univac 494 as a communications processor and an IBM 360/75 as a real-time data processor; both of these were large-scale computer systems.

An MCC ALSEP support system that was independent of but compatible with the Apollo systems was necessary because of the requirements that the ALSEP have sustained support, minimal reliability, and moderate processing; that the ALSEP support have minimal effect on Apollo support; and that the ALSEP systems be as compatible as possible with existing systems. Not only would the compatibility of the ALSEP and Apollo support systems facilitate necessary human and equipment interfaces with existing systems, but systems compatibility also would allow the same hardware and software technology used in Apollo support to be used in the design and implementation of the ALSEP support system.

SYSTEM DESCRIPTION

The resulting ALSEP support system used a medium-scale general-purpose computer system (an IBM 360/50) as a combined communications and real-time data processor. Early in the design stages of the ALSEP support system, it was recognized that maximum and efficient use of the MCC resources would result if the ALSEP could be supported by the existing large-scale IBM 360/75 computer systems constituting the RTCC. Initially, the following basic factors precluded using an IBM 360/75 to support the ALSEP.

1. At that time, Apollo and ALSEP development, testing, and simulations required all available computing resources, including the IBM 360/50 computer supporting the ALSEP.

2. Efficient use of the IBM 360/75 computer in support of ALSEP required the final development and checkout of a multijobbing real-time operating system so that the excess capacity of the IBM 360/75 supporting the ALSEP could be used for other tasks.

The requirements of reduced computing-resource needs and development of a multijobbing real-time operating system were met by the time of the first successful lunar landing (Apollo 11). Consequently, ALSEP support was relocated to the RTCC during the first lunar night after deployment of the early Apollo scientific experiments package, a reduced version of ALSEP. It should be noted that the eventual use of an IBM 360/75 rather than an IBM 360/50 computer had no functional effect on the ALSEP support system. In fact, even the programing system was completely compatible with either machine.

The ALSEP computer provided the necessary interfaces and processing capability to accept, format, process, and transmit ALSEP telemetry, command, display, and control data. After being telemetered from the lunar surface, all ALSEP telemetry data were sent directly from the Manned Space Flight Network (MSFN) through the NASA Goddard Space Flight Center (GSFC) to the MCC data links. This mode of operation allowed the ALSEP support system to function without dependence on the communications processors at the MCC. The ALSEP command and control data were received from the ALSEP display/control subsystem. Then, the commands were verified and forwarded to the ALSEP directly through the MSFN interface or, if the Apollo systems also were using the network, through the MCC communications processors to the network and, ultimately, to the ALSEP. Control data were used to select display formats, to control ALSEP telemetry stream selection, and, in general, to control the ALSEP computer processing. The ALSEP display data were transmitted by the ALSEP computer, under display control, to the ALSEP display/control subsystem. Display data consisted of data for the ALSEP high-speed line printer; digital data for digital event indications; and analog data for seismic drum recorders, analog strip-chart recorders, and analog meter displays. The ALSEP support system and its relation to Apollo support systems in the MCC are shown in figure 1.

As can be seen in figure 1, the goal of an independent ALSEP support system at the MCC was met by implementing a modest (by Apollo standards[1]) facility to be operated in parallel with the existing Apollo systems. However, both the ALSEP and Apollo systems required an interface with and used the MSFN for data communications and command traffic. The design and implementation of the ALSEP interface with the MSFN was a critical step in the evolution of the ALSEP support system.

---

[1] The MCC communications processing facility (communications, command, and telemetry system) was composed of three Univac 494 computers, and the real-time processing facility (the RTCC) was composed of five IBM 360/75 computers.

