Optical Navigation Plan and Strategy for the Lunar Lander Altair

Joseph E. Riedel, Andrew T. Vaughan, Robert A. Werner, Tseng-Chan Wang, Simon Nolet, David M. Myers, Nickolaos Mastrodemos, Allan Y. Lee, Christopher Grasso[1], Todd A. Ely, David S. Bayard

Jet Propulsion Laboratory, California Institute of Technology,
4800 Oak Grove Dr., Pasadena, CA. 91109
joseph.e.riedel@jpl.nasa.gov
[1] *Blue Sun Enterprises, Boulder, Colorado, 80302*

Abstract

This paper reviews the currently planned Altair Optical Navigation (OpNav) system. The discussion includes description of the OpNav camera manifest. The Altair OpNav plan envisions one, OpNav camera assembly, with perhaps a functional backup that includes a wide angle-imager (of 40° to 60° field of view – FOV), and a narrow angle imager (of 1 to 3° FOV) co-mounted on a 2-degree-of-freedom gimbal. Both imagers are assumed to be relatively wide aperture and large dynamic range to provide excellent short-exposure images at mid-latitudes, and adequate images of longer-exposure near the poles. Landmark modeling and tracking methodology is discussed, including the stereophotoclinometry method assumed to be used to obtain high-accuracy terrain maps at lunar landing sites of 1-2m, and 50-100m elsewhere, using the images expected to be obtained from the Lunar Reconnaissance Orbiter (LRO). Characteristics of the OpNav navigation system are discussed and architecture and results from landing simulations presented, showing expected landing accuracies of better than 10m.

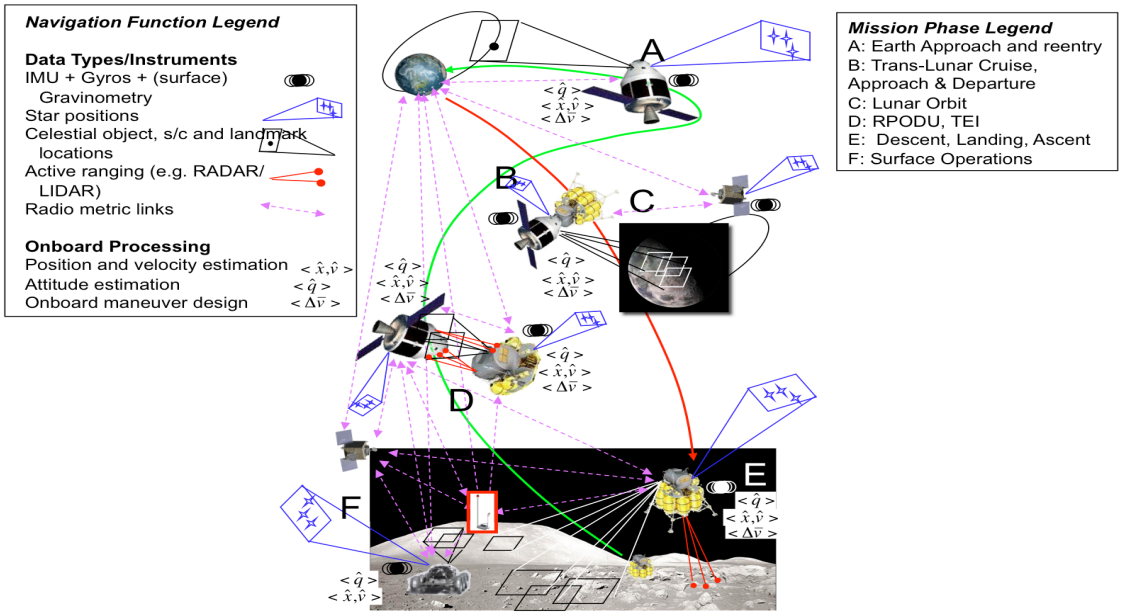

Figure 1. Overall GN&C Schematic Architecture for the Constellation Lunar Missions; *Passive Optical Navigation plays a core role in this architecture, with radio-link-independent onboard data types – such as OpNav - carrying a cruicial performance and safety resoponsibility in all mission phases.*

I. Introduction

Forty years ago, as the Apollo missions navigated their way to lunar landings, they used onboard automated navigation systems that relied primarily upon Earth-based-determined states, and inertial measurement unit (IMU) accelerometer readings; but they also made minimal use of primitive optical navigation. This early use of that data type primarily consisted of crew-member sightings of known targets, and readings of their locations from reticles overlaid on the view-ports. In orbital cruise, such measurements proved to be very powerful, but extremely cumbersome to obtain. During landing such observations were critical to accurate and safe landing but were often extremely difficult to obtain due to crew-member's need to identify previously studied landmarks from new view angles and illumination conditions. When the Constellation Program returns to the Moon with the Altair lander, much greater use of optical navigation will be made. This is because of the advancement of computer technology in the intervening five decades, subsequent mapping missions to the Moon and the development of highly accurate terrain maps, and substantial use of optical navigation aboard unmanned deep space probes, which technology will be applied to Altair.

The Optical Navigation (OpNav) system aboard Altair will consist of dedicated cameras, both wide and narrow angle, that will be gimbal-mounted, to maximize observation flexibility (choice of surface landmarks), and to eliminate the need for turning the spacecraft to obtain OpNav data, which was the case for Apollo. In addition, gimbaling the camera assembly provides for the possibility of re-pointing the instrument for secondary back-up functions, such as star-tracking for attitude determination and rendezvous and proximity phase navigation around the Orion vehicle. Optical images will be processed onboard probably by a dedicated CPU. The nature of this image processing varies depending on the phase of the mission. In Earth-Moon cruise, especially soon after departure, the OpNav system may utilize the star-like images of geostationary or MEO communications and navigation satellites. Looking forward to the Moon, distant or wide-angle lunar images will use limb observations. Larger images of the Moon on approach, in orbit and during descent and ascent will use individual landmarks. A key portion of the OpNav processing is the modeling of and extraction of landmark positions to obtain precise geometric navigation knowledge, and in fact this element of the system is by far the greatest consumer of CPU resources in the navigation system. On ascent, images of Orion will initially be unresolved and star-like, but on approach the Orion vehicle will be mapped using landmark models, exactly as on the surface of the Moon.

Upon obtaining the landmark, limb or satellite image-positions, these data are passed to a general onboard navigation system. But unlike the Apollo navigation system that, at least in the landing case, depended upon an accurate initial state and IMU propagation to determine the absolute spacecraft position, each OpNav observation contains an immediate two- and usually three-dimensional geometric determination of position (GDoP). Heuristically, two consecutive pictures, whose data is differenced, provide velocity information, and thus in a matter of seconds, the 6-dimensional position and velocity space of the spacecraft is spanned. This characteristic of OpNav data can and will influence the structure of the navigation filter onboard, allowing for simplification of the filtering strategy, shortening of data arcs and other aspects that augur for robust performance.

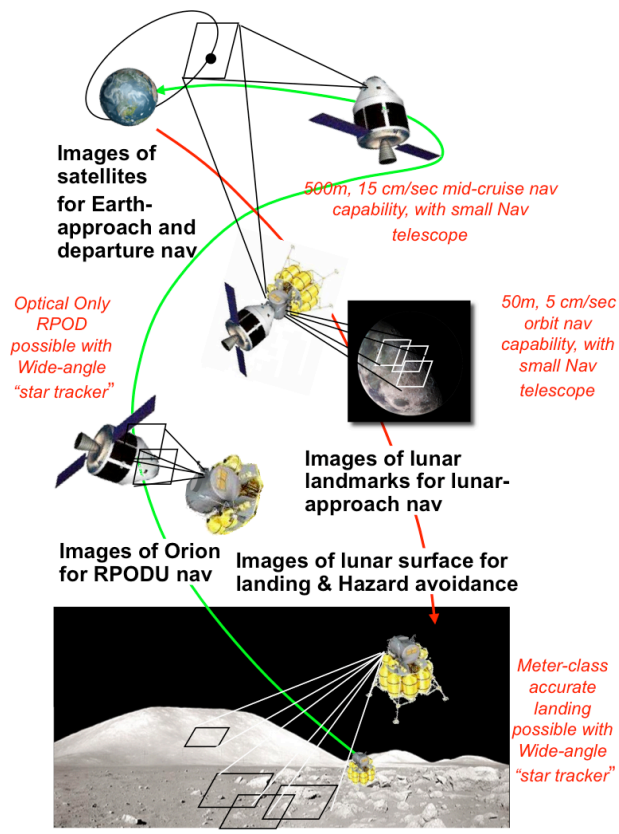

Figure 2. Constellation/Altair Optical Navigation Architecture; *OpNav will be applied throughout all phases of the Altair mission, and will provide high precison navigatin independent of any radio link – accuracies reflect instantaneous kinematic position determination, and do not include effects of non-gravitational or other modeling errors*

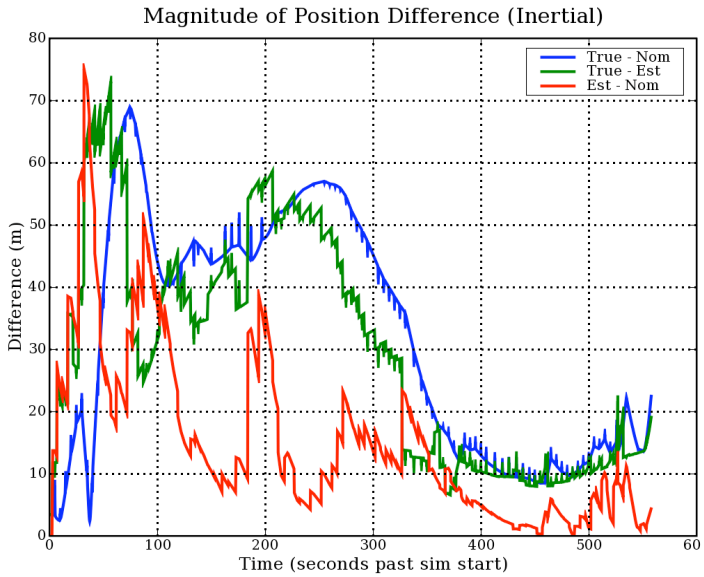

Figure 3. Performance of the AutoGNC Onboard Autonomous Navigation System During an Altair Powered Descent and Landing Simulation; *The plots show the differences between the True, Nominal (planned) and AutoGNC Estimated trajectories. Maximum performance of about 10m knowledge and control is obtained about 70s prior to touch-down at which time the OpNav data loses contact with the surface due to the extreme illumination conditions in the Winter-time Shacleton region*

Beyond formulation details of the Altair optical navigation system there are numerous important architectural and system issues of a navigation system that incorporates OpNav data as planned for Altair, as the system is actively responsible for planning and executing mechanical activations such as gimbal slews, camera shutters, and potentially attitude changes. Ideally, these functions should be combined and coordinated with more general GN&C actions related to image and attitude control. Later sections will describe a general command and sequence strategy based on that of the Phoenix Mars lander that is being used to prototype an Altair navigation system that could be generalized as an overall vehicle command and sequence system. One key aspect of such a system is to handle fault cases robustly, and that is especially important for the navigation system that must, to the maximum extent possible, diagnose and resolve its own problems reliably. Such functions are normally considered beyond the scope of typical navigation problems, but such functions and systems will ultimately determine the reliability of the navigation system itself, and thus the be key contributors to the success probability of the Altair missions and even the safety of the crew.

One of the principal simplifications offered by the use of OpNav is the fact that complex perturbing forces and factors, for example on-board random accelerations, local gravity field perturbations, terrain map frame-tie errors, need not necessarily be comprehensively modeled, as the immediate GDoP provided by adjacent frames can quickly reduce such errors if the current state estimate is adequately (though not necessarily completely) decoupled from the past data and acceleration history. For the most important navigation case for Altair, landing, the most critical state for the navigation system to determine is that relative to the landing site, and that is precisely the content of the GDoP of each of the last OpNav images taken in the final landing sequence.

Upon ascent, the Altair OpNav plan includes a nominal function to once again perform surface relative navigation to match orbit with Orion. On approach to Orion OpNav will provide an auxiliary capability to the coherent radio link between the two vehicles, allowing Altair to approach with sufficient accuracy to utilize the rendezvous LIDAR, and OpNav will additionally provide backup for this instrument providing a means of determining the relative position velocity and attitudes of the vehicles. Expected performance figures of the OpNav system during the rendezvous and docking will be presented below.

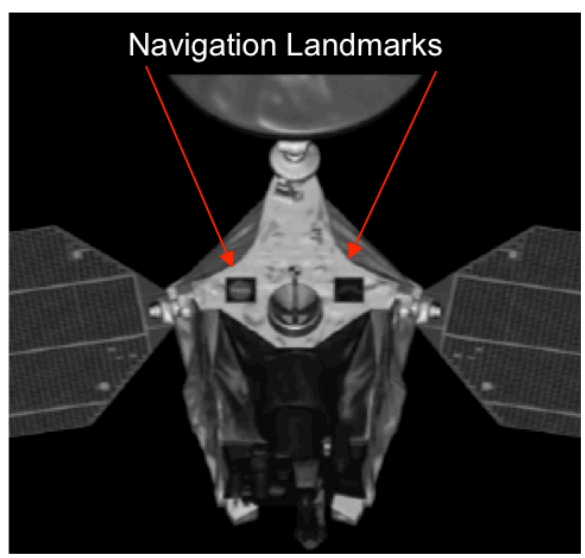

Figure 4. A Simulated Image of MRO Being Used for RPOD Optical Navigation with the Use of "Landmarks;" *Onboard the vehicle, only the "landmarks" need have associated "terrain" models provided for the OpNav system*

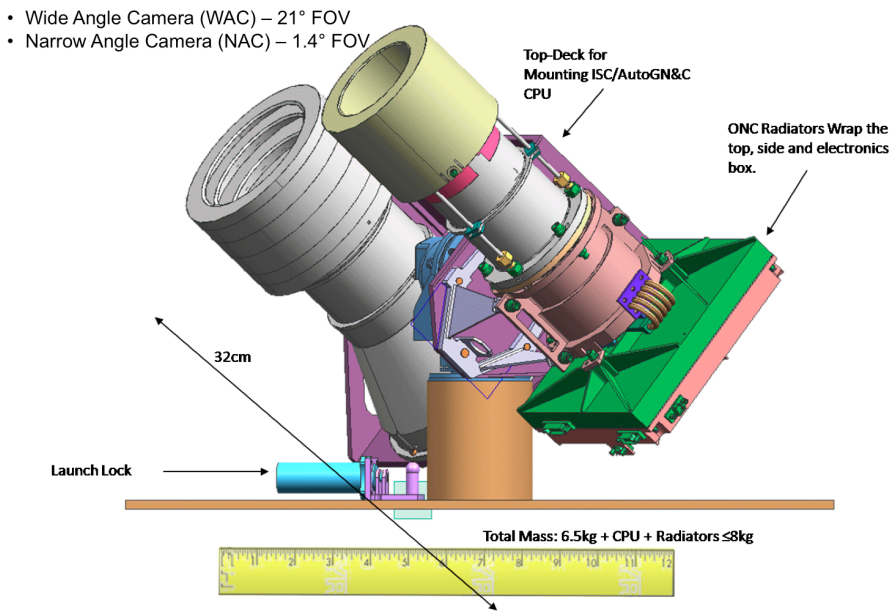

Figure 5. The Optical Navigation Sensor System (ONSS), A Gimbaled Two-Camera Sensor; *This sensor concept is composed of the MRO Optical Navigation Camera (ONC) and the ST6 Inertial Stellar Compass (ISC)*

II. OpNav Architecture and the Optical Navigation Sensing System Instrument of Altair

A. Rationale for the use of optical navigation aboard Altair

Fig. 1 shows the navigation architecture of the Constellation program, including that of both Orion and Altair, and Fig. 2 specifically focuses on Optical data. The Constellation Program (Cx) has wisely required that crewed Cx flight elements be capable of participating as necessary in returning the crew to Earth in the absence of a communications link to the Mission Control Center (MCC). The way that this has been interpreted by the Altair GN&C team has been to maintain that the only likely such scenario is a corruption of the radiometric signal from Altair (and Orion simultaneously) such that conventional Earth-based radio-based navigation is not possible. It is quite possible that radiation from a solar storm could cause such an event (Ref., 1). However, though it is virtually impossible for voice link to be completely lost with Altair and Orion, it is not reasonable to rely on voice transmission of sufficient data to initialize the navigation system aboard Orion to return the crew safely to Orion under exigent conditions. Therefore Altair will be equipped with the ability to independently determine its position without reference to Earth. The only practical means of doing this is with optical navigation, where with cameras Altair will determine its position relative to reference targets such as surface landmarks on the Moon. But there are other potential targets.

