

NAVIGATION FOR THE NEW MILLENNIUM: AUTONOMOUS NAVIGATION FOR DEEP-SPA CE-1

J. E. Riedel, S. Bhaskaran, S. P. Synnott, S. 1). Desai,
W. E. Bollman, P. J. Dumont, C. A. Halsell, D. Han, B. M. Kennedy,
G. W'. Null, W. M. Owen Jr., R. A. Werner, B. G. Williams

Navigation and Flight Mechanics Section
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California; USA; 91109

Phone: 818-354-8724, Fax: 818-393-6388, Email: jcr@radtran.jpl.nasa.gov

ABSTRACT

The first flight of NASA's New Millennium Program, Deep-Space-1, will include a new navigational technology: a fully autonomous optical navigation system. The *DSI* Navigation system will be the first use of autonomous navigation in deep space. The task for this system is to 1) perform interplanetary cruise orbit determination, using images of distant asteroids, 2) control and maintain the orbit of the spacecraft using the ion propulsion system (IPS, another technology never applied to deep space) and conventional thrusters, 3) perform approach orbit determination and control using images of the science targets, 4) perform late knowledge updates of target position during close fast flybys in order to facilitate a high degree of quality data return from 2 targets: asteroid *McAuliffe* and comet *West-Kohoutek-Ikemura*. Additionally, an encounter with Mars will probably be performed with possibly a close flyby of one of the Martian moons, *Phobos* or *Deimos*. In order to accomplish these tasks, several functional components are necessary. These include Picture Planning and Image Processing, Dynamical Modeling and Integration, Planetary Ephemeris and Star Catalog Handling, Orbit Determination Data Filtering and Estimation, Maneuver Estimation, Spacecraft Ephemeris Updates and Maintenance, and general Interaction with the other Onboard Autonomous Systems.

INTRODUCTION

Autonomous onboard optical navigation will be a necessary component of autonomous spacecraft operations for many future planetary exploration missions. Because of light-travel times, there are experiments and even missions that cannot be performed or have limited data potential unless autonomous navigation systems are incorporated. Close orbits around, or very fast flybys of, small poorly characterized objects are examples of such missions. Reducing operational complexity and costs is another goal of autonomous navigation systems. In a not-too-distant future, many small robotic missions may be simultaneously exploring the solar system. To increase the efficiency of these missions, the spacecraft themselves must take on more of the responsibilities of

their own maintenance, including navigation. Adapting many of the techniques proven for optical navigation for *Voyager* and *Galileo*, the New *Millennium DSI* onboard navigation system must autonomously plan picture sequences, perform image analysis, estimate the trajectory and calculate trajectory corrections using the low-thrust Solar Electric Propulsion system. New *Millennium DSI* will be the first planetary exploration mission to autonomously navigate all mission phases. The engineering of such a navigation system poses a number of very significant challenges. An overview of Optical Navigation and how it will be applied to *DSI* is given in Ref. 1.

The presence of an autonomous navigation system onboard a spacecraft imposes certain requirements on the onboard control system, and in turn, the capabilities and function of the control system will influence the architecture of the "Navigator." In fact, one of the more challenging developments of the navigation system is the construction of this interface. The nature of the interaction is to balance the resource needs of the navigation system with those of equally important onboard engineering and mission science objectives. These resources include use of the camera, slew time, mass storage capacity, fuel use, use of the system computer and total time in the sequence of events. The amount of resources devoted to the Navigator will often translate directly into performance of the system.

DSI MISSION ATTRIBUTES

An overview of the New Millennium Program and *DSI* in particular is given in Ref. 2. The *DSI* mission includes a very ambitious and challenging set of mission objectives and activities. There are probably three planetary targets intended for flyby encounters: asteroid *McAuliffe*, Mars, with possibly a close flyby of one of the Martian moons, and comet *West-Kahoutek-Ikamoura*. Currently, it is anticipated that launch will occur in July of 1998. The *McAuliffe* encounter will happen late January of 1999, the Mars flyby in late May of 2000, and the comet encounter about six weeks later. Figure I shows a heliocentric view of the mission trajectory, with

important mission events annotated. The annotations are referenced to Table 1.

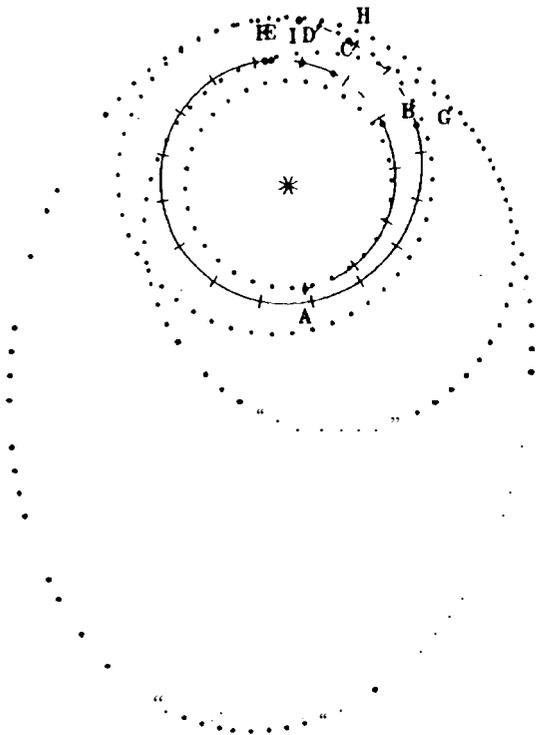


Figure 1. *DSI* Mission Design

At least for the McAuliffe flyby, the *DSI* spacecraft will perform the closest flyby encounter ever attempted in a deep-space mission: 5km from the surface of the asteroid. The encounter parameters of Mars have not yet been determined, but the flyby altitude of the comet will likely be on the order of several hundred kilometers, due to the dangerous environmental conditions near even a relatively inactive comet such as W-K-I.

ID	Time of Event	Description of Event
A	Jul 1, 1998	<i>DSI</i> Launch
B	Ott 24, 1998	End of first principal thrust arc
c	Dec 6, 1998	Beginning of second thrust arc
D	Dec 27, 1998	End of second thrust arc
E	Jan 16, 1999	McAuliffe encounter
F	Jan 20, 1999	Beginning of third thrust arc
G	Feb 8, 2000	End of third thrust arc
H	Apr 26, 2000	Mars encounter
I	Jun 4, 2000	W-K-I encounter

Table 1. Principal *DSI* Mission Events

The aggressive nature of these encounters is enabled solely by the presence of the autonomous navigation system. Performing navigation functions in a closed-loop sense onboard the spacecraft makes possible very late (before encounter) controls of the spacecraft encounter coordinates, and updates of knowledge about those coordinates.

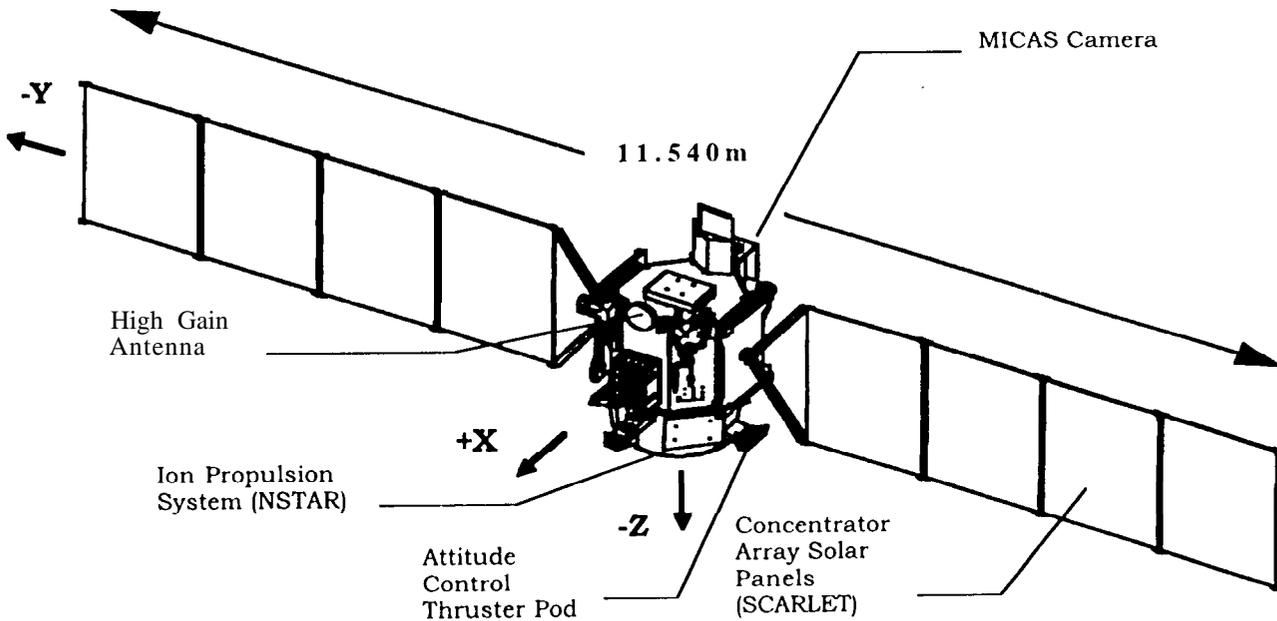


Fig. 2: New Millennium *DSI* Spacecraft

The objectives of the New Millennium Program (of which *DSI* is the first mission) is to develop and demonstrate new technologies which can enable future space exploration missions. The Autonomous Navigation System is one of these technologies being demonstrated. Another such technology, and one that has a fundamental influence on the nature of the *DSI* mission is its solar electric propulsion system. This system is actually composed of two technologies, a 2.5 kilowatt concentrator-element solar-electric array, known as "SCARLET," and an ion propulsion system (IPS) capable of approximately 100 mNt of thrust, known as "NSTAR". The IPS is principally responsible for making the energetically difficult triple encounter mission possible. However, this propulsion strategy seriously complicates the navigation task. Fig. 2 shows a schematic of the spacecraft, with annotations for the prominent solar arrays, the MICAS camera, and the IPS location on the -Z axis.