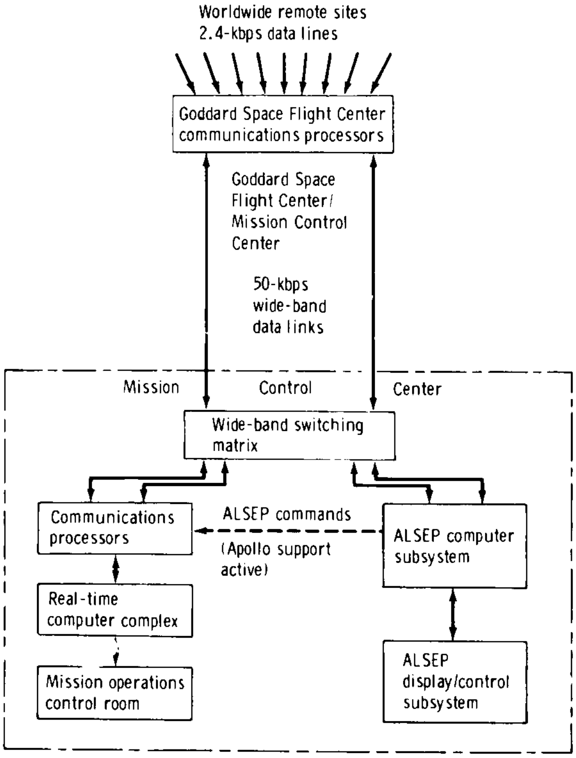

Figure 1. - The ALSEP and Apollo data flow.

NETWORK INTERFACE

A contradiction in requirements was encountered when the functional design of the ALSEP interface with the MSFN was initiated. Redundancy of the ALSEP support systems was explicitly not required, whereas the compatibility of ALSEP support with existing ground-support systems was required. These superficially unrelated requirements conflicted because all network traffic was received by the MCC through redundant 50-kbps wide-band data circuits operating in a primary/alternate configuration. An interface with both wide-band circuits would have required redundant ALSEP interfacing equipment. Because the network interface equipment was required to provide message framing and error-block encoding and decoding and to meet all data-set and computer conventions, the cost for each interface unit was expected to be quite high. The "no redundancy" requirement could not, therefore, be taken lightly.

One possible solution to the network interface impasse would have involved the construction of a single interface unit with the capability to switch between the two data circuits. If the prime circuit should fail,[2] a manual switch to the alternate circuit could be made. This approach would appear satisfactory because system failures required restoration within 2 hours. Certainly, loss of data (failure) could be detected and a switch to the alternate circuit made within seconds or, at most, minutes. However, one additional factor had to be considered.

The second line of the two wide-band data circuits between the GSFC and the MCC was used not only as a backup in case the first line failed but also as an overflow line in case the data-carrying capacity of the first line was exceeded. Therefore, data loss to the ALSEP facility would occur if ALSEP data overflowed to the alternate line while the interface unit was connected to the primary line. However, an Apollo mission ground rule required that mission traffic always be limited to the capacity of a single line. Assuming this ground rule was obeyed, overflow could be expected only when the short-term data rate exceeded the line capacity of 50 kbps while the long-term rate remained below 50 kbps.

---

[2]Circuit failure was assumed when the message originator failed to receive a message acknowledgment from the receiver for five consecutive 600-bit data blocks. Transmission would then switch automatically to the alternate circuit.

Because of a possibility of data loss, a study was initiated to determine the probable severity of data overflow to the alternate circuit and, if severe, to examine and recommend alternative courses of action. It should be noted that overflow could occur only during those periods when both Apollo and ALSEP data were being sent over the network; the ALSEP data alone could not precipitate data overflow. Moreover, if desired and if Apollo systems were not using the network, the data link between the GSFC and the MCC could be used in the "forced" mode; that is, by forcing all data over one circuit.

NETWORK DESCRIPTION

For the purposes of the analysis, the MSFN consisted of worldwide remote sites to acquire telemetry and trajectory data, 2.4-kbps data links to telemeter the digital data to the GSFC, communications processors at the GSFC to reformat and multiplex the data, and two 50-kbps wide-band data circuits to route the multiplexed data to the MCC. This data flow is illustrated in figure 1.

The telemetry equipment at each remote site was capable of transmitting two 2.4-kbps data streams continuously. Each 2400-bps data stream was formatted into 2400-bit blocks. The data rate and block size were fixed for each telemetry site and were not dependent on vehicle or mission. Each fixed remote site in high-speed tracking support produced ten 240-bit blocks/sec on a 2.4-kbps circuit.

The GSFC communications processor transmitted the received remote-site high-speed data to the MCC through the wide-band data links at a data rate of 50 000 bps in 600-bit bursts. Each 600-bit block consisted of a 480-bit block, containing data or fill (or both), and 120 bits of overhead. Therefore, five 600-bit blocks were required to support each remote-site telemetry stream (480 data bits/block × 5 blocks = 2400 bits). A maximum of two telemetry streams for each site was possible. Each fixed tracking site required ten 600-bit blocks/sec to transmit the 10 incoming 240-bit blocks. As of the date of the analysis, no attempt had been made to pack two high-speed tracking formats into one wide-band data format.