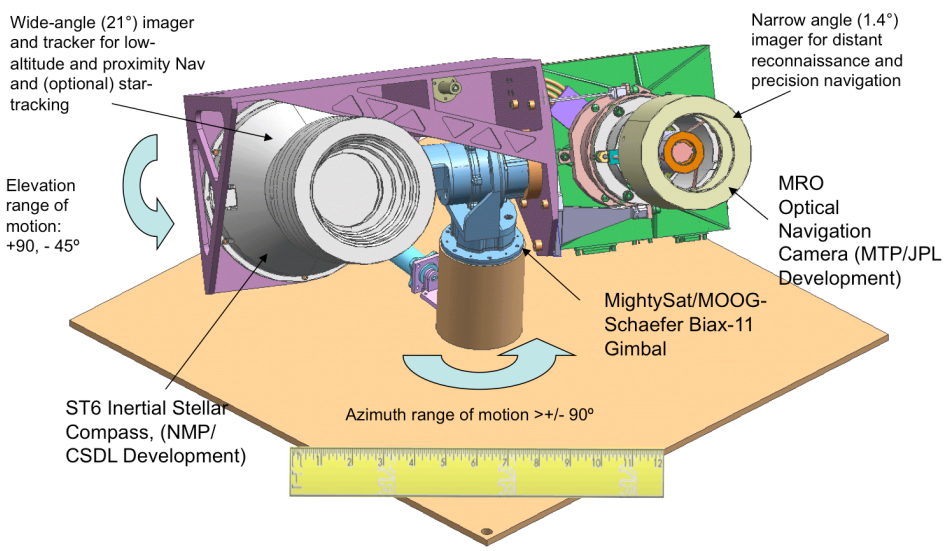

Figure 6. ONSS Showing Wide and Narrow Cameras and MOOG/Schaefer Gimbal; *Gimballing the camera increases the power and flexibility of the ONSS greatly, and relieves the Altair vehicle of the need to change attitude to obtain OpNav data*

Altair has responsibility for safe return of the crew immediately after the Trans-Lunar-Injection maneuver. In this case, Altair is in the immediate vicinity of the Earth, and would likely use the NAC component of the ONSS to measure the positions of orbiting artificial satellites, such as GPS

or telecom satellites in MEO or GEO respectively. Ephemeris positions of these vehicles would be carried onboard, and kept current during the mission as necessary. It is possible to predict the location of these vehicles to the precision of 10's of meters, and as such, in the early Earth-departure phase of the Trans-Lunar cruise, it is possible to perform position measurements to the limiting navigation performance of the NAC, or about 2.5μradians. This implies kinematic state accuracy determination of about 0.5km midway to the Moon. Of course, instantaneous determination of position does not turn into overall navigation performance due to non-gravitational accelerations (Ref. 2).

Table 1: The GDoP of OpNav During the Phases of the Altair Mission

Instantaneous Geometric Determination of Position from Optical Data During Altair Mission Phases (all figures 3σ)					
	LEO	**LEO**	**Cis-Lunar**	**Lunar Orbit**	**Lunar Orbit**
	(10K km range to target)	(30K km range to target	200K km range to Earth/Moon	5m frame-tie errors	100m frame-tie errors
NAC	25m	75m	0.5km	5m	100m
WAC	357km	1.07km	6.5km	5m	100m

Once in orbit, Altair will use both the NAC and the WAC for navigation purposes. Table 1 shows the relative kinematic downtrack/crosstrack and radial power of "Geometric Determination of Position" (GDoP) of the optical data from the two cameras (NAC and WAC). There are two different conditions expressed in the table, where the map frame tie errors are as a result of LRO wide or narrow imaging, where those errors are 100 and 5 m respectively. In general it is expected that the better NAC survey will be available within a few kilometers of a crewed landing site, and probably uncrewed sites as well.

OpNav processing will take place in orbit regardless of the state of the communications link, as the optical data substantially improves the performance of the navigation filter (Ref. 2). Use of the OpNav will continue after PDI, and through the orbital descent. As will be discussed later, OpNav data is amenable to processing with simple filter models due to the fact that it is unperturbed by small accelerations as radio metric measurements are, and can rapidly recover from – if not calibrate – accelerometers biases and drifts. Figure 3 presents the performance of a landing simulation using only OpNav data and the Deep-Impact AutoNav filter during a landing on the Shackleton rim, with a performance peak at better than 10m accuracy. A more thorough discussion of this simulation will be given in Section V, but the salient feature for this brief performance discussion is that position errors are reduced to 10m or so 1 minute before landing.

On ascent, performance of the OpNav data should be equivalent to the descent with respect to landmark tracking accuracy. In fact, the method used for image processing of the rendezvous target images (Orion) is landmark tracking, where the "land"marks are track-able regions on the surface of the target. Figure 4 shows an example of a modeled vehicle being used as a rendezvous target (Ref. 3), in this case, the Mars Reconnaissance Orbiter (MRO). With accurate modeling of the target, the accuracy of RPOD navigation OpNav data will be limited only by the resolution of the camera; of course, ultimate accuracy of RPOD control is limited by many other factors, including non-gravitational acceleration, minimum thrust impulses of the Altair (and target vehicle) relative attitude and rate control ability, and many other factors.

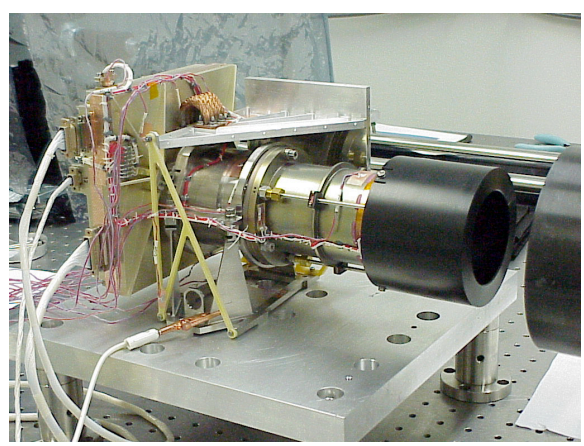

Figure 7. ONC on the Test Bench with Radiators and Insulation Removed

B. OpNav instrument and component description

The Altair project has included in the GN&C equipment list the instrument shown in Figs. 5, and 6. This instrument is a composite of two cameras mounted on a gimbal. The narrow angle camera is modeled after the MRO Optical Navigation camera (ONC). The wide-angle camera is modeled after the Draper Laboratory Inertial Stellar Compass (ISC). The characteristics of these two instruments are shown in Table 2. The gimbal pictured in the instrument suite is based on the MOOG Biax-11.

Table 2: ONSS Instrument Specifications

ONSS Optical Sensor Attributes		
Parameter	**TacSat-2 ISC**	**MRO ONC**
Array Size	512x512	1024x1024
Focal Length	3.5cm	50cm
F-number	1.2	8.3
Aperture	3cm	6cm
Field of View	21deg	1.4deg
Power	0.4W	4W
Mass	1.4kg (w/o CPU)	2.8kg
Sensor Type	APS	CCD

1. The MRO ONC

Of the hardware resources to draw upon for developing an integrated AutoNav instrument, the recently flown and flight-tested MRO Optical Navigation Camera (ONC) is the principal element. Designed over a period of six years with the DS1 AutoNav experience fresh in mind, and the then developing Mars Sample Return mission being vigorously investigated, this instrument was specifically designed to image the dim objects (e.g. stars, distant asteroids, and orbiting sample canisters) in the proximity of bright near-field planetary bodies. As such, it has a large dynamic range and low noise. Table 3 compares the attributes of this instrument with the Cassini imaging system from a navigation standpoint. The MRO ONC is a small, compact, high-power instrument and offers such a wide range of navigation capability, it will likely form the core of the navigation system for a broad spectrum of upcoming missions, and therefore is the ideal candidate to form the center of an integrated AutoNav instrument package. Figure 7 shows the camera being bench tested, and Fig. 8 shows the instrument mounted on the MRO spacecraft, and well demonstrates the diminutive size of this powerful camera.

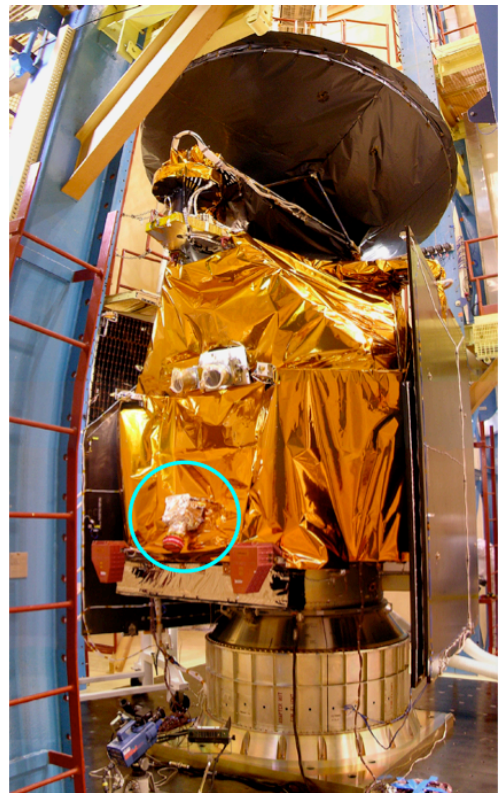

Figure 8. ONC (Circled) Ready for Launch Aboard MRO; *This camera serves as the Narrow-Angle-Camera in Altair's Optical Navigation Sensor System (ONSS)*

2. The CSDL ISC

The Charles Stark Draper Laboratory has developed an instrument, called the "Inertial Stellar Compass", or ISC, Fig. 9. The ISC "is a real-time, miniature low-power stellar inertial attitude determination system composed of a wide field-of-view active pixel star camera and a microelectromechanical system (MEMS) gyro assembly" (Ref. 4). The ISC solves the general attitude estimation, and "lost-in-space," or tumbling spacecraft attitude problem in a compact integrated hardware and software package, of low mass and power (~2.5kg and ~3.5W respectively - ibid). One of the key innovations of the ISC is to combine the ability of a camera to not only determine a fixed attitude, but rates as well, with the MEMS gyro to enable maintenance of attitude knowledge during fast turns when star tracking would not be possible. Another key innovation was the pre-integration of the hardware and software elements, including hardware drivers, into a single unit with simple high-level spacecraft interfaces. The ISC utilizes an ERC-32 CPU, which is a SPARC-based architecture, and hosts on this processor the RTEMS Real-Time Operating System (RTOS). The ISC represents an excellent model for an integrated software and package which potentially greatly eases a mission's development and integrations costs. Additionally, equipped as it is with the CPU and wide angle camera, the ISC represents a

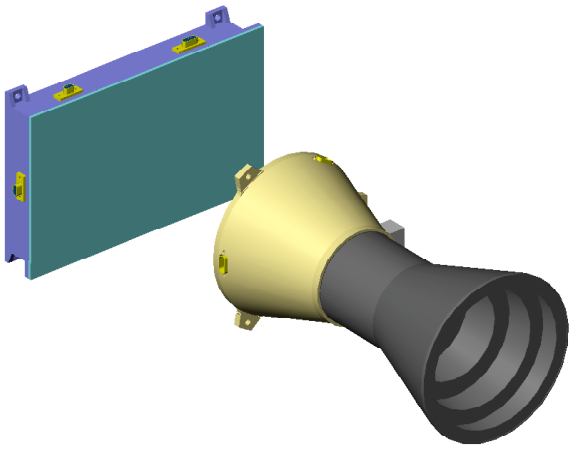

Figure 9. The Charles Stark Draper Laboratory Inertial Stellar Compass (from Reference 17); *The ISC was flown and successfully tested oduring the New Millennium ST-6 Mission aboard the TacSat-2 Spacecraft*

very logical choice for a core element of the integrated AutoNav instrument suite, and will serve in that role as a point design for further discussion and analysis. The ISC has been flight tested as part of the New Millennium space technology program ST6 mission.

3. Integrating and Gimbaling the Instrument Suite

As mentioned above, one of the mission resources that is in conflict when AutoNav is operating is time, in particular, pointing time. Turning the vehicle to orient instruments in order to obtain navigation data, whether radio or optical, almost always presents a conflict with other mission objectives when those instruments are body-fixed. The means of greatly reducing this conflict is to provide a self-pointable camera by mounting the instrument cluster on a gimbal. There is a potentially wide range of capabilities to choose from for a camera gimbal, but the MOOG/Schaefer Model 11 2-dimensional gimbal offers a good strawman design for the ONSS (Ref. 5). This concept will provide Altair with a 9kg 16 watt package of 0.04 cubic meters that provides full navigation and attitude estimation capability for the cost of a typical small science instrument. Table 3 shows vital characteristics of the ONSS. One of the principal costs that a space exploration mission undertakes is the integration of hardware and software elements to each other, and those elements to the spacecraft overall. Guidance Navigation and Control (GN&C) software and hardware elements are often the most difficult, costly and risk-prone to integrate. A pre-integrated suite of imaging instruments, gimbal and AutoNav and attitude estimation software could potentially save substantial funds and retire much mission risk for a wide range of future missions, hence the design configuration of the ONSS.

Table 3: Instrument Power: ONC vs. Cassini NAC

MRO ONC vs. Cassini NAC Navigation Capabilities			
		MRO ONC	**Cassini NAC**
Pixel Resolution		26 μradian	10 μradian
Field of View		1.4°	0.6°
Mass		2.8kg	20kg
Power		3 Watts	20 Watts
Volume		15x15x32cm	25x25x120cm
Noise Level		11 e^-	30-50e-
Full Well		160,000e^-	100,000e^-
Dynamic Range		1000 to 1	250 to 1
Nav Accuracy		2 μradian	2 μradian

C. Optical navigation system architecture considerations

1. The Optical Data type

Figure 10 shows a schematic of the optical navigation data type. The measurement type is basically a two-dimensional angular measurement of a near-field object relative to a known reference frame. The inertial reference frame is best provided by positions of stars measured in the same field of view as the target observations. These star positions provide a direct measure of the pointing direction of the camera. Without this directional fix, the best available information must be invoked, which is usually derived from the resident star tracking aboard the host spacecraft; Altair will be equipped with one such tracker, with the ONSS providing a backup star tracker function. An individual optical navigation data type is in essence a kinematic determination of position

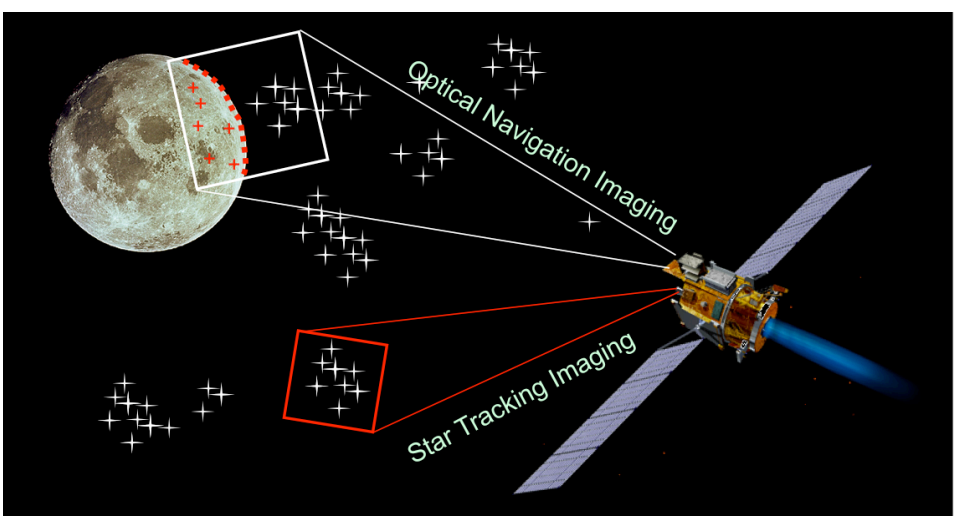

Figure 10. The Optical Data Type: *OpNav is the location of a near-field object (e.g. the Moon) relative to a well-known far-field object (e.g. the background starfield) or relative to well known camera attitude (e.g. from a startracker) - however, with a sufficiently wide-field imager, pointing knowledge can be obtained simultaneous to position knowledge from near-field objects such as surface landmarks.*

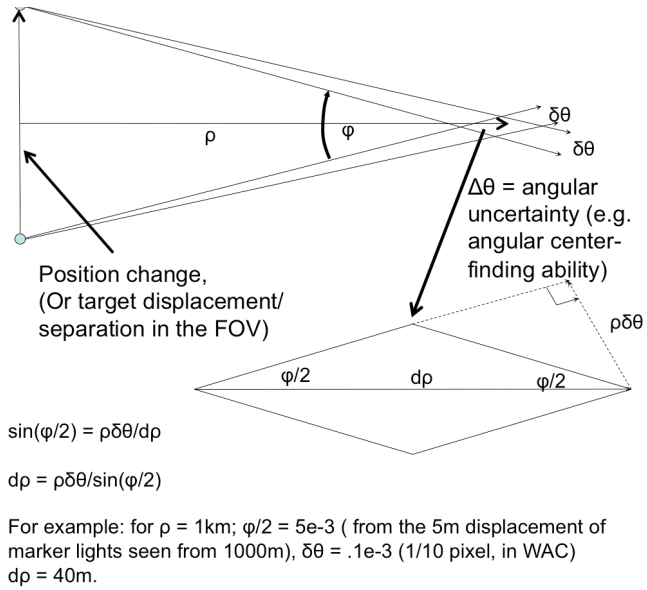

Figure 11. Determining Range from OpNav; *In this example at 1km range, and using the ONSS WAC, 40m range accuracy is obtained.*

in two dimensions, parallel to the camera field of view. But with multiple images in the field of view, such as multiple landmarks, and if the field of view is sufficiently wide, important range-to-target information is obtainable (Fig. 11). Additionally, with multiple objects in the field of view, it is possible to separate pointing from position errors, thereby estimating the camera pointing independently without recourse to the spacecraft base-body estimated attitude. Estimation of the camera pointing – and therefore determination of the position of the inertial reference – from the opnav picture itself is important, because in general this can be done to greater precision than via a star tracker, largely due to the accumulation of alignment errors from the tracker to the spacecraft body, the spacecraft to the gimbal, and the gimbal to the camera.