MISSION DESIGN IMPACTS ON THE NAVIGATION SYSTEM

Ion Propulsion System

The most challenging aspect of the *DSI* navigation task is the low-continuous-thrust, non-ballistic trajectory. This challenge begins with the design of the mission trajectory, which has been detailed elsewhere (Ref. 3). This highly interactive non-linear process is now, at the time of this writing, in its final stages for *DSI*, and is refined almost on a daily basis to reflect changes in the mass of the spacecraft, available power from the solar panels, available launch vehicle capacity and injection conditions, and thrust and efficiency of the engines. Once this design is complete however, it will be made available to the Navigator in the form of polynomial description of engine thrust direction and level as a function of time. A nearly final version of these tables is shown in Figs. 3-5.

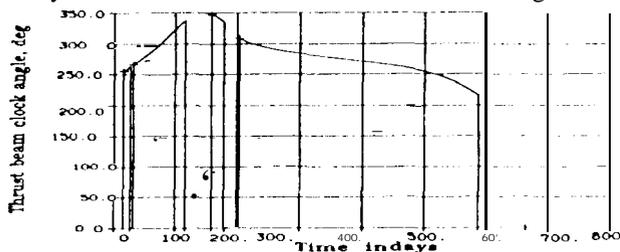


Figure 3. IPS Thrust Beam Angle

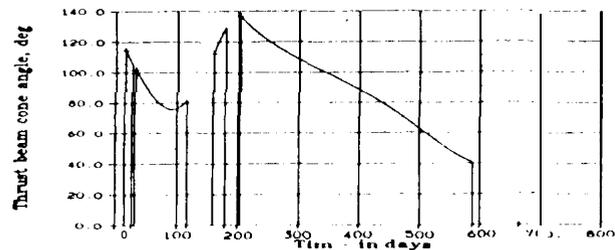


Figure 4. IPS Thrust Beam Clock Angle

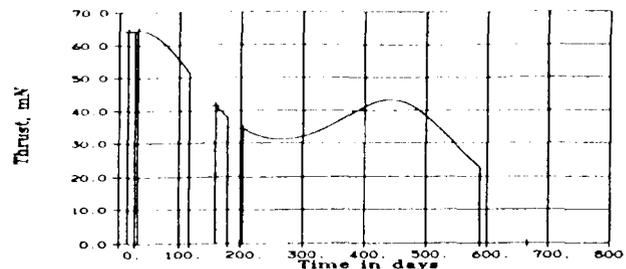


Figure 5. IPS Thrust Magnitude

The next navigation challenge posed by the presence of the IPS is the need to control the engine. It is not sufficient to guide the engine along the pre-computed polynomial functions. There are error sources in the implementation of the nominal design, with accuracies of between 1 and 2 percent expected. Such errors, when combined with normal statistical navigation errors, could if uncorrected, map to millions of kilometers over a seven month trajectory. Necessarily then, the nominal mission design needs to be constantly corrected to account for these errors. Additionally, the presence of the constant thrust of the IPS requires the Navigator to account for this force, and its errors, in the dynamic model of the spacecraft's course, and in the treatment of the optical data.

There is substantial uncertainty with regard to the operability and reliability of the IPS and the software managers for it, all being very new technology. This uncertainty must be reflected in the Navigator design by way of coping with inconsistent operation or outages. Such conditions present themselves as gross deviations from the nominal mission design. To the extent possible, the Navigator must use future control authority to correct for unpredictable and statistically anomalous trajectory perturbations. The spacecraft will be instructed to fly the planned thrust profile, representing thrusting at all available times (typically, about 92% of the time.) If outages occur, the Navigator will attempt to correct the trajectory for them. But if the attempt to linearly cot-rat produces a flightpath to the target which is overly energetically disadvantageous to subsequent encounters, the ground will intervene with a redesigned and optimized mission.

The IPS must be used for dedicated trajectory correction maneuvers during gaps in the mission thrusting, including approach to the encounter targets. The design of these maneuvers is quite different than with the use of conventional chemical thrusters. Since the IPS thrust is much lower (40 mNt vs. 200 mNt) these maneuvers take much longer. As such, the closer the maneuver takes place to the target, the more non-linear is the process to compute the parameters. Additionally, the *DSI* spacecraft is severely constrained in orientation. Some faces of the spacecraft bus cannot be illuminated by the sun, or may be so at only shallow angles, and/or for short periods of time. Use of the IPS constrains the spacecraft to have the

solar panels directly on the sun, with virtually no deviation margin. This combination of constraints (and several others) means that there are significant regions of the celestial sphere at which the IPS engine cannot point. Fig. 6 shows this constraint space, and Table 1 identifies the particular constraints noted. The result is that through communication with the Attitude Control System (ACS) (Ref. 4), the Navigator must ascertain if the computed maneuver direction is in a forbidden region, and if it is, redesign it to be a vector-decomposed maneuver in two directions that are allowed within the constraint space. This process is known as "vectorization"

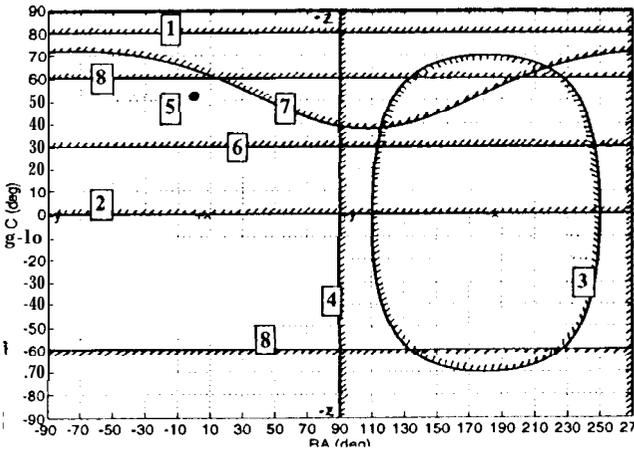


Figure 6. Illumination-Forbidden Regions of Spacecraft Body.

#	Constraint	Cone
1	MICAS Primary Aperture	+/- 10 deg. (+Z)
2	MICAS Optical Bench Radiator	+/- 90 deg. (+Z)
3	MICAS IR Radiator (At all times)	+/- 70 deg. (-X)
4	MICAS IR Radiator (IR in operation)	+/- 90 deg. (-X)
5	MICAS Occultation Port	+/- 1.6 deg.
6	PPU Radiator	+/- 60 deg. (+Z)
7	Star Tracker Boresight	+/- 35 deg.
8	ACS Kinematics Amplification Factor	+/- 30 deg. (+Z and -Z)

Table 2: DSI Constraint Space Magnitudes and Directions

Close Encounters

Another large impact on the Navigator from the rest of the system is the very aggressive nature of the mission. Next to the necessity to control the IPS, maintaining the spacecraft position knowledge and pointing through very close and very fast flybys is the most challenging requirement on the Navigator design. The requirement to keep the encounter target in the camera field of view when possible, created the need to perform the "reduced-state" navigation as discussed below. The close flyby distance of the McAuliffe encounter creates the requirement for an unprecedented control accuracy, necessitated not only by safety concerns, but also because relatively small

perturbations in the flyby asymptote produce serious deviations in target-relative geometry due to the close range, possibly disturbing a carefully constructed observation experiment.

REQUIREMENTS ON OTHER MISSION SYSTEMS IMPOSED BY THE AUTONAV SYSTEM

High Accuracy Imaging Instrument

Potentially, the most obtrusive requirement that the Autonomous Optical Navigation System (AutoNav) places on the spacecraft design is for the presence of a very high quality telescope with which to perform the interplanetary phase of the navigation task. Some periods of the approach navigation also depend upon high quality astrometry, and therefore require a science-capable telescope. Fortunately, most scientifically sophisticated deep-space missions (including DSI) carry a camera capable of providing adequate data for the class of astrometry needed by navigation. An overview of requirements posed by AutoNav, and met by MICAS (the Miniature Imaging Camera and Spectrometer) being flown by DSI is given here:

- 12-bit digitization. This is required to maintain sufficient dynamic range to image bright extended objects and dim stars.
- 0.6 to 2.0 degree field of view. This is required to maintain (with typical chip size, adequate resolution for the cruise optical navigation, typical resolution range is 5 to 40 microradians per pixel.
- 1024 x 1024 pixel array. Such an array size is the minimum standard for quality CCD's, and will determine (via the focal length) the pixel resolution.
- Focused Image to 0.1 pixel. Typical focused optics give adequate point-spread-functions to provide this capability without intentional defocusing.
- 80,000 electron "full-well" with 50c" noise. This is a description of the dynamic range and signal quality of the instrument, which is important to define the effective working span of useable brightnesses.
- Image 13th Magnitude star. This should be possible in a long (smeared exposure) and represents the minimum useable detection of cruise targets, and reflects the presence of accumulated photons/charge from repeated overlays of the drifting image.
- Image 9th Magnitude star. This should be possible in a short (unsmeared) exposure. Such images are the normal mode on approach to a target.

Flight Computer Requirements

The DSI flight computer is a RAD6000 based computer system operating at 33MHz. This computer is a radiation hardened version of an IBM-6000 series work-station computer. There are 96Mega-Bytes of hardened RAM available, which is used as both memory and mass storage. There is 16MB of non-volatile memory from which the computer boots. It is estimated that at least

50MB of RAM will be available for Science and OpNav data storage, and about one-half of the available CPU capacity will be available for Science and OpNav processing during most of the mission.

The computational requirements imposed on the flight computer and data system are relatively modest in most cases. The size of the object code in running configuration, including static variable storage, is about 2 Mega-Bytes. The star catalog, containing about 125,000 stars occupies about 2 MB. The ephemeris file, with the major planets and about 250 minor planets is about 0.5MB, and other miscellaneous files occupy another approximately 0.5MB. The code and data files will be resident in non-volatile memory (EEPROM). The spacecraft system will load the programs and data from EEPROM into RAM at boot time, and those copies will be used for processing. At least once per day, and more often during critical activities, copies of the current data, including currently best-estimated states, data summaries, and the non-gravitational force histories will be written into EEPROM to protect the data from a system failure with associated CPU reboot. At reboot, the latest stored data is recovered, and the Navigator proceeds in a normal fashion.