Because the GSFC message-switching system was designed to "throughput" data with minimum delays, a 600-bit block was formatted as soon as 480 telemetry bits or 240 tracking bits were buffered from a remote site. This design implied the formation of a 600-bit block every 200 milliseconds for each remote-site telemetry stream and every 100 milliseconds for each remote tracking site.

Each of the two GSFC/MCC wide-band data circuits used an output queue (a line of data blocks ready to be serviced) capable of queuing fifteen 600-bit wide-band data blocks. One circuit was designated the prime circuit and the other the alternate circuit. During normal operation, message transmission was on a "message rotary — primary/overflow" basis. Therefore, messages were transmitted on the prime circuit if available and on the alternate (overflow) circuit if the prime circuit was not available. Messages could consist of one or more segments (600-bit blocks), and a message could not be split across the two circuits. A telemetry message contained five segments, whereas a fixed-site trajectory message contained one segment.

The availability of a circuit was determined by the length of the output queue. On receipt of the first segment in each message, the count of the prime and alternate queues was compared and the message transmitted on the circuit with the shortest queue. If only one wide-band data circuit was made available (forced mode), the full 15-segment queue was used with no possibility of overflow. Two additional segment buffers for each queue were effectively "hidden," however, from the routing logic. These hidden-segment buffers resulted from the double buffering of outgoing segments within the wide-band data-network interface units (Univac special polynomial terminal subsystem type 33/22) at GSFC. Actually, the hidden-segment buffers made the routing operate as if overflow could occur only after three segments were in the prime queue instead of one.

NETWORK TRAFFIC

To perform an overflow analysis, a knowledge of the expected traffic to be handled by the network is required. Estimates for the worst-case network traffic during lunar-landing missions were four tracking sites with three two-line telemetry sites during lunar descent/ascent phases and three tracking sites with four two-line telemetry sites during launch phases. Remote-site support estimates of this magnitude resulted in the following traffic-loading predictions for the GSFC/MCC wide-band data link.

1. During descent/ascent phases: 40 segments/sec tracking and 30 segments/sec telemetry, yielding 42 000 bps

2. During launch phases: 30 segments/sec tracking and 40 segments/sec telemetry, again yielding 42 000 bps

The ALSEP data load had to be added to the Apollo traffic. In support of ALSEP (a maximum of two ALSEP sites), each remote site would contribute 5 segments/sec to the data-link traffic. Because the worst-case Apollo phases resulted in equal traffic, the aggregate anticipated worst-case loads on the GSFC/MCC wide-band data links were Apollo mission support with one ALSEP site (45 000 bps) and Apollo mission support with two ALSEP sites (48 000 bps). Although the predicted traffic loads were within a single wide-band data-circuit capacity of 50 kbps, queuing of segments with subsequent data overflow could occur if the segment arrivals to the circuit queue were not spaced equally in time. This "bunching in time" of segments was possible because the remote sites, which ultimately controlled segment generation, were not synchronized and could initiate a message start at any random time (random within a 1-second interval).

CIRCUIT OVERFLOW ANALYSIS BY MODELING

To obtain the probability of data overflow to the alternate wide-band circuit, the following three basic approaches were considered: an analytic probabilistic solution, an empirical solution obtained by actual network operation, and a modeling solution obtained by simulating the network. After the analytical approach was discarded as being essentially insolvable without drastic simplification assumptions and after the empirical approach was discarded as being impractical because of the cost and the difficulty of operating a worldwide data-communications network for test purposes, the simulation approach was chosen.

To simulate data-overflow conditions and to obtain statistics on overflow occurrences, a network model was developed. The model was driven by a variable number of telemetry and tracking sources corresponding to the remote sites. Each telemetry source was represented by the random generation of a message segment having a uniform distribution on the interval from 0 to 1 second and of subsequent segments having a period of 200 milliseconds. Each five successive segments constituted one telemetry message. Each line of the two-line telemetry remote sites was represented by a separate telemetry source. Each tracking source was represented by the random generation of a message segment having a uniform distribution on the interval from 0 to 1 second and of subsequent segments having a period of 100 milliseconds. Each segment constituted one tracking message.

The network model accepted each source segment and determined whether it was a "first of message" segment. If the segment was a "first" segment, the model checked the primary output queue count. If K (the queue-overflow level) or fewer segments were in the primary queue, the message was assigned to the prime queue. If more than K segments were in the prime queue, the alternate queue count was checked and the message was assigned to the shortest queue. The factor K was a variable. Each segment incremented the count of the queue to which it was assigned. Both queues decremented at the rate of once each 12 milliseconds (the transmission time of a 600-bit segment on the wide-band data links). The following statistics were to be collected from the model.