2. Optical Data Navigation Filtering

Though optical navigation images contain no velocity information in and of themselves, use of a dynamic model and multiple measurements obviously supplies such information. But this inherent insensitivity of the data type to velocity is an advantage in the data filtering, as, unlike Doppler radiometric data, opnav data cannot directly perceive the results of the many endogenous accelerations that most spacecraft are subjected to, especially manned spacecraft. These include uncoupled attitude control events, solar radiation pressure, and gas and liquid venting in the case of manned vehicles. This insensitivity makes optical data particularly attractive for automated onboard processing (Ref. 6). The simplest filtering strategy, and one that was used for lunar landing simulations to be discussed in later sections, is a batch-sequential filter, with a sliding window data range, and occasional epoch-state updates (Ref. 6). Since optical data is relatively sparse (maximum reasonable data frequency is probably5 to 10 seconds) keeping a recent data history available for comparison and editing/rejection of data "blunder points" is of substantial advantage. The opnav batch sequential filter noted above does have a stochastic component which accounts for random (uncolored) noise in the camera pointing, since such errors are generally adequately modeled as uncorrelated from picture to picture. The simplest solution from a filtering architecture standpoint then is to solve for the pointing as part of

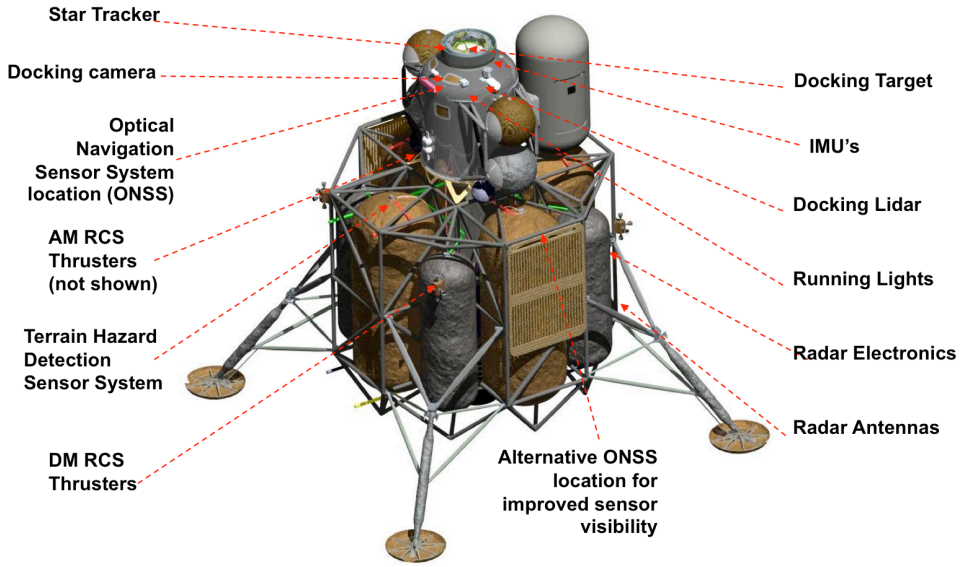

Figure 12. Location of GN&C Equipment on Altair, Noting Two Possible Locations for the ONSS; *Positioning the ONSS on an outrigger clear of the Descent Module will greatly enhance the visibility and flexibility of the instrument*

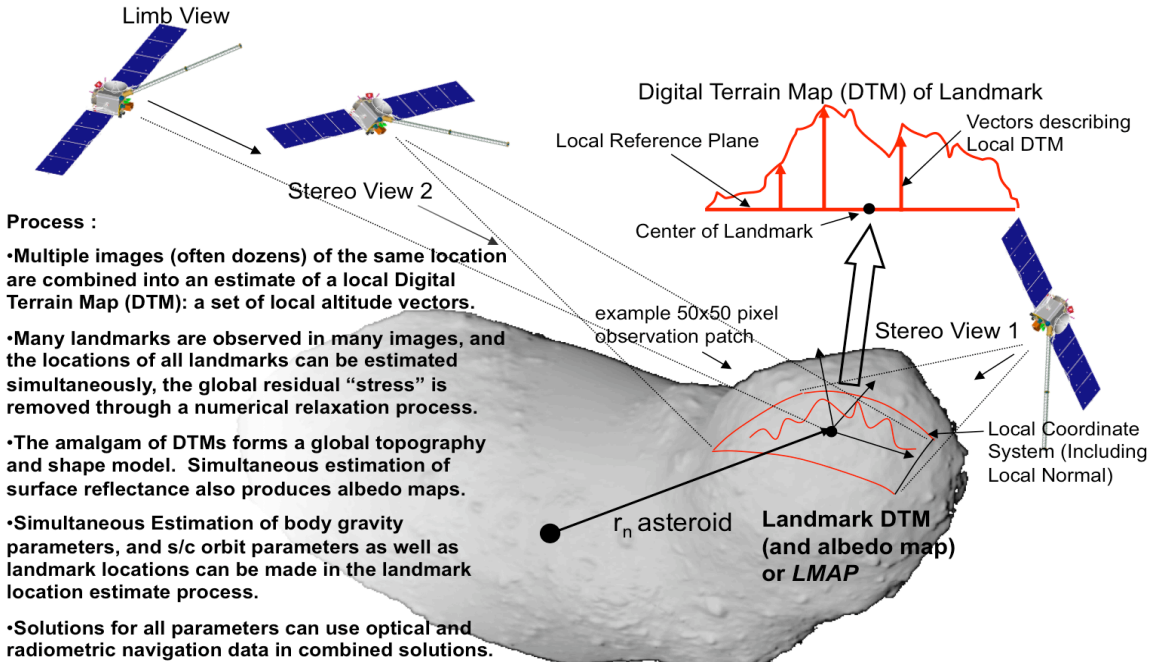

Figure 13. Generation of Stereophotoclinometric (SPC) Terrain Maps (Topography and Albedo); *This map-making process will occur prior to the critical GN&C operations such as descent and landing or rendezvous, and as is aready proceeding as a ground-analysis process for the Moon using LRO images*

the picture processing, and use these values in the dynamic state estimation batch filter. Such a strategy eliminates the need for a stochastic process noise filter in the state estimator itself. Ignoring in the filter the "noisy" accelerations aboard the spacecraft is only possible with an OpNav estimator because of kinematic determination of 2, and usually 3-dimensional target-relative position, and thus velocity as well through the relative position determinations between as few as two pictures. A navigation filter using Doppler radiometric data must infer 3-dimensional position and velocity information from 1-dimension of velocity; similarly, range radiometric data directly measures position in one dimension; use of either –or preferably both- requires accumulation of and carefully modeling of the vehicle dynamics over relatively long data arcs (at a minimum, many minutes). In contrast, two opnav frames over as little as 10 seconds spans the position and velocity solutions space, and as few as 10 pictures gives excellent navigation results, as will be discussed below. So, a simplified filtering strategy that fully allows for the vehicle's complex accelerations would invoke a short sliding-window data arc, a decaying weight filter within that window, and periodic state updates. This strategy satisfactorily allows the epoch state to absorb the effects of the accelerations.

3. Gimbaled Camera Considerations

As noted, the Altair ONSS includes a gimbal. Though the gimbal potentially adds error sources to the navigation data, as noted above, these are negated when an independent inertial reference can be inferred within the opnav frame. The principal reason for including a gimbal in the ONSS is to eliminate the need to turn the spacecraft to point the camera. The gimbal being considered, the MOOG Biax 11 or the slightly larger Biax 22, has excellent performance, including a control resolution of a miliradian, and encoders capable of 0.1mrad resolution. Though not based on pictures, the Apollo lunar-orbiting navigation process included the occasional use of optical data, which were manually observed crater locations measured against ground reticles in the view-ports. Though very powerful, such measurements were operationally difficult, requiring turning of the spacecraft, which itself induced a trajectory dispersion (Ref. 7). The gimbaled ONSS largely eliminates this need, by having up to a 2π steradian field of regard. Figure 12 shows two possible mounting locations of the ONSS, on the pressure vessel/cockpit and on an RCS strut. Investigations are being made to mount the ONSS on an outrigger, probably one of the RCS thruster outriggers, in order to provide better visibility around the Descent Module (DM). From the Ascent Module (AM), visibility of the landing site from the cockpit, even on the extremely shallow nominal 16° terminal glide-slope, is poor, with visibility of the landing site is lost at least a minute before landing. Hence, with an out-rigger mounted

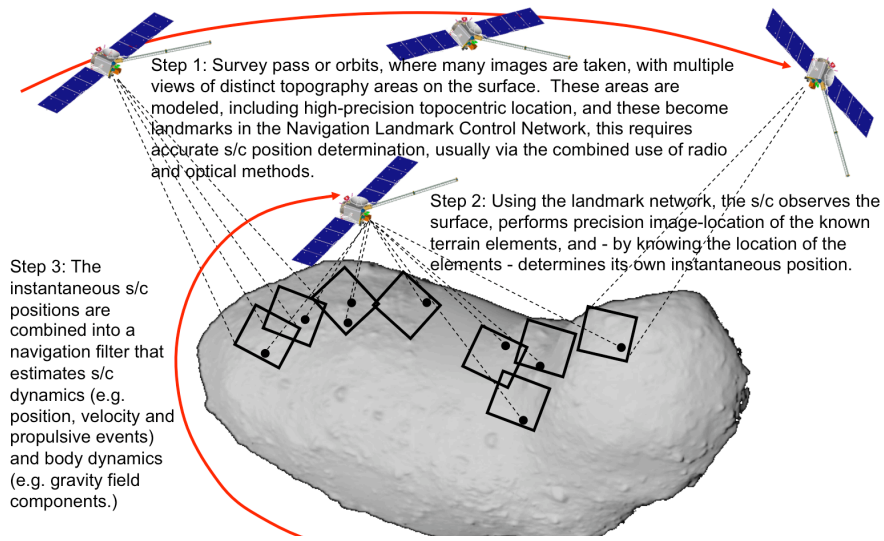

Figure 14. Use of SPC Onboard Terrain Maps; *Is a multi step process, beginning with the generation of the maps – most likely on the ground - followed by subsequent detection and tracking of specific landmarks, followed in turn by use of the landmark locations in the onboard navigation filter*

ONSS visibility around the DM is complete and continuous view of the landing site is possible from ten minutes or more before landing, allowing for a re-optimization of the landing trajectory, as will be discussed below. With such difficult human seeing, it is natural to consider allowing the crew excellent improved visibility of the landing site via electronic vision.

III. Landmark-Relative Position Determination Methodology

A. Terrain maps of the Moon and other target bodies

Optical navigation in orbit at the Moon and during landing will depend on surface landmarks networks to be created with sufficient density, resolution and accuracy to enable OpNav. At this writing, the Lunar Reconnaissance Orbiter (LRO) is surveying the entire lunar surface using its wide angle camera (Ref. 8). Typically, these images will have 75m resolution. Such resolution is sufficient for orbital operations, and cis-lunar transit, but not for precision landing. In an area a few km square around likely landing sites chosen by the Constellation Program, LRO has agreed to perform surveys using the narrow angle camera, which will yield maps of resolution as "high" as 0.5m., and likely no lower than 0.75m. The process of turning these images into maps will take one of two options: either conventional stereography, as has been used for many decades by the USGS for aerial survey and map generation, or stereophotoclinometry (SPC), which is a much newer method developed at the Jet Propulsion Laboratory for the generation of topographic maps from deep space images of small bodies (Ref. 9, 10). Conventional stereography will produce maps of resolution 2-5 times greater than the contributing image resolution due to the discretization of the images into distinct registerable landmarks , whereas SPC can produce maps at or slightly better than the image resolution owing to the fact that SPC will use all available image pixels for map development. NASA's Lunar Modeling and Mapping Project has responsibility for generating these maps, and will apply the appropriate methods needed (Ref 15.)

B. Development and use of landmarks

Figures 13 and 14 show in schematic the method of development of topographic and albedo maps using SPC, and subsequently using the maps for navigation. The figures show the operation for a small body such as an asteroid, but the method is just as valid for a large planetary body such as the Moon, although the magnitude of the mapping task is necessarily much larger. The requirements for performing SPC map development are higher than for conventional stereography, with a minimum of 3 different views of the surface required, with off-nadir angles of greater than 20°. This is as opposed to stereography which can develop maps with only two views, differing angle by as little as 10°. The SPC method, as shown in the figure, models the light absorption and re-emission as a function of reflectance law, albedo and surface normal. Necessarily, the altitude of a particular element of the map, as well as the position of the camera and the pointing angle of the camera are required elements of the model.

The surface of the target body is first divided into sub-maps, referenced to the six faces of a cube. Multiple images (at least three) of each sub-map of terrain are combined in an estimation process. Within each map, elements of altitude and albedo are estimated. For each picture, spacecraft position, and pointing angle are estimated, across all submaps contained within each picture. The creation of the altitude and albedo elements is not an estimation

process per se, but an annealing process, where individual altitude values are represented as a potential-energy work-function, and differences in brightness modeled and observed are reduced by minimizing the potential energy across the field.

Again, Fig. 14 shows the use digital elevation and albedo map landmarks during a mission. The maps will be carried onboard digitally. As pictures are taken by the ONSS, locations of the catalogued landmarks are determined, their maps are retrieved and digitally illuminated under geometric conditions identical to those just observed in the frame. Once illuminated, the synthetic view of the landmarks are correlated with the actual images, and the landmark centers thereby determined. This information is passed to the navigation filter as discussed above.

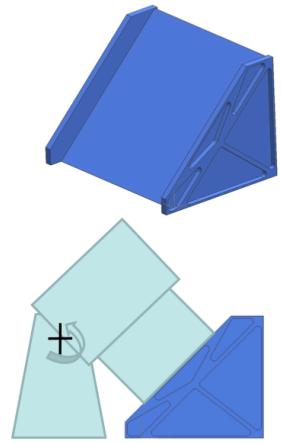

Figure 15. Passive Dust Shield for the ONSS; *This cover makes use of the fact that dust in an airless environment travels on ballistic trajectories*

IV. Expected OpNav Performance for Altair

A. LEO Opnav Operation

As discussed earlier, the principal reason for adopting optical navigation for Altair as a backup to the mainline navigation technique is to provide safe return of the crew during disruption of the com link, specifically corruption of the radiometric measurement signal. To this end, Altair will have responsibility to perform optical navigation virtually from the start of the lunar mission, immediately after the Trans-Lunar-Injection maneuver. From this point until the beginning of the return from the Moon, Altair will have chief responsibility for navigation, and will use optical navigation in the contingent case of a disabled radiometric link. In the actual operations of course, optical navigation is likely to be an ongoing part of the navigation processing stream. Even before leaving Earth orbit, the system is liable to be on and taking images for navigation as part of a check-out process. While in Earth orbit, it is likely that Earth orbiting artificial satellites will be used as navigation targets. The advantage of such targets, such as MEO GPS satellites, or GEO com satellites is that they are large bright targets whose orbits are well known and stable. It is important to note that even if the GPS network is unavailable due to solar-storm atmospheric disruption, or distance from Earth, the large GPS satellites themselves remain excellent navigation optical targets. Table 4. gives a summary of kinematic geometric determination of position (GDoP) by optical navigation data for Altair, showing that for typical LEO to MEO ranges, the in-orbit GDoP power of single image would be on the order of 25m, with multiple images likely to reduce the position error to the level maintainable in the face of non-gravitational accelerations, as discussed earlier. For this type of observation the NAC would be used.

B. Cis-Lunar Cruise

In cis-lunar cruise, either back-pointing images of the Earth satellites, or forward looking images of the Moon

Table 4: Altair OpNav in LEO - Using Manmade Targets

ONSS Earth Orbit Capability and Accuracy with Artificial Satellites		
	Instantaneous pos. and vel.	1 hour ahead pos. and vel.
Navigation of Distant Target (10000km) (LEO)	25m; 2.5cm/s	45m; 2.5cm/s
Navigation of Distant Target (100000km) (GEO)	250m; 2.5cm/s	300m; 2.5cm/s

utilizing landmarks could be used, at the mid-point in cruise both having equal power. One reason both types of measurements are so precise is the presence of stars in the field providing an inertial reference frame. For the Moon images however, this will be somewhat more problematical than for the star-like images of artificial satellites, as the Moon is a resolved object and is very bright. Even with the 1000:1 dynamic range of the ONC, with the low exposures necessary of the lunar surface, stars no dimmer than 7th magnitude are liable to be captured with the sub-solar limb, limiting the navigation frame pointing options. This problem will be addressable to some extent by utilizing landmarks only near the terminator, thereby gaining a magnitude or even more in available navigation reference stars. It is also possible that for the later Altair missions, artificial satellites will be available for

navigation purposes at the Moon, eventually including the Lunar Communications and Navigation System orbiters. As with GPS and com satellites at the Earth, these will present ideal navigation targets, although there certainly will not be as many. Science orbiters of various nations will also likely be candidate targets.