Timing and throughput requirements are not stringent during interplanetary cruise; there is ample time during this phase to plan the images and perform the processing. (A detailed description of the operational activities is given below). When the Navigator has an opportunity to take images, the planning process takes only a few seconds. The processing of each cruise image is estimated to take up to a minute, but since each cruise exposure is about 100sec in duration, it is thought that the precision astrometric processing will keep up with the pace of imaging; especially when considering that several minutes (up to 30) will be required to turn the spacecraft from target to target. Nevertheless, there will likely be room available in the RAM-disk space to hold a number of images if the Navigator, for some reason, is delayed in processing. When finished with image processing, the Navigator will delete the images, or select a small subset for compression and downlink, especially in the early portion of the mission. Additional computational leeway is provided from the fact that during the cruise phase, the information content of the data is not changing quickly, and therefore it is only necessary to infrequently process the reduced image data into a solution of the spacecraft state, a process which can take several minutes.

During the encounter phase of the mission, the timing requirements of the Navigator are much more stringent. In the last 5 minutes on approach to the target, a series of up to 5 OPNAV opportunities occur. These are at increasing frequency, to capture the rapidly increasing information available in the images about range to the target, knowledge of which is critical to keep the asteroid

in the field of view until the last possible moment. Table 3 shows the image times, ranges, and associated spacecraft state knowledge with each of the late pictures. The timing of these frames is very close, and there is not sufficient time to perform all of the normal processing. Therefore a reduced form of the navigation processing is invoked about 30 minutes out, allowing image processing and orbit determination to complete in 10 to 15 seconds. The spacecraft target-relative ACS held ephemeris is then updated with each image, by means of a simple and quick 3-dimensional bias state change to a previously delivered full 6-d ephemeris. Since these updates occur in a matter of seconds, the target can be held within the field of view until the ACS can no longer physically accelerate the spacecraft into a turn at a fast enough rate.

Picture Time (sec)	McAullife Range (km)	Downtrack Error (km)	Crosstrack Error (km)
-20	164	0.8	0.5
-40	328	1.6	0.5
-80	656	3.2	0.5
-160	1312	7.5	0.5
-320	2624	15	0.5

Table 3: Near Encounter OpNav Picture Statistics

Interfaces with ACS, IPS and Sequencing Managers:

A number of interfaces with other flight software subsystems have already been alluded to. The most technically intricate of the inter-system interfaces is with the ACS. This interface is a set of different queries and responses. The Navigator must ask the ACS for a number of types of information: current attitude of the spacecraft; specifications on turns, such as estimated length of time required to turn from one attitude to another; the validity of a specific attitude for a maneuver; and the accumulated velocity due to general RCS (Reaction Control System - a subsystem of the ACS) activity. ACS, in turn, queries AutoNav for information, including: current mass of the spacecraft; and current spacecraft and planetary ephemeris information. Through an indirect sequencing operation (to be discussed below) the Navigator will request the ACS to perform specific operations; for example, turning to a specific attitude, for image taking or IPS thrusting. ACS will also be asked to execute a Trajectory Correction Maneuver (TCM) with the RCS or execute a TCM with the IPS. AutoNav also maintains an interaction with the IPS manager: IPS reports to AutoNav the currently accumulated thrust while the IPS engine is operating; and AutoNav will, through the sequencing interface, request the IPS to go to a specific thrust level and burn for a specific duration. The third principal interface that the Navigator maintains is with the *Sequencer* itself, and this is the simplest major interface. The Navigator will prepare very short sequences (lists of time-ordered commands) to perform specific tasks and ask the Sequencer to start or "launch" them. Additionally, during encounter, the Navigator will be

called upon to launch specific encounter sequences at specific encounter-relative times.

Data Uplink and Downlink Requirements:

Necessarily, the Navigator requires a certain level of information transfer both on the uplink and downlink. This is especially so for this the first flight of the system. The early portion of the mission (the first three or four weeks) will see intense use of the telemetry system to downlink dense data sets pertinent to the evaluation of the new technologies. AutoNav will be among these. Principal among the data to be downlinked in this early evaluation period will be the OpNav images themselves. Other data will include processed results from the Navigator, including reduced image data, (centers of asteroids and stars in individual frames) computed orbit determination results and maneuver solutions. It is anticipated that after a short period of evaluation of the dense telemetered navigation data, that the data can then be r-educed, compressed or stopped. On approach to the asteroid, the first target, there will again be a short burst (a few days) of dense data, to confirm that the Navigator is initiating approach operations properly.

Again, given normal performance of the AutoNav system, uplink requirements should be fairly modest. The largest sets of information likely to be required sent to the spacecraft are new thrust profiles, reflecting newly redesigned mission trajectories, and asteroid ephemerides. It will likely be necessary to redesign the mission trajectory at several points during the mission. The first such time is shortly after launch when the injection errors are known. Although nominal performance of the Delta 7326 launch vehicle is expected, greater than a one-sigma dispersion of about 100m/s will likely necessitate a redesign of the trajectory. The onboard maneuver computation algorithm will not be able to retarget the spacecraft in a fuel efficient manner in the face of such an injection error. Although the maneuver subsystem is tolerant to a certain degree of uncertainty in the engine performance, if the IPS operation deviates from the schedule by two weeks or more, it is again likely that the mission trajectory and thrust profile will have to be redesigned. Finally, it is expected that immediately after the McAuliffe fly-by that the ground operations Navigation team will redesign and uplink the trajectory and thrust profile. The process of optimizing the flight path for fuel use between two flybys is beyond the current capabilities of the flight DSI AutoNav system.

Operational Demands, and Staffing

Despite the expected periodic intervention of ground operations as outlined above, the AutoNav system will exhibit a high degree of autonomy. Operations, such as TCM's and image processing which used to require a significant amount of personnel on navigation and other teams will occur automatically without even the need for the ground to approve the AutoNav system's decisions.

Even in the early part of the mission when extensive analysis of the operation of the onboard Navigator will be taking place, the size of the Navigation team will only be between four and five persons, and this includes at least two performing the validating conventional radio navigation task. This bodes well for future missions using versions, and possibly even extensively expanded versions, of the DSI AutoNav system. It is estimated that a maximum of three persons would be necessary to fully analyze and maintain the operation of the AutoNav system for future missions at least as ambitious as DSI.

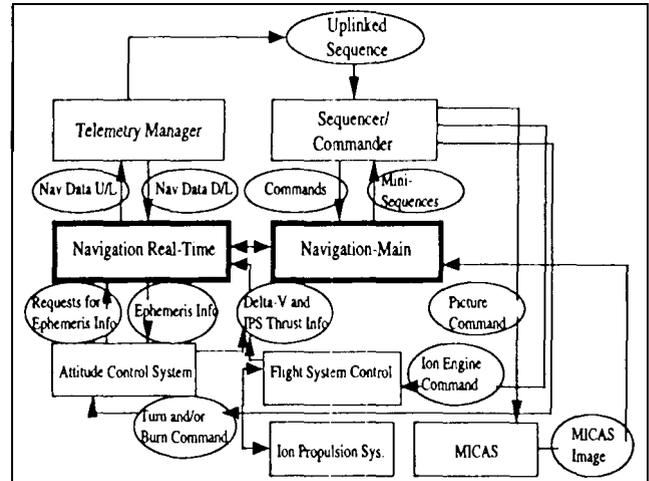


Figure 7: Navigation System Architecture

AUTONOMOUS NAVIGATION SYSTEM DESIGN:

Architecture

The overall DSI software system architecture is shown in Fig. 7. The overall system is based largely on the Mars Pathfinder flight software system. Mars Pathfinder is a conventionally controlled spacecraft, meaning that long series of commands (sequences) are uplinked to the spacecraft for timed execution (Ref. 5). Despite the deterministic nature of the nominal control system, autonomous navigation is still part of the design. This is accomplished by leaving large gaps in the ground-generated stored sequence, in which the AutoNav system is allowed to accomplish autonomous operations; this mode of operations will be discussed in detail below.

The Navigation system is composed of two subsystems, a real-time link, Nav-RT, and a main non-real-time computational link, Nav-Main. The real-time link is responsible for maintaining the ephemeris information for the ACS subsystem and for collecting information about propulsive activity onboard from the ACS and the IPS managers and formatting and relaying it to Nav-Main.

The flow of control through the flight software system and the Navigator is shown in Fig. 7. Normally,

commands to the Navigator come via the Sequencer in an uplinked stored sequence. A summary of the possible commands that the Navigator can process is given in Table 4. All requests for action that the Navigator makes, will also be made through sequences, but these will be short and spontaneously generated onboard by the Navigator itself. In addition to the commands received by and issued from the Navigator, there are a limited number of direct calls to the Navigator and returned replies. These were summarized above.

Command	Navigation Action
[NAV-SET-IPS]	Initialize the IPS thrust arc.
[NAV-IPS-UPDATE]	Update the IPS thrust and vector.
[NAV-DO-TCM]	Perform TCM operations.
[NAV-PHOTO-OP]	Plan and take Navigation Pictures
[NAV-START-ENC]	Start an Encounter sequence.
[NAV-DATA-UPDATE]	Update Navigation parameters.
[NAV-DO-OD]	Perform Orbit Determination.
[NAV-PLAN-TCM]	Compute TCM parameters.

Table 4. Navigation Command Summary

Functional Overview

At the most basic level of description, the AutoNav system uses pictures taken by the onboard camera to determine, *via* a dynamic filter, the spacecraft state. Propagating this state to the target body, retargeting parameters are computed and trajectory correction implemented. During the cruise portion of the mission, pictures of asteroids and stars are the principal data, but on approach to a target, images of that target with or without stars are the main navigational data. In the following sections, these functions will be detailed.

image Planning

The task of the Image Planning subsystem is to provide a schedule of targets for the AutoNav system. These targets include both beacon navigation targets as well as the approach encounter targets. The targets are clustered in time, to enable the planner, when asked, to access a set of viable target-asteroids to use for navigation purposes. The targets are additionally clustered and ordered to minimize turns between. Minimizing the cost of the turn sequences is important to minimize fuel usage. Because of the nature of the illumination constraints on the spacecraft, the beacon asteroids cluster into two discrete groups: those in the “forward” anti-sun half-hemisphere, and those in the “aft” anti-sun half-hemisphere. A fuel and time costly rotation of the spacecraft is necessary to turn from forward to aft, and so at most one such turn is scheduled for each observation opportunity. Within each half-hemisphere, the turns are additionally minimized.