1. Average length of queues

2. Maximum length of queues

3. Number of tracking segments sent to prime and alternate queues

4. Number of telemetry segments sent to prime and alternate queues

5. Number of ALSEP segments sent to prime and alternate queues (The ALSEP was treated as a single-line telemetry source.)

6. Total number of telemetry, tracking, and ALSEP segments generated

7. Extent of prime and alternate circuit use

MODEL IMPLEMENTATION AND OPERATION

The network model was designed and implemented in the IBM general-purpose simulation system (GPSS) language and was executed on the MCC IBM 360/50 and 360/75 computers. The model was fairly short, requiring approximately 100 GPSS operations. As input parameters, the model accepted the queue-overflow level K, the number of telemetry sources (for which each double-line telemetry source was treated as two telemetry sources), the number of tracking sources, and the number of ALSEP sources (treated as single-line telemetry sites). As output, the model automatically provided the maximum length of queues, the total number of segments sent to each queue, the use of wide-band data lines, the number of segments sent to each line, and the various queue and line statistics. In addition to the aggregate queue and line statistics, tables on the queue contents (both primary and alternate) were generated for telemetry, tracking, and ALSEP data. These tables indicated the number of segments in the primary or alternate queue awaiting transmission after each segment was placed in the prime/alternate queue. This information allowed examination of the distribution of queue contents. The model time resolution was 1 millisecond, and the model runs were specified in integral numbers of seconds.

To calibrate and test the network model, several predictable traffic-loading cases were modeled. These cases started telemetry, tracking, and ALSEP traffic at specified times (as contrasted to random times). Because the start times were the only random processes in the model, the results were predetermined and agreed with predictions.

Because the traffic-source (telemetry, tracking, and ALSEP) start times were the only random variables in the model, network-traffic flow became periodic within 1 second (period of telemetry and ALSEP messages). Because of the periodicity, each simulation randomly started all sites, collected data for 1 second, and repeated the process for the period of specified time. The "restart each second" model allowed the maximum random variation of the traffic flow. Several cases of traffic flow were simulated by varying the combinations of traffic sources and aggregate traffic loads. In addition, several cases with varying queue-overflow levels were modeled. The results of these runs are summarized in table I. Cases having a queue-overflow level of two represented the model of the existing routing technique. A queue-overflow level of two was used because the previously mentioned double buffering of each network interface unit effectively resulted in two hidden-segment buffers in each queue. The "queue-overflow level" column of table I includes the two hidden buffers.

Data overflow occurred in all cases that were simulated with an overflow level of two, as seen in table I. Overflow continued at high traffic loads until the overflow level was increased to nine segments. It is interesting to note that the maximum difference between maximum queue contents with overflow levels of two and nine was only three segments, an indication of a minor segment time-delay difference between the different overflow levels.

MODELING RESULTS

With an overflow level of two segments, overflow of ALSEP data to the alternate line was deemed probable at anticipated traffic loads. This conclusion was based on the observed high rate of overflow shown by the simulation runs. The use of the prime line was fairly low (68 to 89 percent of capacity) on the runs with an overflow level of two (table I), an indication of a high overflow rate. An increase in overflow level to five segments significantly reduced the overflow probability but did not eliminate it.

TABLE I. - SUMMARY OF NETWORK SIMULATION RUNS[a]

Data sources			Aggregate traffic, kbps[b]	Queue-overflow level	Maximum queue contents, segments		Average queue contents, segments		Line use, segments		Percent overflow
Telemetry	Tracking	ALSEP			Prime	Alternate	Prime	Alternate	Prime	Alternate	
6	5	2	54	2	6	4	1.69	0.23	0.887	0.189	17.5
6	4	2	48	2	6	4	1.49	.14	.839	.117	13.5
6	4	1	45	2	6	4	1.37	.09	.813	.083	9.3
6	4	0	42	2	5	3	1.27	.07	.774	.062	7.5
4	4	1	39	2	5	3	1.14	.04	.737	.039	5.1
3	4	1	36	2	5	3	1.03	.03	.688	.028	4.0
6	5	2	54	5	9	4	4.00	.09	.990	.085	8.0
6	4	2	48	5	8	2	2.41	.003	.953	.003	.38
6	4	1	45	5	7	2	1.93	.002	.894	.002	.22
6	4	0	42	5	7	1	1.59	.000	.837	.000	.014
4	4	1	39	5	6	1	1.36	.000	.777	.000	.031
3	4	1	36	5	6	--	1.11	--	.717	--	--
6	5	2	54	9	12	4	7.81	.09	.996	.079	7.4
6	4	2	48	9	9	--	2.42	--	.956	--	--
6	4	1	45	9	8	--	1.93	--	.896	--	--
6	4	0	42	9	7	--	1.62	--	.837	--	--
4	4	1	39	9	7	--	1.32	--	.777	--	--
3	4	1	36	9	6	--	1.13	--	.717	--	--

[a]Each run randomly started traffic sources every second and simulated 300 seconds of real time.
[b]Line capacity was 50 kbps.