C. Lunar Orbital Operation

Once in orbit, at the likely 100km nominal altitude, landmarks will likely be the only available nav targets for the majority of the time. Both the wide and narrow angle cameras could be used in orbit for landmark tracking, but the wide will be much more flexible, and provide substantial altitude information in each picture. Although the NAC is capable of obtaining data from this lunar orbit of very high precision, 2.5m, the maps that will be available will in general not be of this caliber. On descent and landing, the WAC will again be principally be utilized, but depending upon the descent trajectory chosen, the NAC would be able to obtain high-precision images of the landing site many tens of kilometers, and several minutes before landing, which would be of substantial assistance to the crew and to the automated GN&C system aboard, as will be discussed below.

D. Descent and Ascent

During the terminal descent, opnav images will be taken not only to adjust the course relative to potential hazards that have been identified at the landing site, but to accurately determine the location of the descent module in its final touchdown position as well. The ONSS will be stowed in its safe position during the landing sequence at approximately 30m altitude due to potential dust contamination. Figure 15 shows the dust shield; this is a purely passive device, as dust raised by the DM main engine travels on ballistic trajectories – in fact much of it goes into orbit. The dust shield needs only to protect the optics from direct and single-bounce ballistic trajectories.

The ascent phase is nominally not as critically dependent upon OpNav as descent and landing, as the nominal ascent is largely target relative, where the target is Orion, and that vehicle will be in view for much of the ascent. Nevertheless, performing landmark-relative tracking and navigation will improve the performance of the ascent navigation, and reduce the statistical propellant consumption for the early rendezvous targeting maneuvers. Additionally, OpNav will be necessary during contingent cases, where descent is aborted, or where an immediate launch from the surface is necessary, before Orion is near to being in view. In such a case, Altair will target a future location of Orion based only on an a priori ephemeris carried onboard, and controlling its inertial state via OpNav until Orion comes into view, at which time the Orion ephemeris can be updated.

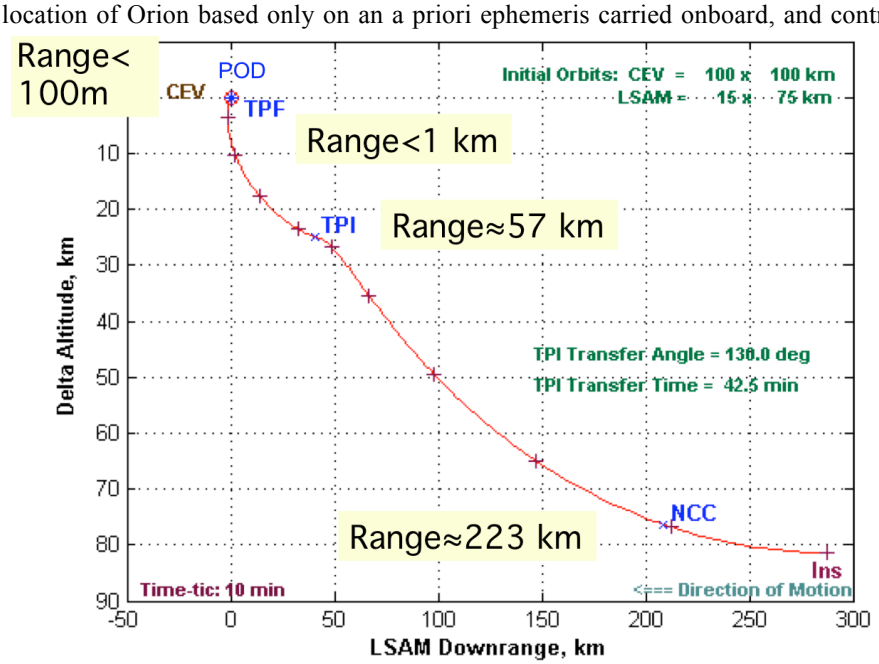

Figure 16. Nominal Geometry of the Approach and Rendezvous of Altair with Orion; *Applicability of OpNav to the RPOD problem is wide, and will be used throughout the ascent and rendezvous.*

E. Rendezvous Proximity Operations and Docking

Figure 16 shows the phases of the ascent and approach to Orion. Table 5 presents the OpNav performance as a function of range during these phases. The table shows both transverse directions and range-to-target determination. Range determination takes two forms in this analysis, single-frame "geometric" range determination as shown in Fig. 11, and "dynamic" range determination where a maneuver of known magnitude, as measured by the IMUs, occurs between sets of OpNavs providing a wider baseline for parallax determination than can be accomplished in a single OpNav frame. In this case, the deterministic components of the approach burns suffice for this purpose, which in general have sufficient cross-line-of-sight velocity-change components for optical range determination. Optical navigation is not the primary data type used for rendezvous, proximity operations and docking (RPOD), but is backup to LIDAR and radiometric measurements. The applicability of the various measurements as a function of range is shown in Fig. 17. Figure 4 above shows the type of onboard modeling that is necessary to simulate and test RPOD optical navigation in close proximity. This model is of MRO. It should be noted however that full-spacecraft photo-quality modeling is not necessary onboard; only the landmarks need be modeled. In this case, several small targets have been placed on the simulated MRO (not on the real vehicle) and two large patches (a NASA "meatball" and a JPL logo) have been placed on the simulated spacecraft, and are rendered in the landmark tracking system for navigation.

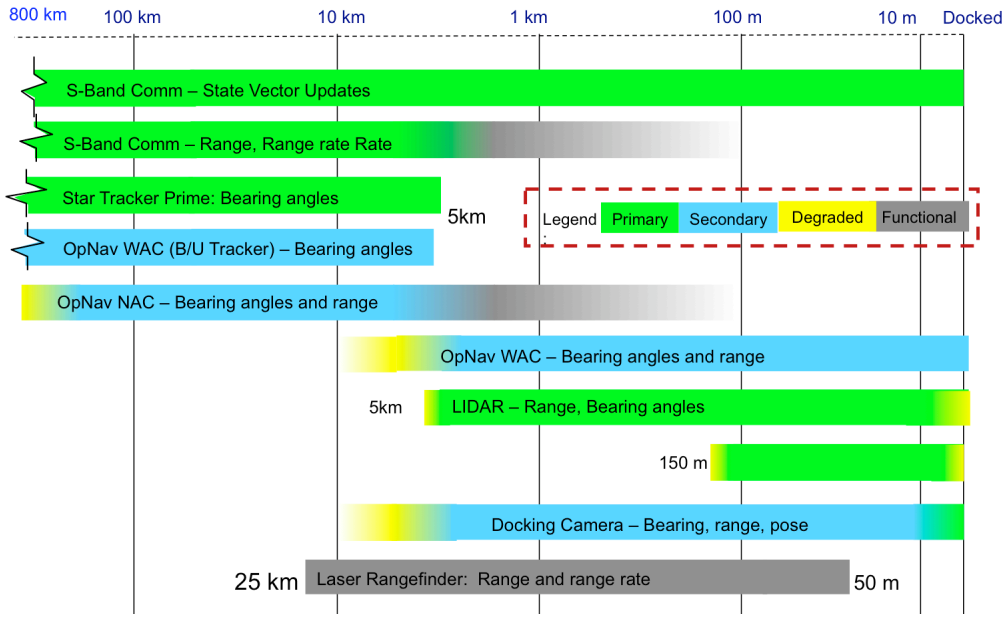

Figure 17. Applicability of Instruments to RPOD, Including Cameras for OpNav; *Passive optical data can be operational and complete the entire approach and docking, and will be the principal means of doing so for the Mars Sample Return Mission*

Table 5: OpNav Capability During RPOD

Phase	Time before docking	Range	Rel Velocity or delta-v	Assumed POS dv	Camera(s)	Range Method	range knowledge*	in-plane knowledge*	notes
NCC	2:46	223km	20m/s rel	12m/s	NAC	dynamic	2km||6m/s	1m||4.5cm/s	5 minute data arc
TPI	2:00	57km	20m/s dv	12m/s	NAC	dynamic	400m||1.4m/s	30cm||4.5cm/s	5 minute data arc
TPF	1:13	1km	n/a	n/a	NAC,WAC	geometric	(WAC) 40m||66cm/s	20cm||9mm/s	1 minute data arc
							(NAC) 1m||1.8cm/s	5mm||9mm/s	1 minute data arc
POD	~0:07	100m	n/a	n/a	WAC	geometric	4m||7.5cm/s	2cm||4.5mm/s	30s data arc

V. A Prototype Optical Navigation Flight Software System and its Demonstration in a Small Body Rendezvous and a Shackleton Landing Scenario

A. Flight Software Architecture for Optical Navigation

1. Virtual Machine Language and Phoenix Landing GN&C Architecture

VML is an advanced procedural sequencing language that simplifies spacecraft operations, minimizes uplink product size, and allows autonomous operations aboard a mission without the development of specific autonomous flight software capability. VML is the command and control software system on a variety of NASA missions, including Mars Reconnaissance Orbiter, Mars Odyssey, the Spitzer Space Telescope, Dawn, and Phoenix. The language is a mission-independent, high level, human-readable script interpreted by flight software. VML sequences are implemented as a set of named functions, allowing operations to be abstracted. Files containing named functions are loaded onto specific sequencing "engines," which serve both to contain the script and to provide a thread of execution. VML features a rich set of data types (including integers, doubles, and strings), named functions, parameterized functions, IF, FOR, and WHILE control structures, data-type polymorphism, and on-the-fly creation of spacecraft commands from calculated values. Statements may be time-tagged with a mix of absolute and relative times. Programmable time tags to delay by a calculated amount or delay until a calculated time are provided. The language also features event-driven sequencing through the use of sequence global variables, which contain spacecraft state information visible to all functions and accessible to flight software through accessor routines. Global variables computed and stored by expert systems accessed by VML such as AutoGNC may be stored and reasserted. These programming and conditional sequencing constructs form a critical resource from which the AutoGNC Exec draws to perform the complex, dynamically varying tasks required of autonomous onboard guidance, navigation and control functions.

The VML system consists of several components working in concert to provide a programmable spacecraft capability (Fig. 18). The user creates functions as text using a text editor. The text is translated by the VML compiler according to mission-specific definition files which are maintained using a command database. The VML compiler also uses mission-specific command translation and time translation tools. The resulting binary is available for use by execution tools like Offline VM (OLVM) and VML flight component (VMLFC) running on the spacecraft or in the software test lab.

The starting point for the AutoGNC Executive function was Phoenix EDL, which also utilized VML. The power of the VML approach to the EDL design was made manifest in the utilization of state-machine diagrams for the 24-phase EDL progression, and the fact that the EDL sequence blocks that comprised the state machines were run from EEPROM loaded before launch. No necessary changes were discovered in flight.

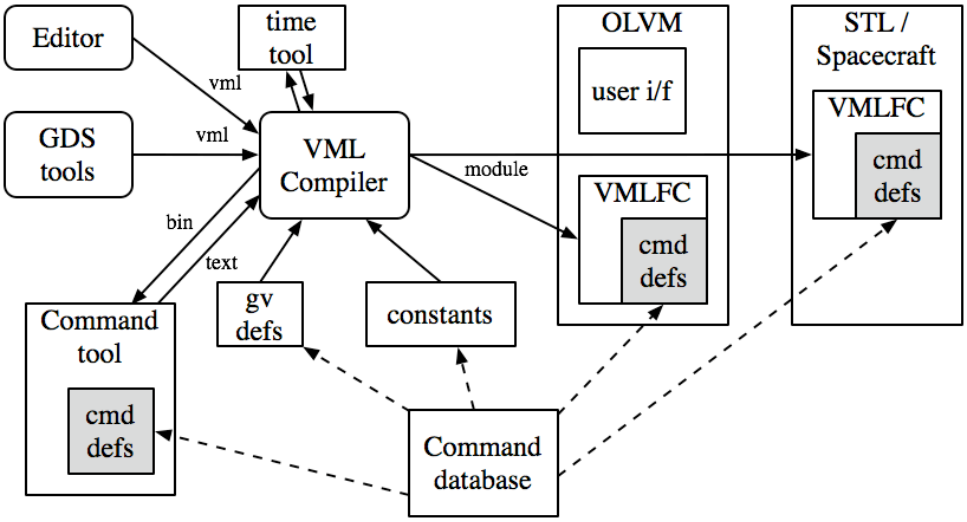

Figure 18: Relationship of VML Elements; *VMLFC is the Flight Component of VML, whereas OLVM is "Off Line" VML used for ground sequence and uplink product development; the VML comiler provides input products for both*

The Phoenix EDL state progression was linear: land or die. There were no branches or aborts possible, the 24-steps of the Phoenix EDL state machine progressed monotonically. The power and utility of the VML state-machine approach to EDL design was made evident in its use in the design and test of the sequences and the fact that the EDL was accomplished with the EEPROM version of the VML blocks, burned in before launch. No needed updates were identified during cruise.

The Phoenix EDL "mainline" worked like a combination of flight director and GN&C process manager (functions that are differentiated in the AutoGNC design). Some characteristics of this "State-machine-like" sequence include:

- ***Detection of spacecraft restart and automatic catch up***
- ***Reactions triggered by events detected by or driven by flight software***
- ***A mixture of timed and repeated elements***
- ***Timeout backups for some events***

Phoenix did feature some manager structures in parallel

Figure 19. Phoenix EDL. *The methodology of the GN&C command and sequence subsystem has become the model for the AutoGNC Executive, which in turn is now a model for the Altair GN&C system.*

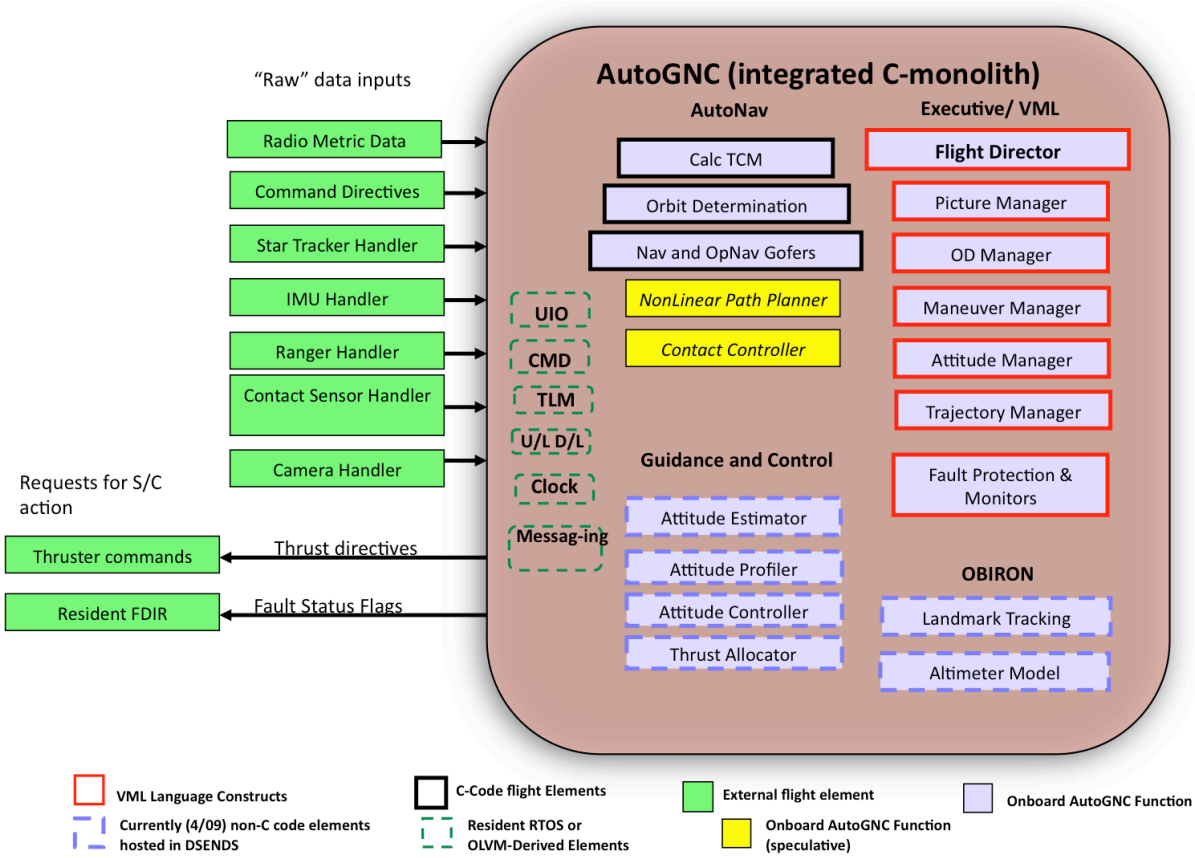

Figure 20: AutoGNC Flight and Simulation Architecture; *This GN&C flight system architecture has been applied to small body proximity operations, landing, and to an Altiar-like lunar landing scenario*