Even though the above considerations are made as part of the ground operations, and possibly even before launch, there is a substantial amount of work for the onboard picture planner to do. Given only a list of asteroid targets, in optimized turn order, the picture planner must assemble a set of specific image requests, including turn

commands for exact pointings in inertial space. Additionally, it must predict the locations of the stars to be seen in the field relative to the target at precisely the time the picture is to be taken. This requires accurate storage and evaluation of ephemerides and star positions. The former will be discussed later, but the latter involves the use of accurately built star catalogs and requisite efficient storage of them. For *DS1*, the onboard star catalog will be based on the TYCHO Star Catalog (Ref. 6) and contains about 125,000 stars. The positions on this tile are accurate to at least 5 micro-radians, a least factor of two greater than is required to not degrade the accuracy of the autonomous OD process.

Image Processing

There are two types of image processing necessary during the mission, long-exposure smeared images of unresolved beacon asteroids, and short-exposure images taken on approach to a target. These latter are pictures of resolved and extended images.

In deep cruise, the need for long exposure images arises from the small size and extreme range of the beacon targets. The consequence of these long exposure times is to cause the ambient motions of the three-axis-intertially stabilized spacecraft to trace the star images of extended parts of the frame. Typical star and asteroid images will be smeared over 20 to 40 pixels. Fig. 8 shows a simulated version of the expected deep-space picture. Frames such as this have been used to test the algorithms and software. Also, simulations of the expected sort of image have been made using an astrometric observing system at the JPL observatory at Table Mountain. A series of these images, made to simulate the unstable characteristics of the spacecraft were made by manually slewing the telescope with its joystick controls. These images were then processed by the image processing subsystem of the Navigator. This analysis is documented in Ref. 7.

The processing system for the smeared cruise images was developed for the *Galileo* mission, and is documented elsewhere (Ref. 8) The theoretical basis of the system is a multiple-cross-correlation algorithm, that uses each of the nearly identically smeared star and asteroid images in a picture as a pattern. Each pattern is then used to locate every other pattern, with the result that extremely complex and often faint patterns can be located relative to one another to high accuracy, usually to 0.1 pixel (picture element) or better.

The actual correlation process can be summarized as a vector inner product. Given a normalized pattern, called a “filter”, that is composed of image elements in a matrix $m \times n$ in size denoted as F , and a sample area $S, M \times N$ in size, of which subset regions of $m \times n$ dimensions are extracted, then a function c_{ij} can be maximized:

$$c_{ij} = F \otimes S_{ij} = \sum_{k=1}^m \sum_{l=1}^n F^{kl} \cdot S_{ij}^{kl}$$

The maximum of c_{ij} represents the position of best match between F and the sample region

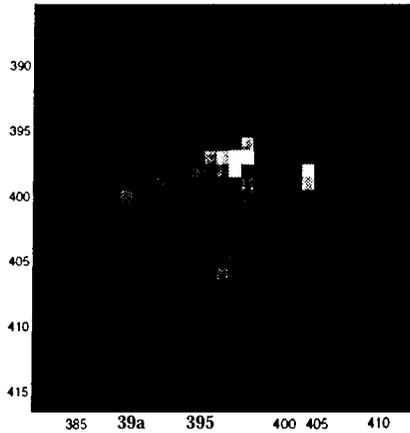


Figure 8. Simulated Cruise Asteroid Image

When the spacecraft nears one of its targets, and the unresolved image resolves, and consequently brightens, the exposure times necessary to image the object necessarily decrease. In fact, the opposite problem faced during the cruise imaging must be dealt with, namely the object becoming too bright to easily image in the same picture with dim stars.

Previous deep-space missions depending upon Optical Navigation (Principally Voyager and Galileo) have taken advantage of very accurate position determination of extended images of targets, namely images of the major body and its satellites. For weeks or months such images were available, and with the addition of reasonably good physical constants models (e.g. shape and size), extremely good position determination was possible. For these missions, a tenth to a quarter of a pixel was normal, translating in the final approach images to a few tens of kilometers (Ref. 9). For *DSI* this situation is quite different. The physical nature of the targets (with the possible exception of Phobos) is poorly known. The resultant uncertainty in the modeled figure presented to the camera results in significantly poorer center-finding. In compensation, the *DSI* targets do not become extended, and therefore subject to mismodeling errors, until the spacecraft is quite close,

It is guessed that the uncertainty in the diameter of *McAuliffe* and *W-K-I* is at least 50 percent, however the uncertainty of the centerfinding process is not nearly this large. The location of the extended images will be determined by a basic brightness centroiding technique. In general, the region in which the body image is located is

predictable to within about one hundred pixels before the picture is taken. Within this vicinity, those areas with brightness greater than background will be used to compute a brightness centroid. The centroid is adjusted for the approach phase angle, *via* the relationship given in the equation:

$$X(a) = \frac{3}{16} \pi R \frac{\sin \alpha (1 + \cos \alpha)}{(\pi - \alpha) \cos \alpha + \sin \alpha}$$

where X is the centroid offset, R is the object radius and a is the solar phase angle. If the approach phase angle were zero, the phase deflection term would be zero, and a brightness centroid measurement of the center of brightness would give an arbitrarily good measure of the geometric center of the body. For the two encounters to be flown where there is large uncertainty about physical constants, the phase angles are about 50 and 90 degrees. A simple differential with respect to diameter reveals the dependence of the phase correction of a diameter error. This relation evaluated for *McAuliffe* approach and *W-K-I*, gives a maximum of less than half a radius, which for both objects is well below a kilometer. As a result this error source does not make a dominant contribution to the overall control and knowledge errors of the *AutoNav* system.

For the late encounter knowledge update process (discussed below) the image processing procedure must be very fast. For this purpose, the precision of the brightness centroid is reduced by a simple process of data compression; the image pixels are merely undersampled. When the body-image is large, and therefore the relative size error as described above is larger, then the inaccuracies of undersampling do not contribute significantly overall to the navigational errors. Fig 9 shows a simulated version of an approach picture to *McAuliffe*. Images such as these are being used to test the algorithms and the flight software.

Orbit Determination

One important advantage of an all-optical-data orbit determination system is the insensitivity of the data type to high-frequency velocity perturbations. This is especially true for *DSI* which for the first time will perform a low-continuous-thrust propulsion strategy. Such systems are presumed to have significant time-varying thrust characteristics. With a velocity-measuring data type such as Doppler, this propulsion strategy poses substantial problems. These problems must be faced by the radio navigation that will be performed as part of the *DSI* operations and validation task, but they will not have to be performed by the onboard *AutoNav* system.

At the *core* of the Orbit Determination (OD) subsystem is the modeled representation of the spacecraft flightpath. This representation defines the nature and extent of the

parameterization and accuracy possible in the system. The Navigator models the spacecraft motion with a numerical n-body integration, using major solar-system bodies as perturbing forces. Non-gravitational perturbations to the spacecraft trajectory included in the model include a simple spherical body solar-pressure model, a scalar parameter describing IPS engine thrust efficiency, and small accelerations in three spacecraft axes. A spherical-body solar-pressure model is sufficient because for the majority of the time, the spacecraft will have its solar panels oriented toward the sun. Even though the spacecraft can maintain this orientation with any orientation of the bus-body about the panel yoke axis, the panel orientation by-far dominates the solar pressure effect.

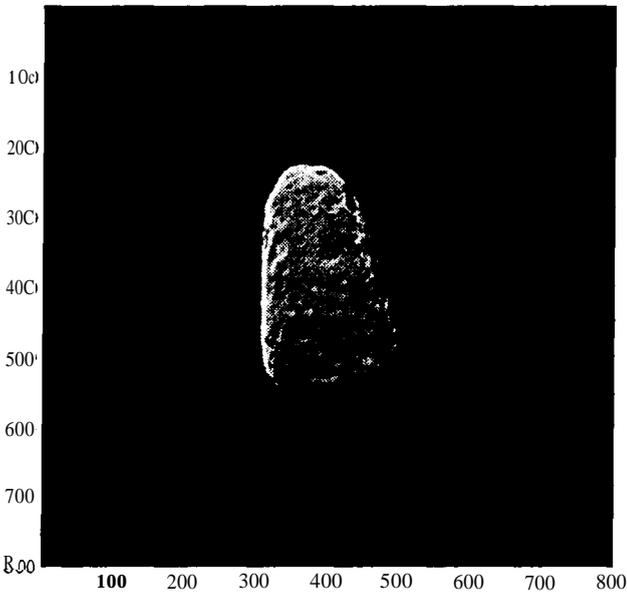


Figure 9: Simulated Asteroid Approach Frame

During the cruise phase, the optical data is typically capable of taking 250km measurements, depending on the available set of beacon asteroids. Over one week's time, that represents the capability of measuring velocity to about 0.4 m/sec, or accelerations to about 1.3 mm/sec². The IPS engine is capable of delivering a maximum of about 0.1Nt thrust, but on average will only be capable of half of that during the mission due to power restrictions. DSI has a mass of about 420kg, and therefore a typical inflight acceleration of about 120 mm/sec². The IPS engine thrust is believed to be predictable to about one percent, or about 1.2 mm/sec². It is clear then, that long-frequency signatures in the IPS performance will be barely perceptible to the optical system. These errors must be modeled, and though the capability of the model will not nearly compare to the requirements of the radio system, with a 0.1 mm/sec velocity sensitivity, and a comparable acceleration sensitivity, coping with the noise in the

engine performance will be the single most complicating factor in the OD algorithms.

The OD filtering strategy is an epoch-state, batch sequential stochastic filter. With the time-constant of the sensitivity to the expected engine performance errors on the order of a week, data batches of a maximum of a week are used. This is especially sensible since for much of the cruise periods, there will likely be only one OpNav observing period per week. The latter limitation is to reduce the on-off cycling of the engine. The data arc will typically be composed of 4 one-week data batches. The spacecraft state at the beginning of the first batch is the principal estimable parameter. Over each batch a random variation in the thrust magnitude is estimated, as well as small random accelerations. A term proportional to the solar-pressure is also an estimable parameter.