At an overflow level of nine segments, no data overflows occurred during all simulation runs with traffic loads below line capacity. This elimination of overflows occurred with a maximum increase of three queued segments over the cases with an overflow level of two (from six to nine segments). It should be noted that an overflow is possible (but not probable) even with an overflow level of nine. Receipt of 10 segments within 12 milliseconds would have caused an overflow at 30 kbps.

Certain combinations of telemetry- and tracking-segment arrival times could have resulted in overflow, even though aggregate traffic loads were below line capacity and a high overflow level was used. Because the number of possible combinations of data-source arrival times ranged from 10^{21} to 10^{36} (assuming 1-millisecond time resolution and start times from 6 to 12 data sources uniformly distributed over 1 second), an exhaustive examination of all possible combinations clearly was impossible. Therefore, the model could randomly select only a small sample of these combinations. Consequently, it was cautioned that, although no overflows occurred with traffic loads below 48K with an overflow level of nine in the model runs, the probability (apparently slight) did exist that overflow could occur.

CONCLUDING REMARKS

The results of the network simulations indicated that a single Apollo lunar surface experiments package facility interface with the network would result in data loss if the existing data-routing techniques were used. Three alternative solutions to the data-overflow problem seemed apparent. The first two proposed solutions, increasing the queue-overflow level to nine segments or recognizing and force-routing all Apollo lunar surface experiments package data on the prime line, had to be eliminated because they would have required modifications to the routing programs at the Goddard Space Flight Center. The risk to Apollo missions that is inherent in any routing-program changes was considered unacceptable. The remaining solution required compliance with network conventions and the implementation of functionally redundant network interface units, one unit for each of the two wide-band data circuits.

The implementation of two interface units was the course of action followed. The validity of this solution was verified when, after the support of the Apollo 11 mission and the early Apollo scientific experiments package, examination of data log tapes indicated significant Apollo lunar surface experiments package data overflow to the alternate line. These data were, of course, received on the second network interface unit.

Lyndon B. Johnson Space Center
National Aeronautics and Space Administration
Houston, Texas, May 6, 1974
921-10-10-00-72

NASA-Langley, 1974 S-404

"The aeronautical and space activities of the United States shall be conducted so as to contribute . . . to the expansion of human knowledge of phenomena in the atmosphere and space. The Administration shall provide for the widest practicable and appropriate dissemination of information concerning its activities and the results thereof."

—NATIONAL AERONAUTICS AND SPACE ACT OF 1958

NASA SCIENTIFIC AND TECHNICAL PUBLICATIONS

TECHNICAL REPORTS: Scientific and technical information considered important, complete, and a lasting contribution to existing knowledge.

TECHNICAL NOTES: Information less broad in scope but nevertheless of importance as a contribution to existing knowledge.

TECHNICAL MEMORANDUMS: Information receiving limited distribution because of preliminary data, security classification, or other reasons. Also includes conference proceedings with either limited or unlimited distribution.

CONTRACTOR REPORTS: Scientific and technical information generated under a NASA contract or grant and considered an important contribution to existing knowledge.

TECHNICAL TRANSLATIONS: Information published in a foreign language considered to merit NASA distribution in English.

SPECIAL PUBLICATIONS: Information derived from or of value to NASA activities. Publications include final reports of major projects, monographs, data compilations, handbooks, sourcebooks, and special bibliographies.

TECHNOLOGY UTILIZATION PUBLICATIONS: Information on technology used by NASA that may be of particular interest in commercial and other non-aerospace applications. Publications include Tech Briefs, Technology Utilization Reports and Technology Surveys.

Details on the availability of these publications may be obtained from:

SCIENTIFIC AND TECHNICAL INFORMATION OFFICE

NATIONAL AERONAUTICS AND SPACE ADMINISTRATION

Washington, D.C. 20546