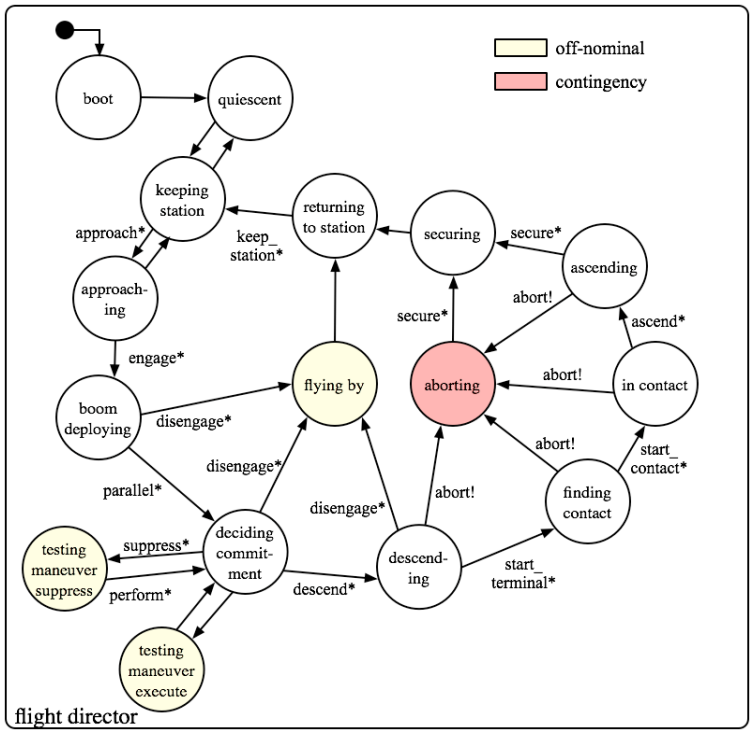

Figure 21: AutoGNC VML Executive Flight Director Mode Transitions; *The Flight Director is the highest level of control in the Executive, and is the state machine which controls the functions of all other Executive managers*

to the EDL mainline, these included:

- ***Communications manager***
- ***CPU utilization monitor***
- ***Landing imaging manager***
- ***Uplink loss manager***
- ***Reserve manager for last-minute changes (Sideline)***

As alluded to above, Phoenix EDL utilized VML 2.0. This was possible using conventional VML blocks and conditional execution due to the monotonic nature of the sequence. This will not be possible with the comet and asteroid sample return missions being simulated, as these Touch and Go (TAG) sequences offer many non-linear opportunities, including multiple sampling events, practice events, flyby runs, and "fly around" aborts. AutoGNC will be continuously evaluating its own health and making decisions on whether continuation is safe. Indeed, the AutoGNC system will autonomously make the decision to commit to the contact TAG event (or to simply fly safely by). These complex autonomy features required capabilities well beyond those of VML 2.0. VML 2.1 was created specifically to accommodate autonomous capabilities such as these in AutoNav and AutoGN&C.

2. Extensions of VML 2.0 into 2.1

VML 2.1 features a reduced memory footprint in order to fit more capability into modestly sized flight processors, and endian-neutral data access for compatibility with Intel little-endian processors. Sequence packaging has been improved with object-oriented programming constructs and the use of implicit (rather than explicit) time tags on statements. Sequence event detection has been significantly enhanced with multi-variable waiting, which allows a sequence to detect and react to conditions defined by complex expressions with multiple global variables. This multi-variable waiting serves as the basis for implementing parallel rule checking, which in turn makes possible executable state machines.

The new state machine feature in VML 2.1 allows the creation of sophisticated autonomous reactive systems without the need to develop expensive flight software. Users specify named states and transitions, along with the truth conditions required before taking transitions. Transitions with the same signal name allow separate state machines to coordinate actions: the

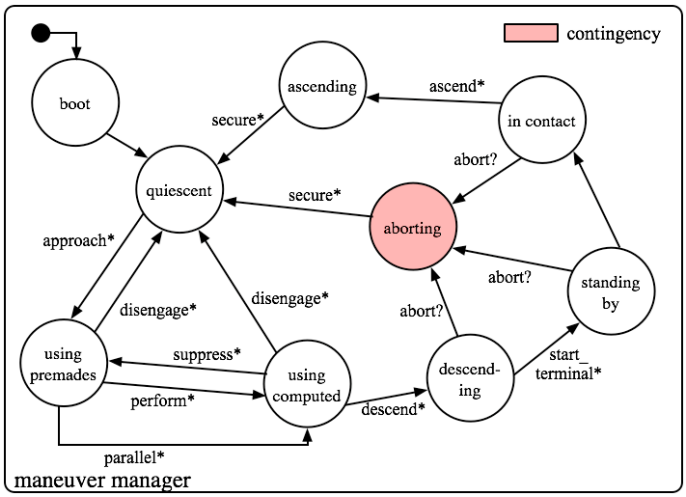

Figure 22: Maneuver Manager; *This manager effects the preparation of trajectory corrections by calling upon AutoNav services to compute appropriate maneuver parameters*

conditions distributed across all state machines necessary to arm a particular signal are evaluated, and once found true, that signal is raised. The selected signal then causes all identically-named transitions in all present state machines to be taken simultaneously. Reactive systems for comet and asteroid sampling missions have been created using the VML 2.1 state machine capability, and tested in a work station environment, connected to workstation-based simulations. These systems have been demonstrated to respond to a wide variety of nominal and off-nominal conditions, protecting the spacecraft while proceeding with the rendezvous, contact, and sampling of the body. Table 6 shows a comparison of these features. Principal among these capabilities are the State Machine constructs and the Select Loop, which is the main internal structural construct of those machines.

3. AutoGNC VML Executive

Figure 20 shows the integration of the AutoGNC Executive into the AutoGNC system. The VML AutoNav Executive is hierarchically structured within VML sequence engines. At the top of the hierarchy is the Flight Director (FD) (Fig. 21). The FD occupies its own VML engine. The structural basis for the FD is the use of State Machines, which are a feature that has been added to VML (version 2.1) specifically for the use of AutoGNC in complex operational scenarios such as comet sample return. The FD is a single state machine, with the states as shown in the diagram, e.g. "boot," "quiescent," "keeping station," "approaching," etc. Between each state is a transition operation. Transition operations are enabled and/or guarded by transition signals, which are labeled. For example, "approach*" is the transition between keeping-station and approaching. The suffix character, "*" and "!", respectively, designate whether the transition is guarded by states and conditions within the managers (to be discussed below) or can be transmitted and forced upon the system at any time by the FD. The FD diagram represents the principal motivation for utilizing VML to control operations of AutoGNC. This high-level representation allows rapid visualization of the operational state of the spacecraft during complex and critical operations, and offers assurances to project planners and operators that transitions between states are available and enforceable at appropriate and needed times, and that by sequestering behavior of the spacecraft within a state to that state only, the correct spacecraft behavior can be warranted virtually by inspection.

Figure 6: Features of the Updated VML

Feature	VML 2.0	VML 2.1
Blocks, libraries, sequences	yes	yes
Processor compatibility	PPC, Sparc	PPC, Sparc, Intel
Flight software attributes change	patch	command interface
Time tags	required	optional
Command completion detection	via WAIT	Implicit (like MER)
WAIT for simple condition	yes	yes
WAIT UNTIL compound condition		yes
SELECT LOOP rules		yes
Objects, attributes, methods		yes
State machines		yes
Collections		yes
Modeling		yes
Compiler creates globals file		yes

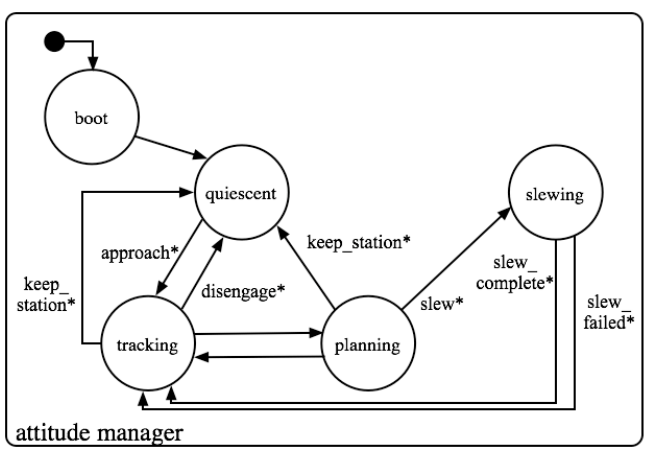

Figure 23: Attitude Manager; *This manager controls the actions of the Attitude Profiler, and effects either attitude hold or changes*

Of course, a state designation in and by itself accomplishes no GN&C functionality. The AutoGNC work, as well as other spacecraft work closely associated with GN&C functions, is abstracted within managers, which send commands to the C-language functions of AutoNav and Guidance and Control subsystems. The managers are themselves state machines and occupy their own VML engines. Following is a discussion of the several managers that comprise the AutoGNC executive, and a description of how they are arrayed to provide GN&C function.

4. AutoGNC VML Executive Managers

OD Manager: In the OD manager there are only two states, enabled and quiescent, and these states are transitioned by the FD depending on whether

AutoNav orbit determination results are needed. The current logic within the OD manager is such as to command an orbit estimate immediately after a picture is shuttered and sent to AutoNav for processing, or more than 10 seconds have passed since the previous OD was performed.

Maneuver Manager: The maneuver manager (Fig. 22) has more states, depending on a number of conditions, including whether AutoGNC is quiescent, computing maneuvers, using pre-computed maneuvers, executing the terminal ascent burn, or being asked to perform a contingency abort maneuver. The ascent burn (executed in the "ascending" state) is a special case maneuver in that it enjoys a non-deterministic execution time. The two states "using premades" and "using computed" refers to the use of premade maneuver designs vs. those from AutoGNC. The "abort" transitions are annotated in this chart with a "?" to indicate that this transition is always enforceable by the Flight Director. Transitions marked "*" are synchronized among all state machines – all must agree on the conditions of change. The "parallel" transition is a euphemism for "flying-by" (as in parallel course to the target.)

The Trajectory Manager: This manager is responsible for executing the velocity changes computed by the maneuver manager, and is toggled between executing and suppressing maneuver requests from AutoGNC, the latter being primarily a safety and/or testing condition. Also contained in the Trajectory Manager is the notion of slewing to an attitude appropriate for a maneuver. The states associated with the slew are coordinated with the attitude manager, and these are "slew", "slew complete", and "slew-failed.".

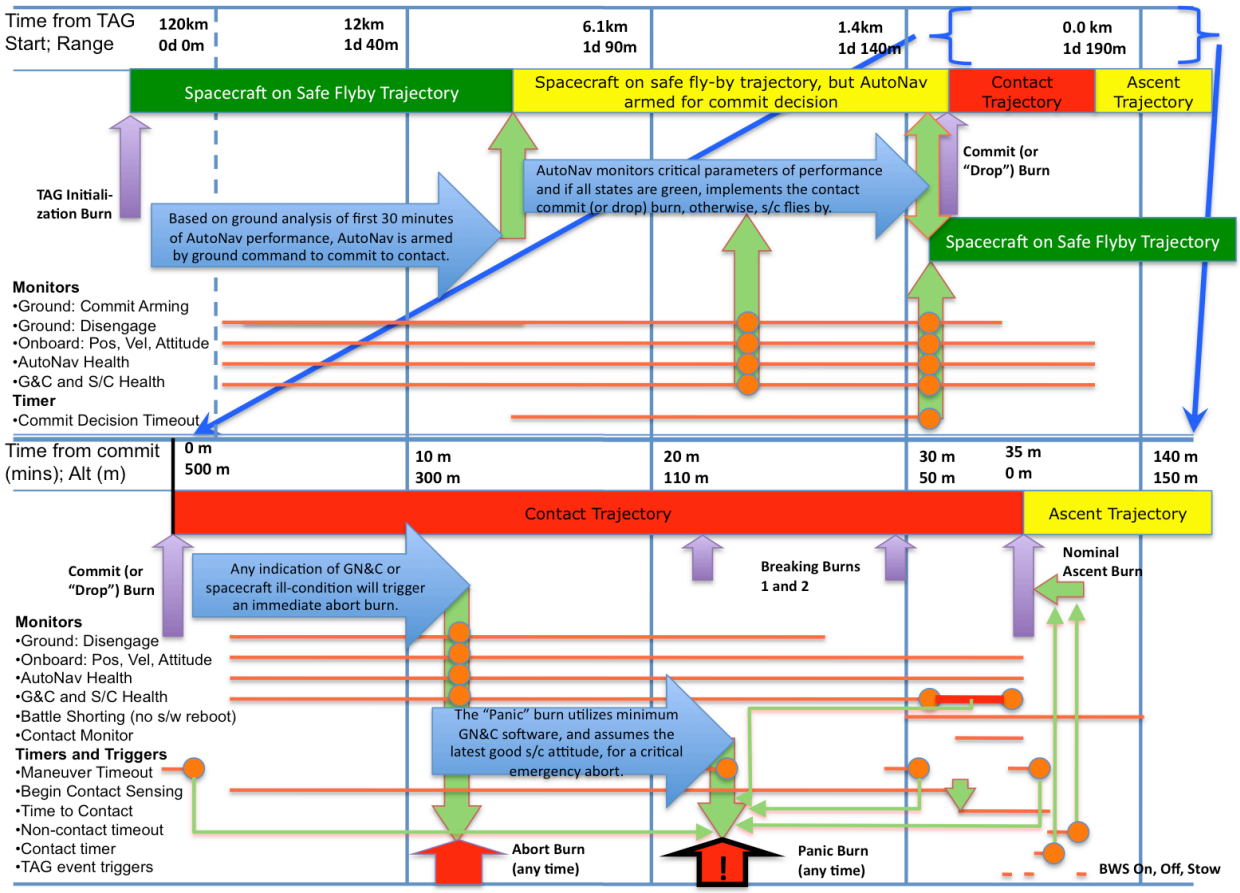

Figure 24. AutoGNC Monitors as Invoked for a Touch and Go GN&C Scenario, as a Model for Altair Landing; *Onboard autonomy as enabled by the AutoGNC VML Executive offers a rich space of options for fail-safing AutoGNC operations, integrating the reactive functions of managers with the sensing abilities of subsystem monitors*

Attitude Manager: The attitude manager is responsible for calling the computational elements responsible for planning and executing turns (Fig. 23). The planning function entails computing the needed time to execute the requested turn for purposes of letting various managers (e.g. the maneuver manager) know how soon to command a turn for a burn. The tracking state is that which is achieved after a target crequest change, which in general does not imply a constant inertial attitude.

Image Manager: The image manager controls the rate of images, and provides a framework for adjusting the frequency of image taking as a result of conditions. This manager has three conditions under which it will adjust that frequency, being far from the target, near to the target, or in an ascending condition. Other conditions that will be added will include no-imaging during a slew, or during a maneuver that has required a slew.

5. AutoGNC VML Executive Monitors

A discussion of Executive monitors raises the issue of fault protection within AutoGNC; VML-based monitors are the method of implementing fault tolerance in AutoGNC. Each of the individual monitors is responsible for assessing the health of a certain aspect of AutoGNC. Figure 24 shows diagrammatically the monitors and timing logic within AutoGNC as applied to a TAG ("Touch And Go") scenario, which has great similarity to many other GN&C scenarios including lunar landing.. In general, the timers are contained within the monitors and are part of the monitor logic. Figure 24 refers to the TAG scenario of a comet sample return mission. The spacecraft will depart from a holding station 24h before contact. The spacecraft will close on the comet nucleus on a safe flyby trajectory, while AutoGNC systems are engaged, and operate. The ground will monitor the health and success of the AutoGNC systems and give AutoGNC authority to commit to the "Drop" burn, which places the spacecraft on a potentially impacting trajectory. The decision to commit to the Drop Burn will be based on all AutoGNC systems being in a "green" status, and if there is a long-range (e.g. 5km) radar altimeter, the green status will depend upon agreement between that sensor and AutoGNC's estimate of altitude. If the monitors have not approved of the descent commitment by the time the "commit decision timeout timer" expires, then the spacecraft will abort the descent opportunity.

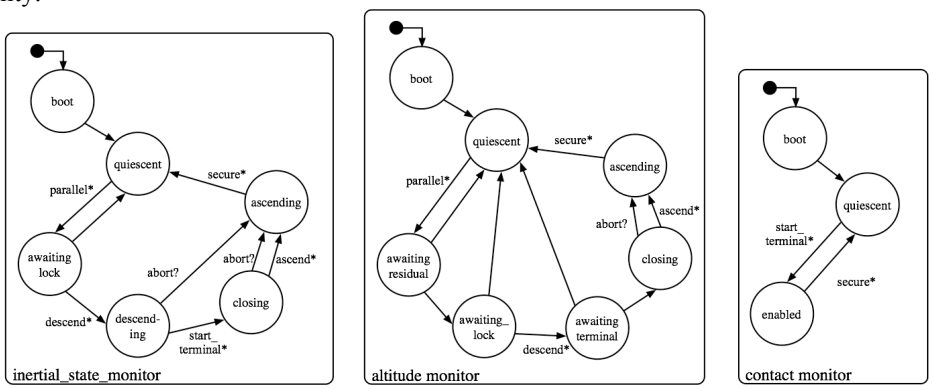

Figure 25: Inertial, velocity, and contact monitors. *The first two monitors provide #green, #yellow, and #red conditions regarding spacecraft inertial state. The contact monitor records the touchdown time.*

But once the Drop burn is executed, the general monitors will continue to assess AutoGNC health, and if any state or subsystem fails, an abort maneuver will be triggered. As the spacecraft descends to the surface, several terminal descent timers are in play, providing "life-boat" triggers for critical events in the TAG scenario. The first trigger is a countdown enable the final "time-to-go" calculation. Once enabled by this trigger, AutoGNC will compute the time to contact using the estimated spacecraft state, the known (via previous over-flight survey) position of the landing site, and the known local normal at the landing site. Armed with this information, the time-to-go is computed, and a second timer is set which will provide a back-up to the contact sensor in that after the nominal time-to-go duration plus margin, an ascent maneuver is triggered. This is to allow for the case that the contact sensor malfunctions and doesn't detect the surface, or highly unusual surface conditions cause the sensor not to respond – such as a very light and powdery regolith.