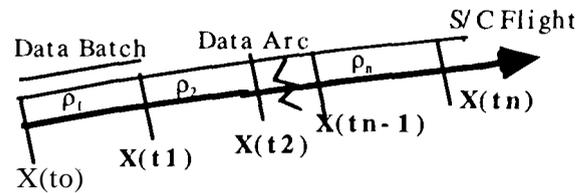


Figure 10. Schematic of Orbit Determination Data-arc Structure

Fig. 10 shows the subdivision of the data arc into batches over which an estimate parameter set is constant. $X(t_0)$ is the spacecraft state at the start of the data arc, $X(t_1)$ at the start of the second batch, etc. ρ_n is a scalar parameter describing a proportionality factor on the nominal IPS thrust magnitude in the spacecraft +Z direction. For any observation made at time t within batch one, the filter must integrate the state $X(t)$, and the state transition matrix. The later has two components, for the state itself $\partial X(t)/\partial X(t_0)$ and for the dynamic force parameters: $\partial X(t)/\partial X(\rho_1, S)$ where S is a vector of other force models, including solar pressure and small bias accelerations active across the data arc; these later model the small components of the thrust error which project in the cross directions from Z. For this observation at time t, and for subsequent observations a measurement matrix A can be formed:

$$A = \begin{bmatrix} \partial O_1 / \partial q \\ \partial O_2 / \partial q \\ \vdots \\ \partial O_n / \partial q \end{bmatrix}; \quad \frac{\partial O_{2 \times 1}}{\partial q} = \frac{\partial O_{2 \times 1}}{\partial X} \cdot \frac{\partial X}{\partial q}$$

On is the observation vector for observation n, and is a 2x1 vector, (pixel and line). The formulation of $\partial O/\partial X$ is documented elsewhere (Refs. 9, 10). q is a vector of

estimable parameters, and for batch 1, $\mathbf{q} = [\mathbf{X}(t_0), \rho_1, \mathbf{S}]$. \mathbf{A} is combined into a covariance matrix referenced to, Γ_{t_0} , via a UD factorized orthogonalization procedure (Ref. 11) an example of which is known as the Householder transformation. To process data in batch 2, an additional parameter must be added to the estimate vector, namely ρ_2 the thrust proportionality error for batch 2. Thus for batch 2, $\mathbf{q}_2 = [\mathbf{X}(t_0), \rho_1, \rho_2, \mathbf{S}]$ and the filter will integrate \mathbf{X} from t_1 to t_2 , as well as $\partial\mathbf{X}(t)/\partial\mathbf{X}(t_1)$ and $\partial\mathbf{X}(t)/\partial\mathbf{X}(\rho_2, \mathbf{S})$. The state partials for a time t in batch 2 relative to the solve-for epoch t_0 and those with respect to ρ_2 are given by:

$$\frac{\partial\mathbf{X}(t)}{\partial\mathbf{X}(t_0)} = \frac{\partial\mathbf{X}(t)}{\partial\mathbf{X}(t_1)} \cdot \frac{\partial\mathbf{X}(t_1)}{\partial\mathbf{X}(t_0)},$$

$$\frac{\partial\mathbf{X}(t)}{\partial\rho_1} = \frac{\partial\mathbf{X}(t_1)}{\partial\rho_1} \cdot \frac{\partial\mathbf{X}(t)}{\partial\mathbf{X}(t_1)}$$

And in general, for batch n , where $\mathbf{q}_n = [\mathbf{X}(t_0), \rho_1, \rho_2, \dots, \rho_m, \dots, \rho_n, \mathbf{S}]$:

$$\frac{\partial\mathbf{X}(t)}{\partial\mathbf{X}(t_0)} = \frac{\partial\mathbf{X}(t)}{\partial\mathbf{X}(t_{n-1})} \cdot \frac{\partial\mathbf{X}(t_{n-1})}{\partial\mathbf{X}(t_0)}, \text{ and}$$

$$\frac{\partial\mathbf{X}(t)}{\partial a_m} = \frac{\partial\mathbf{X}(t_{n-1})}{\partial a_m} \cdot \frac{\partial\mathbf{X}(t)}{\partial\mathbf{X}(t_{n-1})}$$

where ρ_m is an arbitrary thrust error vector from an earlier batch. When all of the data from all of the batches is combined into \mathbf{A} and Γ_{t_0} , an estimate of the parameters can be made:

$$\begin{bmatrix} \mathbf{X}_{t_0} \\ \rho \\ \mathbf{S} \end{bmatrix} = \Gamma_{t_0} \mathbf{A}' \mathbf{W} \Delta \mathbf{y},$$

$$\Delta \mathbf{y}_{1 \times 2N} = \mathbf{O}_{2 \times N} - \mathbf{C}_{2 \times N}$$

where $\Delta \mathbf{y}$ is the residual vector formed as the difference between the observation vector \mathbf{O} and the computed predicted value \mathbf{C} . \mathbf{W} is the observation weighting matrix. N is the total number of frames taken, and $2N$ is the number of data (pixel and line for each). Iterations are performed on this solution, repeating the solution one or more times with the improved integrated ephemeris and force models from the previous solution. When the solution is converged, the elements of \mathbf{p} are not equally well determined; \mathbf{p}_1 is the best determined, as all of the data in the data arc influence a measurement of \mathbf{p}_1 , whereas \mathbf{p}_n is the poorest, as only the last batch has an influence on its solution. To get the covariance to start

the next solution cycle the covariance at t_0 must be mapped forward in time:

$$\Gamma_{t_{n/2}} = \mathbf{D} \Phi_{t_0}^{t_{n/2}} \Gamma_{t_0} \Phi_{t_0}^{t_{n/2}}$$

where $\Phi(t_0, t_{n/2})$ is the state transition matrix from t_0 to the midpoint of the data arc. \mathbf{D} is a de weighting matrix to allow for errors accrued due to unmodeled perturbations.

The decision has been made to entirely reinitialize the solution process for each data arc. Operationally, this process typically has the following events:

- 1) A solution is performed for a four batch data-arc, with an epoch-state at the beginning of the first batch. This solution uses effectively no *a priori* constraint, relying on the data arc for virtually the complete state determination.
- 2) Data is accumulated beyond the last batch, into what is the "new" batch.
- 3) The estimated state from step 1 is integrated to the beginning of the second batch. This integrated state becomes the reference or epoch-state for the next solution.
- 4) A solution is made using the data in the new batch, but excluding the old (original "first") batch. The process repeats starting at step 1.

In this approach, the rationale for completely redetermining the state using the data arc only, without any pre-constraint, or forwarding of information from previous solutions is two-fold. First, there is sufficient information in a month's worth of optical data (four typical batches) to sufficiently determine the position and velocity of the spacecraft. Second, the earlier data (earlier than about a month) are sufficiently decoupled from the current data arc via the random non-gravitational accelerations so as to contribute little or no information to the solution.

Integration and Ephemeris Services

The characteristics of the spacecraft dynamic models are discussed above, but the actual mechanism used to perform the integration is a separate issue, as is the representation of the spacecraft integrated trajectory, and the ephemerides of the major and minor planets, including the encounter targets.

The numerical integrator used is a Runge-Kutta 8th order. This integration algorithm, while not computationally the most efficient available, represents the best compromise between speed and accuracy (Ref. 12). The heritage of the algorithms chosen to be incorporated into the flight Navigator was an important aspect of that decision. The precise coded version of the RK-8 chosen has a history of use in diverse orbital applications of more than twenty years. This integrator has a manually set maximum and minimum integration step size, and automatically ranges

between them based on the current level of dynamic perturbation. The accuracy achieved when operating under flight conditions, is several tens of meters over a seven-month ballistic cruise, with full dynamic perturbations in force. This comparison is against the JPL Orbit Determination Program (ODP) principal integration routine (Ref. 13) which sets the standard for deep space navigation accuracy. The RK-8 subroutine will be used to integrate the spacecraft position and the partial derivative equations for purposes of state and parameter estimation.

As stated earlier, *DSI* is a complex mission from the standpoint of expected dynamic perturbations. In order for the trajectory integrator to provide sufficient accuracy to the system, information about actual onboard propulsive activity is provided to the Navigator. This information comes from two sources, the IPS manager and the ACS. From the IPS device-manager comes a constant tally of accumulated thrust time and thrust level. By monitoring voltages and currents in the ion engine, the IPS manager is able to compute an estimated thrust magnitude. Over a span of about a minute, the IPS manager tallies this thrust, and then reports to the Navigator the accumulated thrust and time since the last message. This process continues whenever the IPS is in operation and thrusting.

The ACS also reports all propulsive activity to the Navigator, in a somewhat different manner. The ACS is constantly inducing propulsive events, but of varying magnitude compared to the IPS. In the maintenance of the spacecraft attitude, the ACS is inducing small limit-cycling turns with a frequency of roughly ten seconds when doing precision imaging (e.g. navigation observations) or tens of minutes during ballistic cruise. Additionally, ACS is responsible for implementing TCM's. These can implement several m/sec of velocity change in a matter of minutes. Every turn of the spacecraft is a propulsive event, since only in one axis (the roll -Z- axis) are the thrusters balanced, and each turn can impart roughly a mm/sec of velocity to the spacecraft. Attitude maintenance maneuvers will approximately average to zero delta-v, due to their short extent (asymmetries in the thruster performance will not however) but large turns will in general not average to zero. Even a few mm/sec when accumulated and mapped over a one month-long data arc is many kilometers of spacecraft displacement. This is very observable to the Navigator, and therefore must be tallied. During all periods of operation therefore, the ACS Velocity Estimator is monitoring ACS activity and computing accumulated velocity. When an accumulation of more than a mm/sec is achieved in any of the three inertial directions, a report is sent to the Navigator. If some fixed time, (usually 10 minutes) passes without the minimum accumulation, a report is sent nevertheless. The Navigator accumulates both types of information, and condenses it into a record of propulsive activity over the past. This record is kept for approximately five weeks,

more than enough to cover the past integration history over the longest expected data arc. The trajectory integrator then reads this record to integrate an accurate propulsive history from the epoch-state to the end of the data arc.

The planet, asteroid and spacecraft ephemerides are represented as Chebyshev function polynomials of varying order. This follows the standard representation of the planetary ephemerides in the ground navigation software. The accuracy of the stored planetary and asteroid ephemerides (relative to their generating values) is .01 km, using a 10-30 coefficient model, effective over about 5 days. The spacecraft model, with a similar representation accuracy, uses 25 coefficient representation over 1-2 day intervals.

IPS Control, Maneuver and TCM Design

Perhaps the most crucial function of the Navigator is the control of the IRS. A deep space mission has never been flown whose trajectory was not composed of long ballistic cruise segments, punctuated by planetary gravitational assists and virtually instantaneous velocity changes. This, the first deep-space low thrust mission, compounds the challenge, by requiring control of the ion engine to be performed autonomously.

The design of a low-thrust mission is a specialized technology of its own (Ref. 13), independent of the navigation function. And clearly this design process proceeds well in advance of the stage of the mission requiring autonomous navigation. The results of the design are provided to the Navigator in the form of a time-history of thrust level and direction (Figs. 3-5). The form of storage onboard of the direction profiles is by first order polynomial in time, with each week having a separate set of coefficients. The thrust levels are stored as discrete integer levels for each week.

As will be discussed below, during typical cruise operations, the Navigator will be called upon to perform weekly evaluations of the thrust profile. Part of this evaluation will be to use the current best estimated state to determine what changes to the upcoming week's thrust profile are necessary to return the spacecraft to an intersecting trajectory with the target. As discussed earlier, the changes that are possible to the designed mission trajectory are limited, due to constraints of spacecraft body orientation. Also, there is limited time to implement the mission thrust-arcs, and the existing design already uses most of the time available on the first leg, to McAuliffe. Therefore, the corrections that are possible are constrained, and represent relatively small and linear (or nearly so) corrections to the nominal designed mission.

The strategy to be used for updating the thrust profile is to treat one or more of the upcoming weekly thrust periods as an individual maneuver. Corrections to the nominal

thrust polynomial can be considered the parameters of a maneuver to be estimated. Details of the algorithm used to accomplish these corrections are recorded elsewhere (Ref. 14). Briefly, it is based on a linear estimate of control parameters, s which have varying dimension, depending on the number of adjacent control segments being adjusted. A trajectory miss vector AX is computed in the 3-dimensional encounter asymptotic coordinates. The parameters s are small changes in direction in each segment, and a change in duration of the overall burn arc. In order to obtain the solution that minimizes the corrections to the nominal thrust arc, the minimum-normal solution for s , is formed *via* the equations:

$$As = K'(KK')^{-1} AX,$$

where :

$$K' = \begin{bmatrix} \frac{\partial X}{\partial \alpha_1} & \frac{\partial X}{\partial \delta_1} & \frac{\partial X}{\partial \alpha_2} & \frac{\partial X}{\partial \delta_2} & \dots & \frac{\partial X}{\partial \alpha_m} & \frac{\partial X}{\partial \delta_m} & \frac{\partial X}{\partial \tau_m} \end{bmatrix}$$

$$\Delta s' = \begin{bmatrix} \Delta \alpha_1 & \Delta \delta_1 & \Delta \alpha_2 & \Delta \delta_2 & \dots & \Delta \alpha_m & \Delta \delta_m & \Delta \tau \end{bmatrix}$$

and

$$AX' = [\Delta B \ R \ AB \ . \ T \ \Delta t_{of}]$$

$\Delta[B \cdot R, B \cdot T, t_{of}]$ are the target relative asymptotic coordinates, representing two cross-track directions, and the along-track direction at closest approach. The solve-for parameters, $\Delta \alpha_n$, $\Delta \delta_n$, and $\Delta \tau$ are changes in a series of n thrust segment directions, and the end time of the final thrust arc. This solution is performed iteratively until converged. In this way, the solution process is actually a non-linear one, but will only succeed if a solution exists near the linear region.

As the IPS thrust arc progresses, and variations in engine performance and minor (or major) outages in thrust time relative to the nominal plan occur, the spacecraft trajectory will deviate from the designed-to nominal trajectory. The targeting strategy outlined above will return the spacecraft to the specified target conditions, but in so doing, will alter the velocity vector of the encounter asymptote. Enough of a change in this vector could cause a potential problem in maintaining the next legs of the mission to potentially Mars and WKI. If it is determined that sufficient changes to the asymptote have occurred, the trajectory will be reoptimized on the ground, and the corresponding thrust profiles will be uplinked to the spacecraft. With a redesigned mission will be a new projected mass-usage profile, associated with propellant consumption. The accuracy of this profile will effect the dynamics of the onboard integration, and therefore will be uplinked with the thrust profile.

During periods of non-thrusting, and in the twenty days before encounter conventional TCM's will be performed. These will use the IPS with the exception of the final 2 maneuvers, which will be executed using the hydrazine thrusters of the ACS. Table 5 shows the TCM schedule, with expected and associated OD errors mapped to encounter at each TCM for the final 20 days of approach to McAuliffe. The algorithm used to compute these maneuvers is the same as used for the IPS control algorithm. Necessarily however, the maneuver solution is for only three parameters: the three components of delta-velocity. Another important difference between a RCS TCM and an IPS control, is that the former occurs in a relatively short period of time; whereas IPS controls can take hours or days. In most cases the applied maneuvers are expected to be small, on the order of one m/s or less, which for the IPS will still be a maneuver of less than two hours.

Time to Encounter	Range to McAuliffe (km)	Downtrack Error (km)	Crosstrack Error (km)
-20d	12.6E6	570	660
-10d	6.3E6	138	27.3
-5d	3.1E6	69	5.5
-2.5d	1.6E6	54	2.5
-1.5d	0.0E6	44	1.5
-1.0d	630E3	42	1.2
-12h	315E3	40.2	0.89
-6h	157E3	40.1	0.25
-3h	72E3	40.1	0.20

Table 5: Approach TCM Schedule with Associated OD Performance Statistics

The nature of the bus-body illumination constraints has been discussed earlier, as has the need to constrain the direction of TCMS accordingly. The need to perform maneuvers in any direction of the sky persists however, as statistical variations in the orbit process do not observe the constraints of onboard instruments. Any direction of propulsive maneuver (using either RCS or IPS) can be accomplished by vectorally splitting the maneuver into two parts, whose vector sum equals the original design (Fig. 11). Through interaction with ACS, the Navigator determines if a particular maneuver request is allowed, and if not, decomposes the TCM into two parts. The precise nature of the interaction necessary to accomplish this will be discussed below.

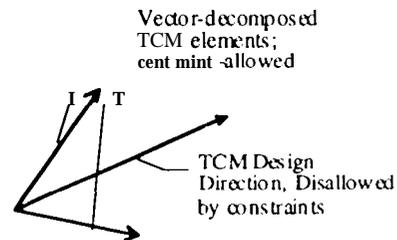


Figure 11. TCM Vector-decomposition

Although still under development, it will likely be necessary to provide an opportunity to perform an emergency deflection maneuver within the last six hours of approach. There is substantial uncertainty about the albedo and therefore the size of **McAuliffe**. Since the planned flyby altitude is so small (5km) the spacecraft could be jeopardized in the event that the body is greater than 3 times its estimated size. Therefore, in the final 6 hours of approach, the Navigator will periodically perform a crude estimate of body size based on incoming image data. If the estimate shows a major axis greater than an acceptable safety margin, the Navigator will make an emergency change of the target conditions, and use a contingency TCM opportunity to deflect the spacecraft trajectory.

Late Knowledge Update

The final control of the spacecraft trajectory will occur at about 6 hours prior to encounter. Subsequent to that maneuver, the full navigation picture processing and OD estimation process will be in force. But at approximately 30 minutes from closest approach, normal navigation operations will cease. Because of the very short timescale of activities at encounter, the Navigator must initiate simplified processes. The principal technical feature that enables the simplified processes is the fact that for the final few minutes of the approach, the Navigator can acquire no additional useful information about the velocity of the spacecraft. This being the case, the data filter reduces dramatically to a 3-state estimate of instantaneous **spacecraft position only**. The estimates occur from picture to picture, and each solution is conditioned by the covariance obtained from the previous picture. Over so short a time-span, the absence of any process noise, or other attenuation of the accumulating information does not pose difficulties due to **mismodeling**. This is especially true due to the rapidly increasing power of the data as the spacecraft approaches; any modeling errors in previous images would be overwhelmed by the later pictures. The picture processing used during this final stage of the approach has been discussed briefly above.

OPERATION OF THE NAVIGATION SYSTEM:

The operation of the Navigator, though largely an autonomous function, is managed in a gross sense by ground commands. These commands are imbedded in a conventional stored sequence. Typically, a ground directive is given to the Navigator, followed by a period of uncommitted time in which the Navigator is allowed to perform autonomous action. Following are detailed descriptions of the major Navigation actions.

Navigation Imaging Opportunity

The simplest period of activity during the mission is ballistic cruise (non-powered cruise). During this period of time, the only regular navigation operations that occur

are the taking and processing of navigation frames. Such an event is triggered by a **Nav-Photo-Op** spacecraft command. Though this operation happens during all phases of the mission, it will be discussed here in the context of a non-thrusting (ballistic) portion of the trajectory. For most of the mission, this operation will occur once per week. At one point in the sequence, a **Nav-Photo-Op** directive is issued to the Navigator by the ground-generated stored sequence. Associated with this command, is a period of time allocated to the Nav function to accomplish picture planning, execution and processing.

Before the **Photo-Op** session begins, it is the ground system's responsibility to put the spacecraft in a state that is possible to command turning and imaging operations. This preparation activity includes turning the camera on, and changing whatever camera states are necessary, and doing so with sufficient lead time to insure readiness when the **Photo-Op** begins. If any ACS states need setting, this must also be done. Additionally, the ground must insure that no operations occur which conflict with imaging and turning commands during the extent of the **Photo-Op**.