For most TAG scenarios, there will be several targeting burns during the descent. For the scenario pictured in Fig. 24, there are two. For each of these, as there are for Trajectory Correction Burns (TCMs) prior to the Drop Burn, there are timers (mediated by the Maneuver Manager) that determine if the maneuver has timed out before successful completion. The nominal duration used for this health determination is computed onboard. The final maneuver of the TAG scenario is the ascent burn. This burn is not calculated onboard but pre-computed on the ground, and stored onboard in a VML block. Two alternative ascent burns are available to be triggered in lieu of the nominal ascent burn, an "Abort" burn and a "Panic" burn. The Abort burn is performed very much like the nominal ascent burn, and utilizes all of the GN&C machinery (e.g. attitude profiler, controller, and thrust allocator, as well as a maneuver time-out monitor). The Panic burn on the other hand, is triggered by those events that are indicative of G&C malfunction, including time-out of a maneuver (including the nominal Ascent or Abort burns) and does not use higher G&C functions, but instead makes the tacit (and probably accurate) assumption that the spacecraft attitude is reasonably good (since in the worst case, attitude control was lost a moment previously) and a "safe-enough" ascent burn can be implemented by pre-programmed firing of thruster valves through manual command in a sequence. This extreme exigency is assumed necessary as indicators are in this circumstance that higher G&C functions are not reliable, and the Panic burn is the last measure available to prevent probably mission-ending contact of the spacecraft with the surface.

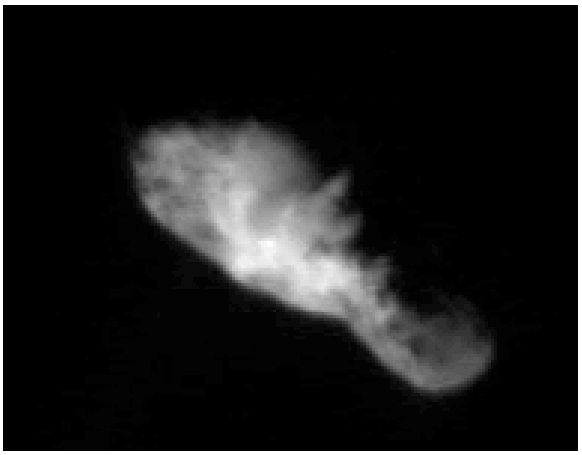

Figure 26. Deep Space 1 Picture of Comet Borelly as Captured by AutoNav, September, 2001; *AutoNav was first flown aboard DS1 as one of 12 new space technologies*

Inertial State Monitor: This monitor compares the current estimated position and velocity, against the nominal trajectory. If AutoGNC is functioning properly, the estimated spacecraft state should be within margin of the desired trajectory. If not, that fact constitutes a fault condition and the monitor detects that. This monitor in addition assesses attitude and attitude rate differences from commanded. There is not (necessarily) a pre-computed attitude profile in the form of an ephemeris that the spacecraft follows, instead, a series of commands requests that the spacecraft point in a particular direction to perform a burn, or to take images of the landing site. The attitude commander then computes the desired attitude quaternion and rate required for each moment, and the inertial state monitor differences these with the current estimated by the attitude estimator. When outside of a given range, which is time or altitude dependent, the inertial monitor declares a fault state.

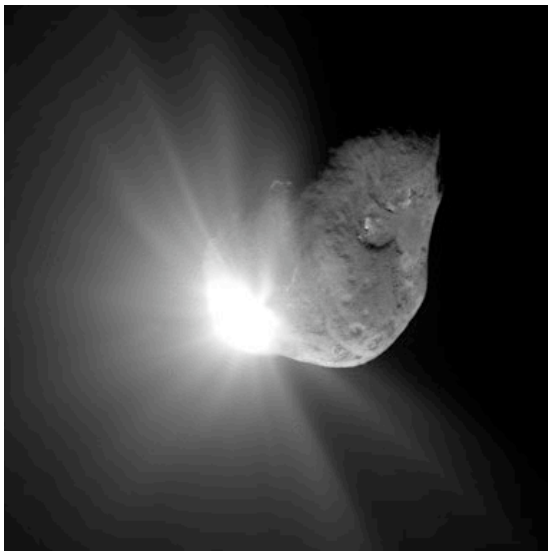

Figure 27. Results of the Deep Impact Impactor Spacecraft AutoNav Trajectory Targeting on July 4, 2005; *AutoNav operated for 24h prior to impact, and autonomously computed delivery maneuvers at 90, 25 and 12 minutes prior to impact.*

Altitude Monitor: It is the Altitude Monitor that performs the critical function of assessing the agreement of the altimeter instrument with the estimated state from AutoGNC. The altitude monitor also provides for the triggering of the calculation of the time-to-go timer discussed above.

Contact Monitor: Finally, the Contact Monitor is the simplest, of the monitors, but encompasses the very critical function of monitoring the contact sensor, and keeping the contact time-out timer trigger active. This and the previous two monitors are pictured in Fig. 25.

6. *Autonomous Navigation (AutoNav) and AutoGNC*

JPL's autonomous onboard navigation system (AutoNav) was initially developed beginning in 1996 for the Deep Space 1 (DS1) Mission (Ref. 6), which was the first flight of the New Millennium Program. New Millennium was specifically chartered to develop and fly new technologies. AutoNav was one of those technologies. DS1 fully checked out AutoNav and eleven other new technologies, and then proceeded to get the first detailed, close-up images of a comet nucleus, Borelly, in September of 2001, during a close flyby (Ref. 11, Fig. 26). Elements of AutoNav were used in the Stardust Mission flyby of comet Wilt-2 in 2004, capturing more images, and the very successful flyby and impact of comet Tempel 1 in 2005 (Ref. 12, Fig. 27). AutoNav is an ensemble of C-language routines that accomplish a number of primarily optical-navigation-related functions, including image processing and data editing, orbit determination and maneuver computation. For DS1, the AutoNav system was also responsible for implementing extensive periods of autonomous action including picture taking sessions, maneuvers, low-thurst engine activations, and computing orientation of the spacecraft during the flyby event. The DS1 system was implemented before the existence of VML, and therefore consequently all of the autonomously generated sequences to implement these activities were created with extensive mission-specific code in the AutoNav Executive that was discarded and re-written for each subsequent use of AutoNav(at substantial expense), despite the fact that the missions themselves shared similar characteristics. Four years ago it was decided to invest in an upgraded AutoNav system that would incorporate two important and fundamental changes: First, the Executive would be completely re-built in VML to allow for multi-mission adaptability of the overall system while preserving the flight heritage computational elements of AutoNAV, and second, attitude control and guidance capability would be added to the navigation functionality of AutoNav, providing for a generic and full 6-degrees-of-freedom knowledge and control computations. The latter development was in anticipation of much greater demands upon the onboard GN&C system anticipated in upcoming missions as will be discussed below.

The core of the AutoGNC system is the Deep Impact-derived AutoNav system, composed of three subsystems: Image Processing, Orbit Determination, and Trajectory Correction. The architecture appears in Fig. 20. Measurements for the OD filter come from a legacy image processing tool called OBIRON (On-Board Image Registration for Optical Navigation) [17][18][19]. OBIRON has space mission heritage as ground software for the Japanese Hayabusa mission to the asteroid

Dynamical equations of motion

- **Includes central body acceleration, 3rd body perturbations from other planets, solar radiation pressure, thrust from the ion engines, and miscellaneous accelerations**
- **2nd order differential equation modeled as two 1st order differential equations**

$$\dot{\mathbf{r}} = \mathbf{v}$$

$$\dot{\mathbf{v}} = -\frac{\mu_s}{r^3}\mathbf{r} + \sum_{i=1}^{n_p}\mu_i\left[\frac{\mathbf{r}_{ri}}{r_{ri}^3} - \frac{\mathbf{r}_{pi}}{r_{pi}^3}\right] + \frac{AG}{mr^3}\mathbf{r} + \frac{k}{m}\mathbf{T} + \mathbf{a}$$

where

$\mathbf{r}$ = the heliocentric cartesian position vector of the spacecraft

$\mathbf{v}$ = the heliocentric cartesian velocity vector of the spacecraft

$\mathbf{r}_{pi}$ = the heliocentric cartesian position vector of the *i*th perturbing planetary body

$\mathbf{r}_{ri}$ = the position of the spacecraft relative to the *i*th perturbing body

μ_s = the gravitational constant of the sun

μ_i = the gravitational constant of the i_{th} perturbing planet

n_p = the number of perturbing planets

A = the cross -sectional area of the spacecraft

G = the solar flux constant

T = the thrust vector from the ion engine

k = the thrust scale factor

m = the spacecraft mass

$\mathbf{a}$ = miscellaneous accelerations acting on the spacecraft

Given $\mathbf{q}^*$, the nominal trajectory parameters, as

$$\mathbf{q}^* = \begin{bmatrix}\mathbf{r} & \mathbf{v} & k & \mathbf{a}\end{bmatrix}$$

Filter estimates corrections, $\mathbf{q}$, to nominal trajectory parameters

$$\mathbf{q}(t) = \begin{bmatrix}\Delta x & \Delta y & \Delta z & \Delta\dot{x} & \Delta\dot{y} & \Delta\dot{z} & \Delta k & \Delta a_x & \Delta a_y & \Delta a_z\end{bmatrix}$$

The correction at time t is a linear mapping of the correction from time t_0

$$\mathbf{q}(t) = \Phi\mathbf{q}(t_0)$$

where Φ , the state transition matrix, is defined as

$$\Phi(\mathbf{t}) = \frac{\partial\mathbf{q}^*(t)}{\partial\mathbf{q}^*(t_0)}$$

The Partial derivatives of the observed pixel and line locations, ***p,l*** with respect to the state at time ***t*** is:

$$H(t) = \begin{bmatrix}\partial p/\partial r & 0_{1x7}\\ \partial l/\partial r & 0_{1x7}\end{bmatrix}$$

This can be mapped back to the epoch, t_0, via the the state transition matrix:

$$\tilde{H}(t_0) = H(t)\Phi$$

The minimum variance least squares solution to the epoch state corrections is:

$$\hat{q} = \left[P_0 + \tilde{H}^T W\tilde{H}\right]^{-1}\tilde{H}^T WY$$

Where:

P_0 is the a priori covariance of the state parameters.

W is the weighting values of the pixel and line observables.

Y is the residual vector between the observed pixel/line locations and their predicted values.

Figure 28. Mathematical formulation of the AutoGNC/AutoNav Orbit Determination Filter; *AutoNav utilizes a batch-sequential, least-squares, U/D factorized sliding-window data filter*

Itokawa, as well as NASA's NEAR mission to the asteroid Eros. OBIRON is currently being configured as a flight element for NASA's Autonomous Landing and Hazard Avoidance Technology (ALHAT) development that will support manned lunar landing projects. OBIRON will also be the means of performing ground-based optical navigation for the Dawn mission. OBIRON combines preloaded surface elevation maps, camera configuration, estimated camera location and attitude, and the sun illumination direction in order to predict what the camera should see in the vicinity of specified surface landmarks. Landmarks are then registered with the image through a correlation process, resulting in sub-pixel coordinates for the landmark locations in the camera plane. These coordinates are then passed to the filter as raw measurements for orbit determination.

The Orbit Determination element of AutoGNC uses a batch-sequential least-squares estimation strategy rather than a current state filter. This choice was made to allow for the comparison of data in a data arc for purposes of data rejection. The data arc is determined by a sliding window, and thus periodically the time of the estimated position and velocity must be advanced along with the estimated covariance. The latter step is performed using a numerically integrated state-transition-matrix. This formulation is shown in Fig. 28. Currently, data from the IMU is not incorporated directly into the filtered spacecraft state estimate, but is used as a "truth" acceleration when the spacecraft state is integrated, and thus IMU readings of propulsive events and the contact event are incorporated indirectly into the spacecraft state estimate. Altimeter measurements are not currently included in the estimated state, but are used as a "sanity check" against AutoGNC results. The trajectory correction component of AutoGNC uses a linear estimate of required velocity change to achieve a desired position or velocity state. The current very successful targeting strategy for the TAG scenario is to use position targeting only, and to target such as to achieve the desired state at the time of the next trajectory correction, which will then (under perfect conditions) correct the velocity back to the desired nominal course. In practice there are always errors, but many scenarios including TAG, this series of corrections represents a convergent algorithm. In the case of the TAG the approach readily achieves the required error accuracies of less than 5 m and 2 cm/sec. AutoGNC includes a set of utilities and services, chief among which is the Ephemeris Service which provides positions and velocities of celestial bodies, their attitudes and the position and velocity of both the nominally planned and the currently estimated spacecraft trajectory.

The attitude control and guidance system of AutoGNC is derived from several flight sources, including Cassini, DS1, and MER. An important recent upgrade to AutoGNC was to create a flight-prototypical commanding method for spacecraft attitude that was sufficiently general to put the spacecraft into any desired attitude, but also sufficiently comprehendible to be usable by typical flight teams. As part of this attitude specification system is an automated means to re-orient the spacecraft in a directed manner, specifying maximum rates and accelerations. The system also includes provision for attitude constraints, for example sun-exclusion zones for cameras and other instruments. The system just described constitutes the "Attitude Profiler", which is the first of three Guidance and Control elements of AutoGNC, followed by the Attitude Estimator and the Attitude Controller.

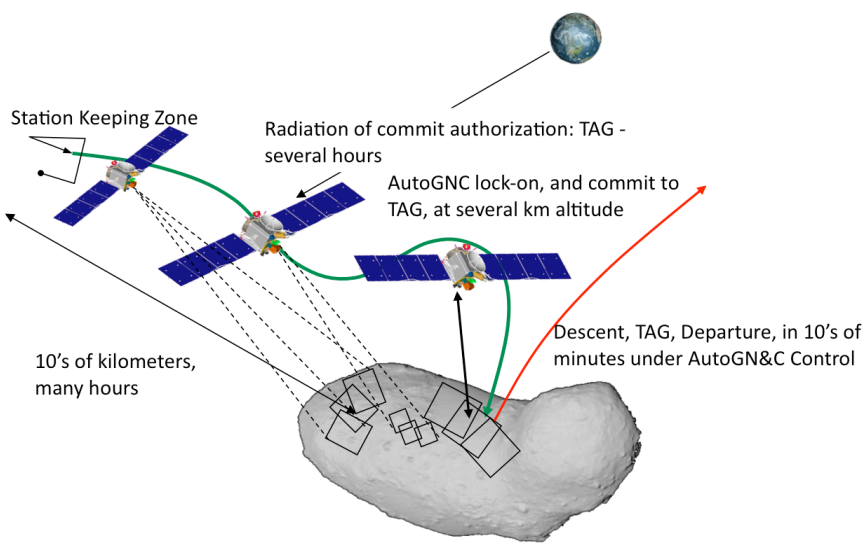

Figure 29. Touch and Go GN&C - Rendezvous, Contact and Sample Recovery and Return: *Many if not all of the operations required for TAG are required for Altairs descent, and precision landing.*

The attitude profiler is a high-fidelity software element of Cassini heritage with a high NASA Technical Readiness Level (TRL). The profiler has

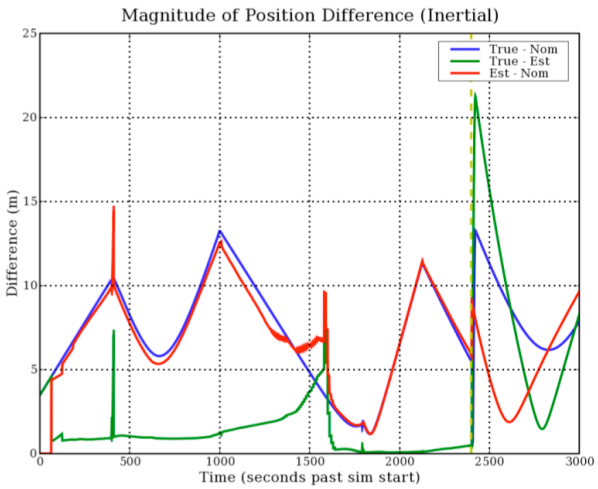

Figure 30. TAG OpNav Performance of AutoGNC - Position Accuracy; *"True" is achieved, "Nom" is nominal desired, "Est" is AutoGNC-estimated trajectory*