Very little information is necessary to pass to the AutoNav system with this directive, but it is necessary to inform Nav how much time is available to obtain its images. When the "Nav-Photo-Op" directive is issued, the following operations take place:

- 1) Nav determines what the current attitude of the spacecraft body is, in order to be able to return to that attitude after imaging if requested. Otherwise, ground operations can specify a different terminal attitude.
- 2) AutoNav identifies the set of navigational targets that are appropriate for the current time of the mission.
- 3) A target is selected, in order, from the list starting at the beginning of this period. Each of the lists has been optimized so as to minimize the extent of the turns between targets.
- 4) Nav determines from ACS how long a turn from the current attitude to the requested attitude will take. Additionally, The ACS planning expert is asked how long it will take to turn from the target attitude to the *a priori* attitude. If the sum of these is less than the time remaining in the AutoNav session, then the sequence of operations continues, other wise a branch to the end procedure (step 7) commences.
- 5) AutoNav prepares a small file onboard which contains a sequence. This sequence effects the following operations:
 - An ACS commanded turn to the OpNav target,
 - A Pause for sufficient time to finish the turn,
 - A Command to take an image - with automatic notification of image completion being sent to AutoNav.
- 6) AutoNav then launches the constructed sequence file, using one of the eight available sequence strings. With receipt of the image complete notification from the launched sequence, the main **Photo-Op** events continue,

with a branch back to event 3) and a selection of the next target in the list.

- 7) Begin the termination process for the Photo-Op, with the construction of a minisequence to turn the spacecraft back to [he starting or other requested attitude.
- 8) Launch of the final turn mini-sequence, and this marks the end of Photo-Op.

IPS Control:

During the months of continuous thrusting, there are periods of time when the IPS must be shut down for short periods. These interruptions include time for navigation data taking, for downlink of data, and possibly for technology validation experiments. Also, on a regular basis, at least once per day, the direction of the engine thrust must be updated by the AutoNav system.

As with the Nav-Photo-Op directive, use of the commands to enable the AutoNav system to operate the IPS, require the ground operating system to prepare the spacecraft for the autonomous operation of the navigation system. In the case of a "NAV-SET-IPS" command, the ground generated sequence turns on and otherwise conditions the IPS engine. From a cold start, there is a considerable amount of preparation necessary, taking up to an hour. However, since these activities are well known, repetitive, and well calibrated in terms of time required, the mission operations team uses a fixed sequence, called a "block" and as part of normal invocation of Navigator, this will be routinely done.

To begin autonomous IPS operations then, the ground first issues the "IPS-PREPARATION" **block command** leaving the ion engine in a state ready for the AutoNav system to issue a simple "thrust-on" command. Then, after leaving sufficient time in the sequence to complete the preparation cycle, the sequence issues a "NAV-SET-IPS" command. In response to this command, the AutoNav system begins a series of tasks:

- 1) A computation is made of the necessary thrusting over the next day. The direction of engine is determined, as is the duration of the burn.
- 2) The ACS planning expert (APE) is queried to determine the length of time required to turn the spacecraft to the desired position.
- 3) A mini-sequence is constructed to accomplish several tasks:
 - Turn the spacecraft to the desired direction
 - A delay necessary to guarantee completion of the turn.
 - A directive to the IPS manager to turn on the thrust grids of the ion engine, and to leave the thrust on for a maximum of 1 day, or for a shorter duration if specified.
- 4) The mini-sequence is launched.

Other than at the end of a long mission burn, the duration of thrust given to the IPS manager will be one day. To accomplish the necessary updates to the thrust vector, the

ground-generated sequence will include periodic requests of AutoNav to update the direction. Although it would be possible for the AutoNav system to autonomously provide update vectors, in order to do so, AutoNav would have to become aware of other scheduled events on board the spacecraft which would cause a change in the status of the engine, such as telecommunication events. Since it causes little impact on the ground system to issue the NAV-UPDATE-IPS command, AutoNav will rely on this method. On receipt of this command, the Navigator will construct and launch a new minisequence to update the thrust direction and duration. These directives will go to the ACS attitude commander and IPS manager respectively.

At the end of a mission thrust segment, the navigator will, in response to an IPS-UPDATE command, issue a directive to the IPS manager with a thrust duration of less than a day. The IPS manager will keep track of the amount of time that the IPS has been thrusting since a thrust-on or thrust-update directive from Nav, and if this duration is met, the manager will shut down the IPS.

As stated earlier, the timings of events that shut down the IPS, such as navigation picture taking and telecom sessions is not known *a priori* onboard by the Navigator, being carefully scheduled by the ground. Therefore, the AutoNav system must cope with the otherwise unscheduled shut-down of the engines at any time. This is accomplished *via* the design of the IPS control, involving continued monitoring of the accumulated thrust from the engine. At any time, the Navigator is prepared to evaluate the thrust accumulated thus far, and to thereby reevaluate the necessary duration of thrust given to the IPS manager in a command. Therefore, the ground control system may shut down the engines at any time, and the Navigator will adjust to the circumstance.

Such a shutdown is simply implemented. The ground-generated sequence commands the thrust to turn off, then commands the engine to whatever shut-down state is required. The Navigator is made aware of the shutdown implicitly via the lack of "engine-on" status messages from the IPS manager.

Trajectory Correction Maneuvers:

With conventionally navigated spacecraft, the implementation of a TCM required a major effort for the ground control team. With the AutoNav system, ground control is relieved of all responsibility for the TCMS except for scheduling. Much as with the OpNav image taking, the ground merely schedules a time-gap in the sequence in which the AutoNav system may place its autonomous operations. In this case, the operations are to turn the spacecraft and operate the engines: either the RCS thrusters or the IPS.

During an extended mission burn, no dedicated TCM's are necessary, as continual corrected control is taking place. However, after a mission burn, or during a ballistic cruise, or especially on approach to an encounter target, dedicated opportunities to correct the trajectory are advantageous, and can be scheduled frequently with no additional ground costs. For *DSI*, it is anticipated that the spacecraft travel no more than a month between TCM opportunities, and that they occur much more frequently on approach to a target, as has been discussed earlier.

The ground implementation of a TCM is as follows. Previous to issuing any command to the Navigator, ground operations must insure the readiness of the RCS system or the IPS (or both), depending on which is to be mandated. Obviously, if the navigator will be given the option of using either. Such preparations might include turning on the IPS, or activating the TCM RCS thruster heaters. When the preparations are complete the ground-generated sequence issues a NAV-PERFORM-TCM command. This begins a series of activities:

- 1) The Navigator will perform an orbit determination calculation based on the latest data, to determine the current spacecraft state and its propagation to the encounter target.
- 2) The velocity change necessary to take the spacecraft to the target is computed.
- 3) The ACS vectorizer is queried as to whether this TCM needs vectorization, and if so, what are the components into which it can be broken down, (Fig. 11).
- 4) The APE is consulted as to the extent of time required to implement the turn(s).
- 5) The Navigator constructs a mini-sequence to accomplish a series of tasks:
 - A: Direct ACS to turn the spacecraft to the requested attitude,
 - B: Wait the required amount of time to implement the turn,
 - C: Direct ACS to implement the delta-v.
 - D: If an unvectorized turn, proceed to E, otherwise, complete steps A through C for the second leg of the TCM,
 - E: Direct ACS to turn back to the *a priori* attitude, or a requested terminal attitude.
- 6) The Navigator then starts the mini-sequence, to accomplish the above activities, and this completes the implementation of a TCM.

These activities are constrained to take place in a given amount of time. This constraint is enforced by two methods, first by a hard limit in the total length of time provided in the sequence. If the Navigator hits this limit in constructing its mini-sequence, this constitutes an error. To prevent this error from occurring, the Navigator is initially constrained from implementing TCMS of greater than a certain magnitude. The magnitude of this limit will correspond to a 3-sigma maximum expectation value of statistical delta-v. If this limit is surpassed, the Navigator will implement the maximum magnitude in the

computed direction. The allocated sequence time will correspond with this expected maximum time with some additional appropriate buffer.

Encounter Operations:

The precise activities of the *DSI* encounter will be determined well in advance of the encounter itself. These operations will be encoded into a series of sequences stored onboard the spacecraft, and triggered into operation by the Navigator.

At least for the McAuliffe encounter, the dependence of the scheduled sequences upon the precise location of the spacecraft relative to the target does not become strong until the last five minutes of approach. The important dimensional dependence is upon the down-track dimension, as this direction remains poorly determined until very late. Consequently, the final approach sequence is subdivided into 4 short sub-sequences, each with increasing sensitivity to time-of-flight (down track position) errors, and each positionable with greater accuracy by the Navigator.

For the approximately five hours following the final TCM, previous to the start of the McAuliffe-Encounter operations, images are being taken by the spacecraft and passed to the Navigator for processing. Throughout this "Far Encounter" period, the Navigator is updating its estimate of the spacecraft encounter coordinates, including the time of closest approach (TCA.) Since the timing of these events is not dependent upon a precise determination of TCA, these can be scheduled in the sequence in a completely deterministic way.

The first of the asteroid encounter sequences (AE 1) begins 260 seconds before closest approach at a range of about 2000 km. The first action of this sequence is to take an OPNAV image, at E-240sec. This image is immediately sent to AutoNav for processing. As the science activities of the encounter sequence proceed, the AutoNav system is reducing the data and obtaining a new encounter state estimate. The science activities of AE1 will include infrared and ultra-violet observations of McAuliffe. Since the combined processes of data readout, image analysis, and state estimation take approximately 12 to 15 seconds, there is time in AE1 for the Navigator to process several pictures if the science sequence allows. Each update of the target-relative ephemeris is automatically reflected in improved pointing accuracy. This is so because the ACS system is regularly querying the Nav system for the latest ephemeris information. All science observations are specified as target relative (vs. absolute inertial directions) and thus are improved in accuracy whenever the Navigator improves the accuracy of the ephemeris. It should be emphasized again however, that once the sequence is started, the time of a specified event is deterministic and cannot change. AE1 will end at E-175sec.

The second encounter sequence (AE2) will begin at about 160 seconds before closest approach. As with AE1, the first action of the sequence will be to take an OPNAV image, in this case, at about E-155 seconds. There is a gap of about 15 seconds between AE 1 and AE2 which will allow the Navigator to move the start point of AE2 to correspond to updated estimates of the time of closest approach. As with AE1, there will be opportunities for multiple OpNav pictures to be taken and processed, and the estimated spacecraft ephemeris updated before the end of AE2 at E-90 seconds.

The third encounter sequence (AE3) will begin at E-80 seconds, and as previously, the first activity is to take an OPNAV image at E-75 seconds. Additional OPNAV images may be taken in AE3 using the other visual frequency imaging system, the APS (Active Pixel Sensor), before the sequence ends at E-40 seconds.

The final encounter sequence (AE4) begins at E-35 seconds. The final OPNAV image is taken with the CCD sensor at E-33 seconds, and the final target-relative ephemeris is made available to ACS at about E-23 seconds. From this time until the spacecraft can no longer accelerate its slew-rate to keep the target tracked, at about E-15 seconds, science images with the APS and CCD will be taken. Even when this limit is reached, several images may still be taken over the next few seconds, as the asteroid (then over three CCD fields of view in apparent diameter) sweeps out of view. AE4 will continue taking IR images of the asteroid as it sweeps out of view, and turn the spacecraft to view the retreating asteroid on departure. This turn should be complete within about a minute, whereupon science imaging (but no OPNAV imaging) will continue, until AE4 ends at E+240seconds.

Following AE4, conventional deterministic sequencing will resume, with final science views of the asteroid. Within five days or so, AutoNav operations will also resume, with periodic beacon-asteroid images, and autonomous control of the IPS.

PRELIMINARY SIMULATION RESULTS

Although the development of the navigation flight system is not yet complete, preliminary simulations have been run with the software to assess its performance. This simulation uses the current baseline trajectory obtained from mission design, which assumes a launch on July 1, 1998 and flyby of the asteroid McAuliffe on January 16, 1999. Covariance analysis was performed on the last 30 days of this cruise prior to asteroid encounter to determine OD performance in both an interplanetary cruise and small body flyby scenario. The analysis assumes no a-priori knowledge on the state at the E-30 day epoch. Data scheduling during this time frame is shown in Table 6. Note that up to around E-12 hours, observations are taken

of multiple beacon asteroids to fix the heliocentric spacecraft trajectory. Subsequent observations up to the encounter are solely of the target asteroid to accurately determine the target-relative spacecraft state, in particular, the time-of-flight or downtrack component.

The resulting performance is graphically displayed in Figures 12 and 13. These show the semimajor and semiminor axes of the 3-dimensional positional uncertainty ellipse mapped to the encounter as a function of time. Figure 12 shows the dramatic improvement in position knowledge in all three dimensions gained from the data from E-30 to about E-7 days. The largest dimension of the ellipse has a value of about 70-80 km at this time, and represents the best knowledge of the downtrack uncertainty of the spacecraft position relative to the target obtainable from the beacon and target asteroids. The two other dimensions of the ellipse however, have about the same values and are an order of magnitude better than the largest component. This is due to excellent crosstrack information obtained from observing the target asteroid with optical data. By the time of encounter, these components will be known to the 100-200 m level.

Time to Encounter (days)	# of observations	IAU Catalog # of asteroids used
29	13	5,15,46,126,132,163,183,270,313,398,696,1036,3352
22	13	5,15,46, 126,132,163,183,270,313,398,696, 1036,3352
15	12	5,15, 126,132,163,180,183,270,313,398,1036,3352
13	13	5,15, 126,132,163,180,183,270,313,398,1036,3352
10	12	5,15, 126,132,163,180,183,270,313,398,1036,3352
8	12	5,126,132,163,180,183,270,313,347,398,1036,3352
6	12	5,126,1 32,163,180,183,270,313,347,398,1036,3352
4	12	5,1 26,132,163,180,183,270,313,347,398,1036,3352
3	12	5,126,132,163,180,183,270,313,347,398,1036,3352
2	12	5,126,132,163,180,183,270,313,347,398,1036,3352
1	12	5,126,132,163,180,183,270,313,347,398,1036,3352
0.4	12	5,1 26,132,163,180,183,270,313,347,398,1036,3352
0.4 -0.0	39	3352

Figure 13 shows an expanded view of the last hour prior to encounter. Note that the semimajor axis of the uncertainty ellipse (representing the downtrack error) which had not shown much improvement from E-7 days has a sudden dramatic drop at about E-1 hour. This is caused by the changing geometry as the spacecraft flies by the asteroid. The cross line-of-sight measure of the spacecraft position relative to the target is rotated into the downtrack direction, thereby improving the estimate of

this component. This clearly illustrates the need for late observations of the target, and why it would be impractical to process this important data on the ground due to light-time considerations. Only by processing this information onboard can the improved knowledge from late observations be taken advantage of for science purposes.

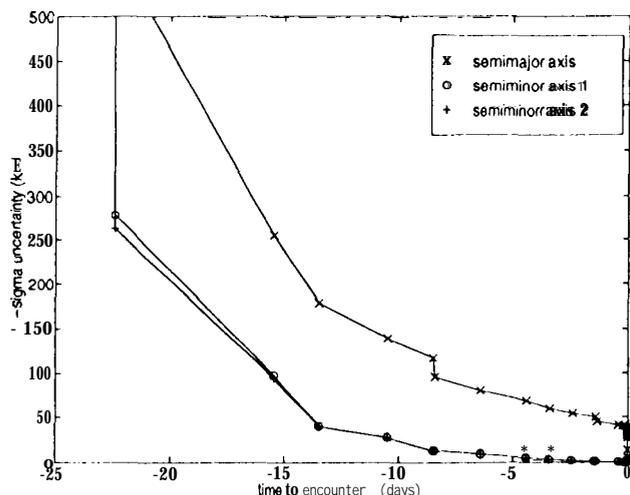


Figure 12: Autonomous Navigation System Orbit Determination Performance, Far Encounter

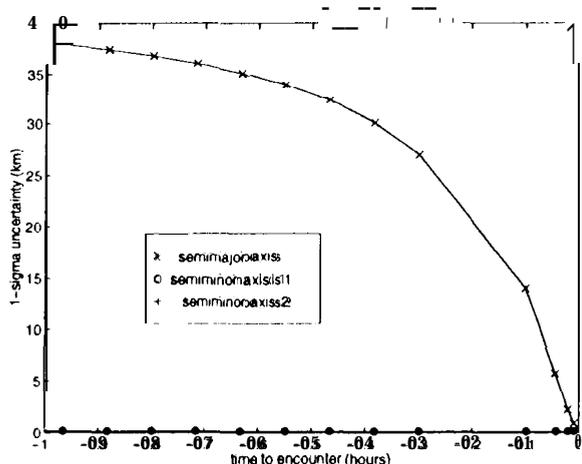


Figure 13: Autonomous Navigation System Orbit Determination Performance, Near Encounter

ACKNOWLEDGMENT:

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

Special thanks are due to a number of individuals at JPL: Marc Rayman, the *DSI* Mission Engineer, under whose auspices much of the *DSI* Navigation development has occurred, Abdullah Aljabri who lead the *DSI* Flight Software Team during its most important formative

stages, Steve Williams and Dennis Byrnes for leading the mission search and design process and Ted Drain for consultation and assistance in implementation details of the flight system, especially the ephemeris server subsystem.

REFERENCES:

1. Riedel, J. E., Bhaskaran, S., Synnott, S. P., Bellman, W. E., Null, G. W.. "An Autonomous Optical Navigation and Control System for Interplanetary Exploration Missions", IAA paper IAA-L-0506, Proceedings, Second IAA International Conference on Low-Cost Planetary Missions, April 16-19, 1996 Laurel, Maryland, USA.
2. Rayman, M. D., Lehman, D. H., "NASA's First New Millennium Deep-Space Technology Validation Flight," IAA paper IAA-L-0502, Proceedings, Second IAA International Conference on Low-Cost Planetary Missions, April 16-19, 1996 Laurel, Maryland, USA
3. Coverstone-Carroll V., Williams, S. N., "Optimal Low Thrust Trajectories Using Differential Inclusion Concepts", AAS/AIAA Space flight Mechanics Meeting, Cocoa Beach Florida, USA, Feb. 1994
4. Lisman, S., Chang, D, Hadaegh, F., "Autonomous Guidance and Control for the New *Millennium DSI* Spacecraft" AIAA paper 96-3817, July 29-31 1996, San Diego, Ca., USA.
5. Stolper, S. A., 'Embedded Systems on Mars: The Mars Pathfinder Attitude and Information Management (AEM) Subsystem', Proceedings, Embedded Systems, Conference East, Vol II, Miller Freeman, Inc, San Francisco, 1997.
6. Hog, E.; "The TYCHO Project on-board the HIPPARCOS Satellite," IAU Symposium No. 141, "Inertial Coordinate System on the Sky", pg. 307-310
7. Bhaskaran, S.; Riedel, J.; Synnott, S.; "Autonomous Optical Navigation for Interplanetary Missions,"
8. Vaughan, R. M., Riedel, J. E., Davis, R. P., Owen, W. M., Synnott, S. P., "Optical Navigation for the Galileo Gaspra Encounter," AIAA paper 92-4522, AIAA/AAS Astrodynamics Conference, Aug. 1992, Hilton Head, SC, USA
9. Riedel, J. E., Owen, W. M., Stuve, J. A., Synnott, S. P., Vaughan, R. M., "Optical Navigation During the Voyager Neptune Encounter", AIAA paper 90-2877-CP, AIAA/AAS Astrodynamics Conference, Aug 1990, Portland, OR.

10. Synnott, S. P., Donegan, A. J., Riedel, J. E., **Stuve**, J. A. "Interplanetary Optical Navigation: Voyager Uranus Encounter," **AIAA** paper 86-2113, **AIAA** conference, Aug. 1986, Williamsburg Va.
11. Bierman, G. J., *Factorization Methods for Discrete Sequential Estimation* (Academic Press, New York), 1976
12. Battin, R.; "Resolution of Runge-Kutta-Nystrom Condition Equations through Eighth Order", **AIAA** Journal, Vol. 14, August 1976, pp 1012-1021.
13. Ekelund, J; "The JPL Orbit Determination Software System", **AAS** paper 79-111, **AAS/AIAA** Astrodynamics Specialist Conference, Provincetown, Mass, USA, June 25-27, 1979
14. Desai, S.; Bhaskaran, S.; Bellman, W.; Halsell, C.; Riedel, S.; "The Autonomous Navigation System: Autonomous Control of Low Thrust Propulsion Systems", **AIAA** paper 97-3819, **AIAA** Guidance, Navigation and Control Conference, Aug. 1997, New Orleans, USA.