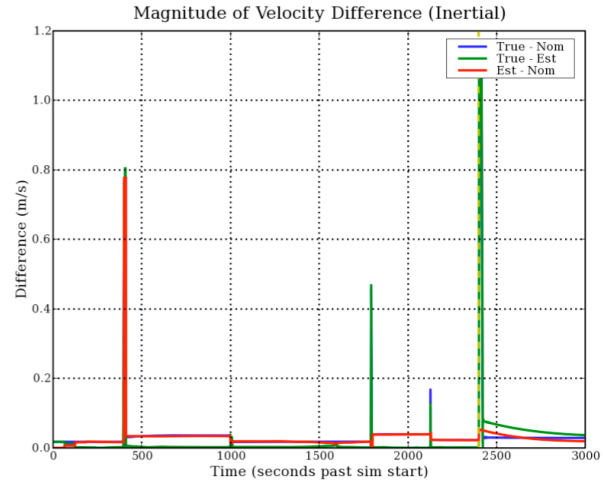

Figure 31. TAG OpNav Performance of AutoGNC - Velocity Accuracy; *"True" is achieved, "Nom" is nominal desired, "Est" is AutoGNC-estimated trajectory*

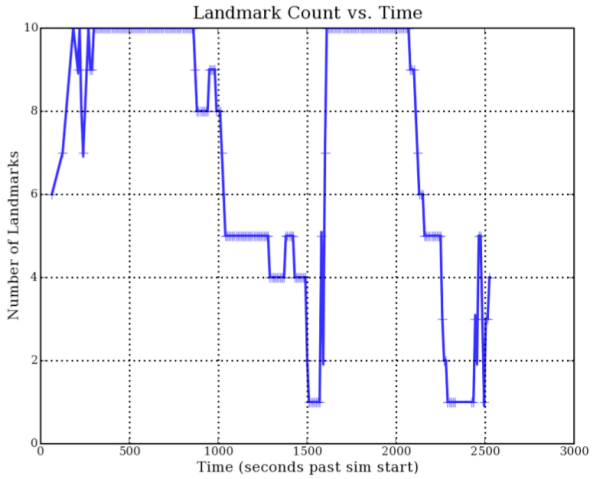

Figure 32: Landmarks Processed by and used in AutoGNC OpNav Filter as a Funciton of Time in the TAG Event; *The number of landmarks has a ceiling of 10 to match the expected CPU throughput and AutoGNC allocation expected*

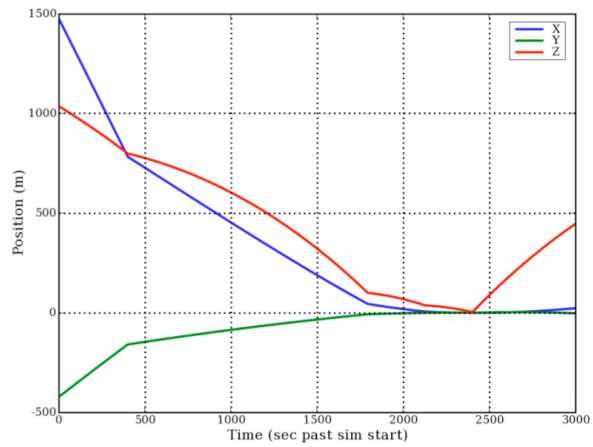

Figure 33. TAG Landing Site Frame Trajectory Nominal Positoin: *Z is nomral to mean surface, X, Y parallel*

access to ephemeris files for all relevant target bodies, and uses a trapezoidal rate profile to plan a turn from the current attitude to the desired target. Primary axis targets can be provided as body centers, body surface coordinates, latitude/longitude coordinates, inertial coordinates to a star, or a desired inertial thrust direction. Secondary axes are typically specified to be nearest another vector, such as the Sun direction or the pole of the target body. When not actively commanding a turn between targets, the profiler operates in a tracking mode using the current ephemeris files, allowing the spacecraft to perform operations such as science imaging or communications with Earth. Constraint avoidance algorithms are available in the software but have not yet been tested in simulation.

The attitude estimator is a Kalman filter based on fixed pre-determined gains. The estimator processes IMU and star tracker measurements to compute an estimate of the attitude quaternions and IMU biases. Although the IMU measurements are used to propagate the attitude quaternions at 10Hz, the Kalman filter itself runs at 2Hz. The attitude controller computes error from the desired and estimated quaternions and rates. It then computes commanded torques from the errors using a non-linear proportional-derivative control law.

The commanded delta-V and torque are combined and transformed into thruster commands using a pulse-width modulation thruster allocator. The allocator assumes a fixed thrust level for each thruster. Setting the thruster on/off time every 0.1 second control period modulates the thrust. The commanded delta-V for the duration of the maneuver is divided into increments corresponding to a commanded impulse to be applied in a control period. At each control period, that commanded impulse is transformed into a commanded force, and combined with the commanded torque produced by the attitude GN&C system. The thruster allocator is provided with information like thrust magnitude, thruster location and thrust direction for each thruster. It is also given a list of thrusters to use for each mode of operation. The thrust allocation problem is then solved using an optimizer to find the set of thruster commands providing the commanded force and torque while minimizing fuel consumption.

VI. Utilization of AutoGNC for a Small Body Sample Return "Touch and Go" (TAG) Mission Scenario, as a Viable Analysis Model for an Altair Lunar Landing

A. TAG Mission Simulations

One of the first mission challenges for the newly minted AutoGNC system is likely to be a sample return from a small body, either an asteroid or a comet, or perhaps a moon of Mars. These missions require tight coordination between navigation and attitude guidance and control. In particular, the scenarios being most closely studied utilize a "Touch and Go" strategy, where the spacecraft flies close by the target, and drops to the surface, making glancing contact to collect a sample, and then rapidly ascends to safety. The close interaction between Navigation and Guidance occurs during contact, where often severe torques and forces are applied to the spacecraft at a time with tightly prescribed attitude and contact force is mandated for the sampling mechanism to properly function. Shortly after first contact (sometimes as little as 1 second), the spacecraft is then going to begin the ascent thrusting event, during which time the spacecraft attitude is also tightly constrained so as to achieve the correct trajectory.

1. Typical timeline of activities for a small-body TAG mission

In general for such mission scenarios, the spacecraft is in a holding orbit before commencement of TAG activities. Depending on the mass of the body, this holding trajectory may be at a range of a few to several 10's of km. Still under ground-control (i.e. deterministic sequence) the spacecraft will depart on a low altitude flyby trajectory which has very low impact probability. For the first 12-18 hours of the approach to the target (a period of time that is generally constant despite target mass and size) the AutoGNC system is operational, determining spacecraft position velocity, and making course adjustments, and thereby allowing ground control sufficient time to assess AutoGNC system health. If health checks are passed, then the ground will grant authority to perform the Drop Burn, placing the spacecraft on an impact trajectory. If the descent commitment time expires, the spacecraft

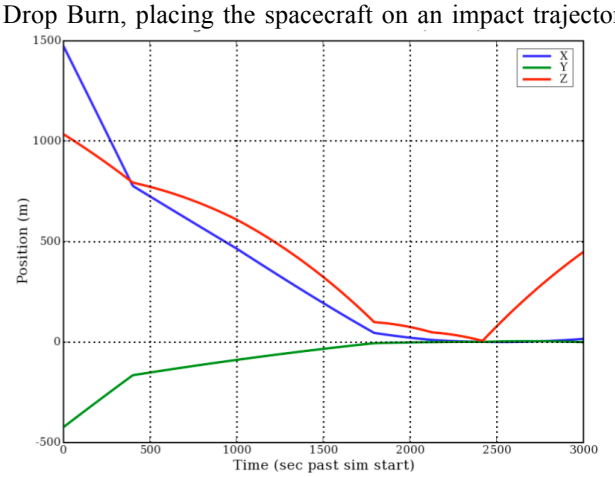

Figure 34. TAG Landing Site Frame Trajectory Truth Positon: *Z is normal to mean surface, X, Y parallel*

Figure 35. TAG Event Velocity in Landing Site Frame; *Z is normal to surface, X, Y, parallel to mean surface, desired contact velocity is 20cm/s*

skips the descent opportunity and flies safely by (Fig. 29). A critically important feature of the VML AutoGNC Executive is the ability to self-assess the health of AutoGNC subsystems, and in the period before the Drop Burn, the Exec itself will be able to wave off the TAG attempt. Prior to the Drop Burn the the solar arrays are feathered for protection from contact damage and to lessen dust deposition, and there occur several targeting burns. These events are triggered by the Executive and auxiliary manager functions as discussed in later sections.

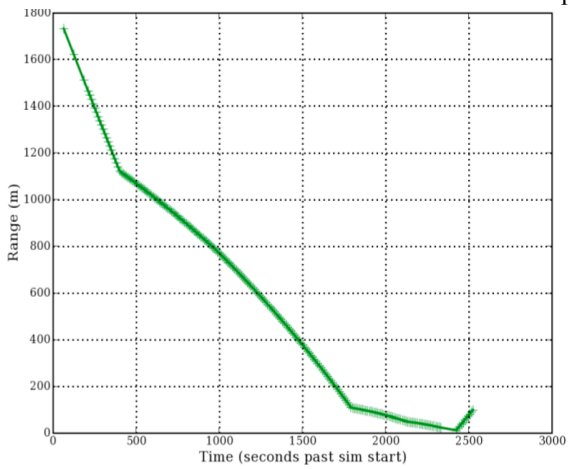

Figure 36: Range to the Landing Site During TAG Event; *To be compared to altitude as shown in the next figure*

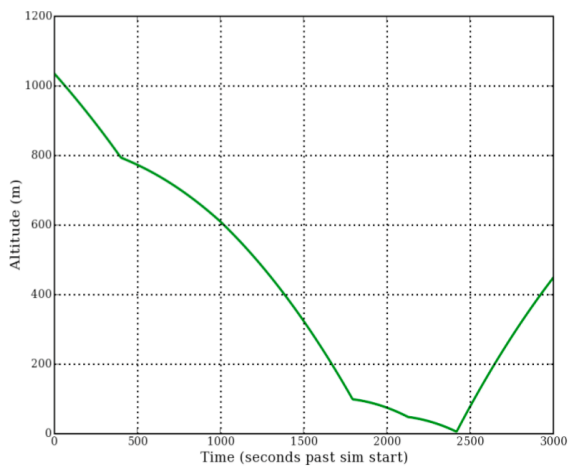

Figure 37: Spacecraft Altitude Relative to Surface Throughout TAG Event

Once the Drop Burn executes, the spacecraft is wholly dependent upon the AutoGNC Executive system to maintain the mission's survival, and internal Executive health-checks continue and if any subsystem fails, an abort maneuver is triggered to send the spacecraft on a safe trajectory away from the target. The period of time for this dependence on the autonomous system is roughly an hour. A wide range of factors are monitored by the Exec, including the spacecraft position, velocity and attitude parameters being within desired ranges from the nominal plan, that the orbit determination system is receiving sufficient data, and that the trim maneuver computations are successful; there are from 1 to several course correction burns during the descent that may have severe restrictions on burn direction in order to minimize surface contamination. Several terminal descent timers come into play during descent, providing "life-boat" triggers for critical events in the TAG scenario. This approach compensates for a variety of different faults, ranging from a failed contact sensor that doesn't detect the surface to highly unusual surface conditions (like a very light and powdery regolith) which cause the sensor not to respond. Below a certain altitude, laser range finding provides an estimate of the time to contact. If contact-sensor readings determine contact is made, the spacecraft remain in that state for a predetermined number of seconds and the sample is collected. After that, or after the time window for attempting contact expires, an ascent is initiated. Errors during the descent, contact, or ascent phases result in an abort burn. The Executive determines two genre of abort burn, a nominal abort which utilizes the conventional attitude control computations, and a "Panic" abort, wherein the Executive directly commands a series of thruster events that have been precomputed to result in an appropriate thrust direction. This latter choice is made only under the most extreme conditions, when the Exec has determined that the AutoGNC system is failing in some way, and its results are not to be trusted; the core assumption in this "panic" implementation is that AutoGNC was successfully maintaining attitude until just a moment before the "panic" and therefore an reasonably descent ascent burn can be achieved without active attitude control. This extreme contingency is the last resort attempt to salvage the spacecraft, because no action in this event will lead to the spacecraft colliding with surface within minutes.

2. Results from a Comet TAG operation – inertial frame performance

Figure 30 presents the first of several presentations of the performance of AutoGNC in the TAG scenario. It shows the magnitude differences between the desired nominal, the estimated, and the true trajectories. The knowledge of the AutoGNC estimator (green line) is very good. Even as the true trajectory deviates from the nominal trajectory (blue), orbit determination is able to use optical observations and nongrav information during maneuvers to keep the knowledge error small. Velocity errors are shown in Fig. 31. Knowledge errors appear in several places, including at the start of the simulation run where initial conditions are not perfect, and during maneuvers due to a finite-burn modeling discrepancy which has been subsequently corrected. Additional knowledge error is introduced due to the accelerometer bias during the maneuvers. Note that from the time of the second maneuver at 1000 sec to about 1600 sec, the knowledge error climbs steadily until it is suddenly removed. Compare this section of the plot to the same time range in the landmark count plot in Fig. 32. The loss of knowledge is a direct result of the lack of landmarks passed to OD, the reason for which is explained below. When the number of landmarks is restored to the maximum of 10, the knowledge error is driven back down.

Nominal contact is at 2400 sec, and contact detection is actually first made at 2413.98 sec. There is a knowledge error at contact due to contact forces and the fact that the contact time is not exactly what the filter was given a priori. In this run, no calls to OD were scheduled after contact, so any knowledge errors from contact and ascent are simply propagated to the end of the run. In future runs, OD will be scheduled to run after contact to monitor the status of the ascent. Control errors in this plot (blue) are typically introduced by execution error. There are also contributions from finite burn modeling errors, and from the knowledge error. AutoGNC calculates each corrective maneuver based on a recent estimated state, and errors in the estimate will result in imperfect maneuver solutions. The apparent control error at contact is the result of the timing error in contact time. Since the nominal trajectory is built "on the ground," the ascent maneuver time is fixed and will deviate from the true trajectory, where the ascent maneuver is based upon the contact sensor.

Again, Fig. 32 shows the magnitude of velocity differences between true, nominal, and estimated trajectories. Velocity data come from the exact same sources (and at the exact same times), as the position data in Fig. 30. It is interesting to note in simulation runs with initial velocity errors that it takes OD two batches of landmarks separated in time to properly solve for the velocity of the spacecraft. After the second image and OD, the velocity error is essentially removed.

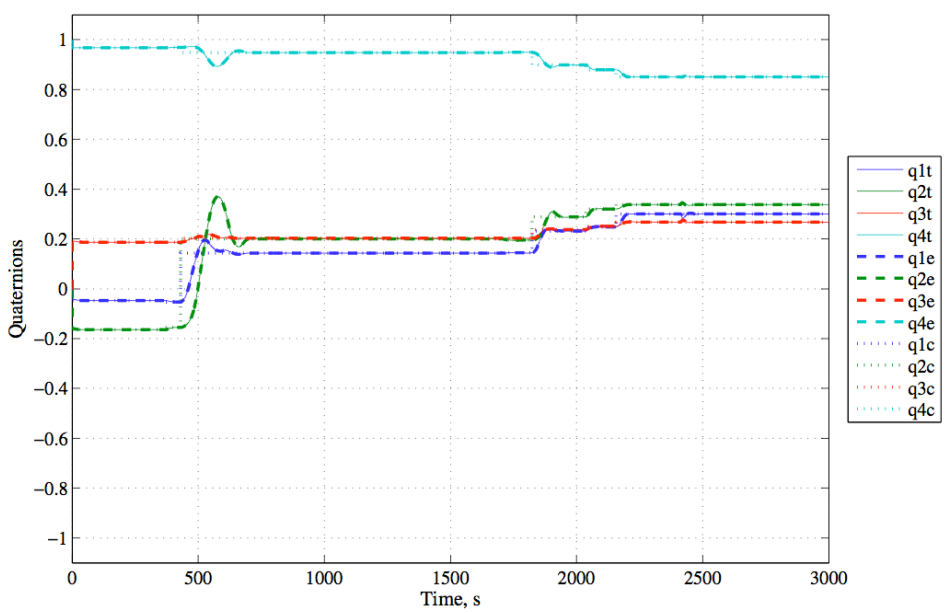

Figure 38. Truth (t), Estimated (e) and Commanded (c) Spacecraft *vs.* Inertial Attitude Quaternions Throughtout TAG Event; *Notable characteristic of the attitude performance is the near identity of the 3 sets of attitudes, except immediately after a commanded turn – these divergences are due to the early version of the non-planning attitude commander.*

There are several spikes in the velocity error plot that result from the difference between how the maneuvers are modeled and how they are executed. Currently, the maneuvers are modeled in AutoNav as impulsive burns, but are executed by DSENDS (the dynamic simulation subsystem) as finite burns. Thus there is a velocity error during the execution of the burn, and for long burns there can be a small residual velocity error after the burn is complete. As mentioned above, the accelerometer bias also introduces a velocity knowledge error.

The nominal contact time is indicated on the plot by a dashed yellow line. The large knowledge error at contact is again due to the timing of contact. The OD filter a priori expects the ascent maneuver to happen at a nominal time. If it does not happen at that time, then the estimated state will deviate from the true state until OD is run again and the state is updated with current nongrav history information. This results in the brief spike in the green line.

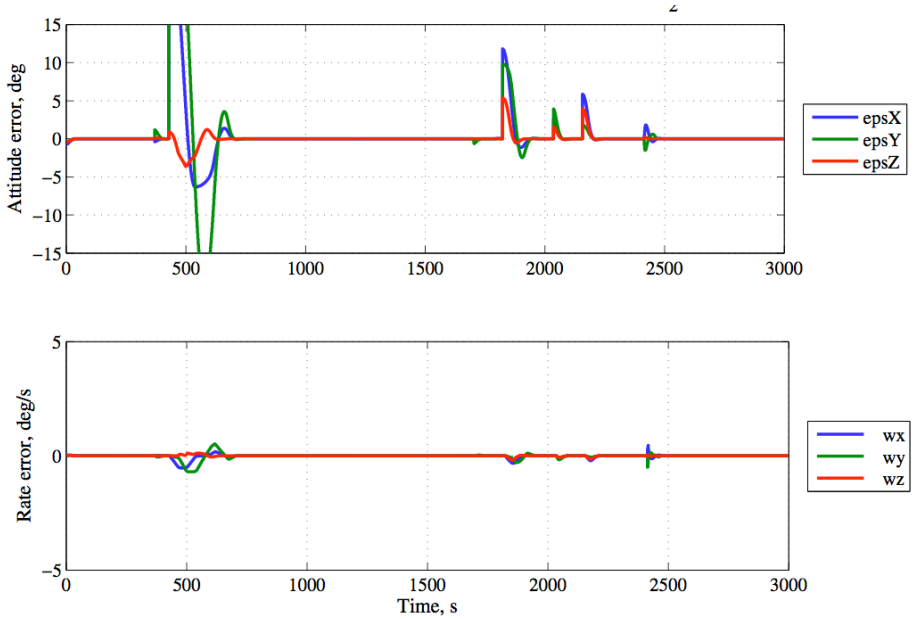

Figure 39: Spacecraft Attitude Errors Throughout the TAG Event; *Measured as deviations in angle and rate from desired profile in principal spacecraft coordiantes X,Y,Z*

3. TAG performance, landing-site relative.

Figure 33 shows the nominal/planned position of the spacecraft in the landing-site-local-vertical-horizontal frame. Several things can be observed in this plot. Since it is the nominal position data, it describes the trajectory that is intended for the rendezvous spacecraft. At the time of this simulation, maneuvers in AutoGNC were currently planned and computed as impulsive burns (as was practiced during the DS1 and Deep Impact missions), so distinct corners are evident in the trajectory. TAG is evident as all three components of the plot converge on 0; however, it can be noted that the X and Y components approach zero long before the Z component. This is because the last portion of the trajectory is a vertical descent onto the landing site. The nominal ascent maneuver has the effect of rapidly increasing the Z component of the landing-site frame trajectory, but the X and Y components are largely unaffected. This demonstrates that the ascent maneuver, which is planned and executed by the VML executive, is in fact in a direction normal to the landing site.

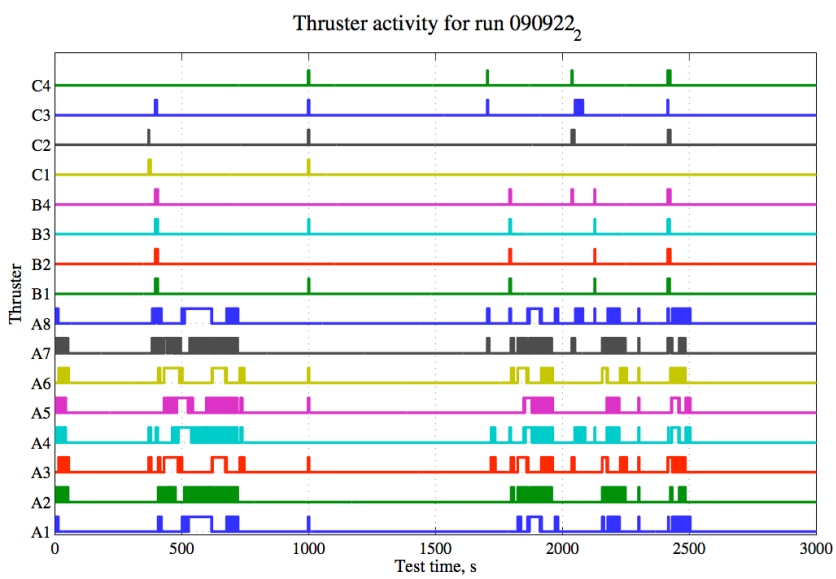

Figure 40. Commanded Thruster Activity During TAG Event; *3 sets of thrusters are utilized, 8 1N attitude control thrusters (A's), 4 upper-facing 20N TCM thrusters (C's), and 4 downward pointing 200N main ascent engines (B's)*

Figure 34 shows the truth position in the landing-site frame, to be compared and contrasted with Fig. 33. Ideally, these plots will appear to be nearly identical. Because the DSENDS executes the maneuvers as finite burns, the maneuvers can appear 'rounded', instead of having sharp corners like the nominal plot. This also can sometimes be difficult to see, since the length of each maneuver

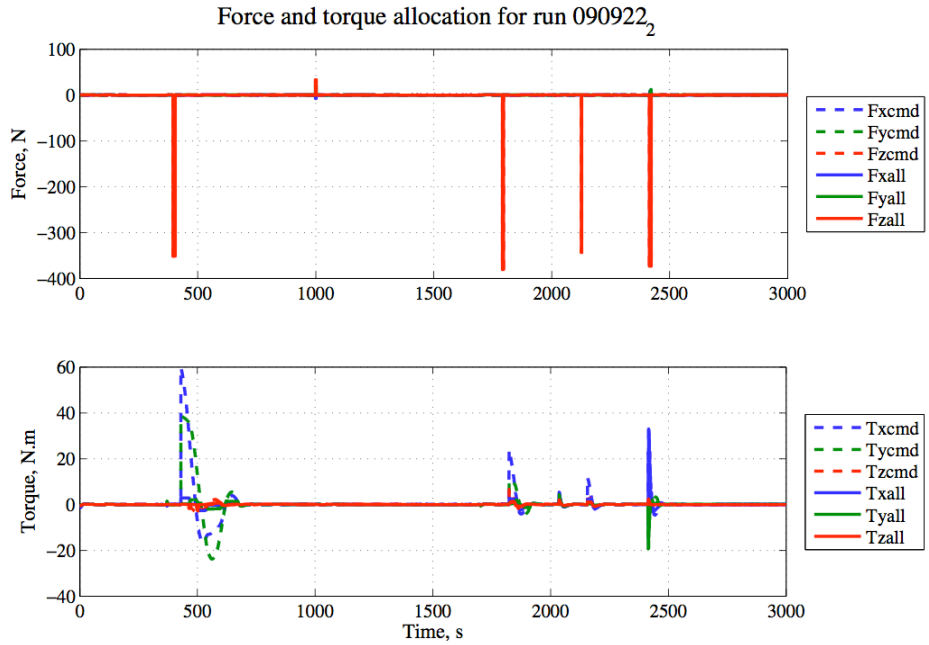

Figure 41. Force and Torque Commands and Allocaitons in the TAG Simulation; *The full-6DOF TAG simulation includes a realistic spacecraft and engine model, and uses a flight-derived attitude controller and allocator to maintain attitude, "cmd" denotes "commanded," and "all" denotes "allocated."*

is small compared to the length of the whole trajectory. Although difficult to see on the scale of the plots presented, Figures 33 and 34 together help show why the actual contact happened 13.98 seconds late in this run. First, recall that the nominal trajectory is generated on the ground. It can be see from Fig. 34 that the spacecraft velocity in the Z direction reverses at 2402 sec. This represents a nominal contact at 2400 sec, then 2 seconds on the surface, and then the ascent burn. At the time of nominal contact, the altitude in the Z direction is 4 m. The X and Y components are small, 4 cm and 12 cm respectively. In Figure 8, however, the Z component is 9.4 m, while X and Y are 64 cm and 36 cm respectively. So the resulting maximum trajectory control error of 5.4 m is almost exclusively in altitude, which translates into time-of-flight. Thus, the spacecraft is late to touch down, but lands on the desired target within half a meter.

Figure 35 shows the nominal/planned velocity in the landing-site frame. The nominal velocity plot in the landing-site frame has several interesting qualities. Impulsive maneuvers are easily recognizable as discontinuities in the lines. The acceleration due to gravity is evident in the Z component as a slope in the line, tending to increase the spacecraft velocity in the -Z direction. This is punctuated by the braking burns, which slow the rate of descent. The contact velocities can be read from the plot; vertical contact speed from the Z component, and horizontal speed from the X and Y components. The intention is for the horizontal components of velocity to be nearly zero, while the downward component is 20 cm/s. Figure 36 shows the truth velocity in the landing-site frame to be compared with Fi. 35. As with the position plots, the finite maneuvers that are executed by DSENDS can be seen in the slopes between the discontinuities in the curves; these lines are strictly vertical in the nominal plot. This plot also indicates the actual contact velocity in the vertical and horizontal directions, compared to what was intended in Fig. 35 above. The vertical component of the velocity from the truth plot is 22.5 cm/s, or 2.5cm/s faster than the nominal velocity. It can also be noticed that there appear to be more maneuvers in the truth plot

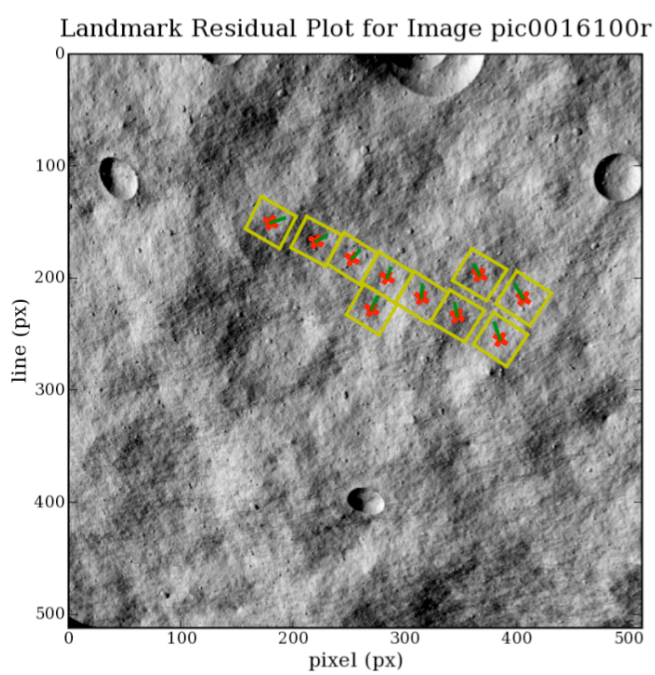

Figure 42. Sample OpNav Picture from TAG Simulation; *Shown are the set of landmarks used in this frame, and the post-processing landmark residuals scaled by x10*

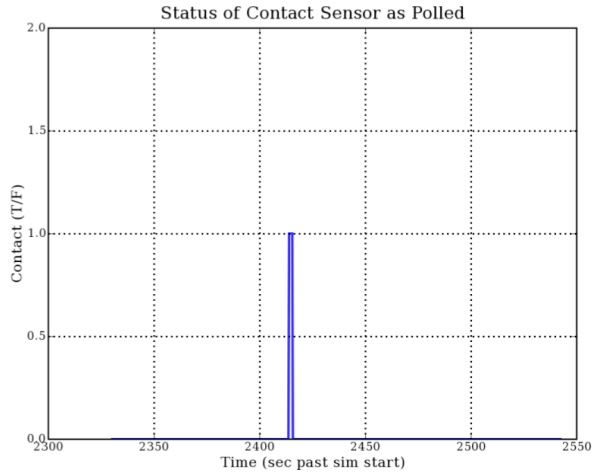

Figure 43. Contact Sensor Readings Throughout TAG Simulation; *Accurate emulation of physical contact, and appropriate sensing/response to contact detection is a critical component of the TAG simulation, as failure could well result in loss of mission*

than in the nominal plot. This is, in fact, true. The nominal trajectory is designed with several placeholder maneuvers in place that have a zero magnitude; in flight, these "statistical" maneuvers are implemented to correct control errors between "deterministic" maneuvers.

4. Range and Altitude Performance

Figure 36 shows the range to the landing site as a function of time. The data here is recorded at image times, and those times are indicated with marks on the plot. Note that this plot does not necessarily indicate a height above the surface, except in the case of a vertical descent; rather, it is the distance from the camera to the landing site (image target), which is useful in computations of camera pixel size, and also in the selection of which surface model to use when rendering images. Note that the ascent is much steeper than the descent, indicating the speed of the approach versus that of the ascent. Figure 37 shows the altitude of the spacecraft with respect to the landing site. This figure may appear very similar to Fig. 36, however this plot does actually represent vertical altitude instead of range. The altitude is calculated at image times, and with respect to the surface target landmark. It assumes that the landing site is an infinite flat plate and does not attempt to account for the curvature of the body.

5. Attitude Control Performance

Figure 38 shows the spacecraft attitude quaternion as a function of time. Three sets of lines indicate the true attitude (solid), the estimated attitude (dashed), and the commanded attitude (dotted). A slight curvature at time zero indicates a small initial attitude offset that is promptly removed by the controller. Several discrete steps in the commanded attitude are quickly matched by the controller with a small amountof 'ringing' as the desired attitude is reached. The discrete steps in the commanded attitude are a result of the completion of a deterministic maneuver, which results in the attitude profiler changing the pointing direction to match the next maneuver direction. Note also that the spacecraft attitude is constant after the final maneuver, through contact and beyond. Figure 39 shows the spacecraft pointing error and angular rate error in the body frame as a function of time. Large discontinuities are present after the maneuvers, when the pointing profiler discretely changes the desired pointing to the next maneuver. Small disturbances in the attitude can also be seen during the maneuver, as the allocator is forced to allow some torques in order to satisfy large thrust demands on the system.

Figure 40 shows the activity of each thruster as a function of time. A given thruster is saturated when there is a solid plateau in the appropriate curve. When the curve appears to be solid, there is frequent pulsing of the thruster but it is not saturated. It can be observed that the translational thrusters (sets B and C) are used almost exclusively during translational maneuvers, which are then followed by RCS activity to recover from induced torques and slew to the next commanded attitude. There is also RCS thrusting at the very beginning of the simulation as the attitude controller removes the initial attitude error. Figure 41 shows the commanded and allocated force and torque as a function of time. In both subplots, two sets of lines are present: the allocated thrust (solid) and the commanded thrust (dashed). It is interesting to note that during the large discrete changes in attitude after maneuvers, the requested torque is far above what can actually be allocated by the system. In general, attitude control is good, including the period during contact starting around 2414 sec.

6. Other measures of performance

Figure 31 above shows the landmark count as a function of time. These are the landmarks that are passed to AutoGNC as measurements for use in orbit determination. The data for this plot is naturally generated at image times, so the image frequency can be noted from the markers, in addition to the number of landmarks. The maximum number of landmarks allowed per image in this simulation was 10. It was mentioned above that a drop in the landmark count between about 1000 seconds and 1600 seconds is responsible for an increase in the knowledge error. The cause for this dip comes from how the landmarks are generated, which is "on the ground," and by hand. Two batches of landmarks have been created for this scenario, all centered around the landing site. On set has a 1m/pixel resolution, while the other set has a 10cm/pixel resolution. OBIRON is configured to compare the resolution of the landmark with the resolution of the image, and will only attempt to correlate a landmark if the ratio of the resolutions is less than 6 (in either direction). The angular extent of a pixel in the wide angle camera used by this scenario is approximately 1mrad, so the 1m landmarks are valid at ranges between 6km and 167m, while the 10cm landmarks are good at ranges between 600 m and 16.7 m. The range plot in Fig. 36 indicates that the range during the period in question (1000 sec – 1600 sec) is about 800 m down to about 350 m. This is a range where both sets of landmarks are at the limits of their usefulness. While some landmarks are available for use, they are severely scaled and apt to be deleted. When the range decreases to a point where the 10 cm landmarks are beginning to correlate well, the landmark count improves, and with it, the navigation. This is something to consider in future simulations and the actual mission. Figure 42 is an example of one of the first images where the high resolution landmarks are correlated and used by OD. Note that all 10 landmarks are used and the residuals indicate the error that OD is quickly able to solve for. A drop in the number of landmarks used can be seen near contact; the landmark count then grows again after contact. This effect is caused by the very near proximity of the surface at contact. The surface has been tiled with landmarks a priori by the Opnav team, but at very close range there are very few landmarks in each frame for the system to use in landmark tracking. Additionally, the a priori landmarks have inherent resolutions that can become incompatible with the extremely high resolution images taken near contact (see above, the limit for the high resolution landmarks is approximately 16 m altitude).

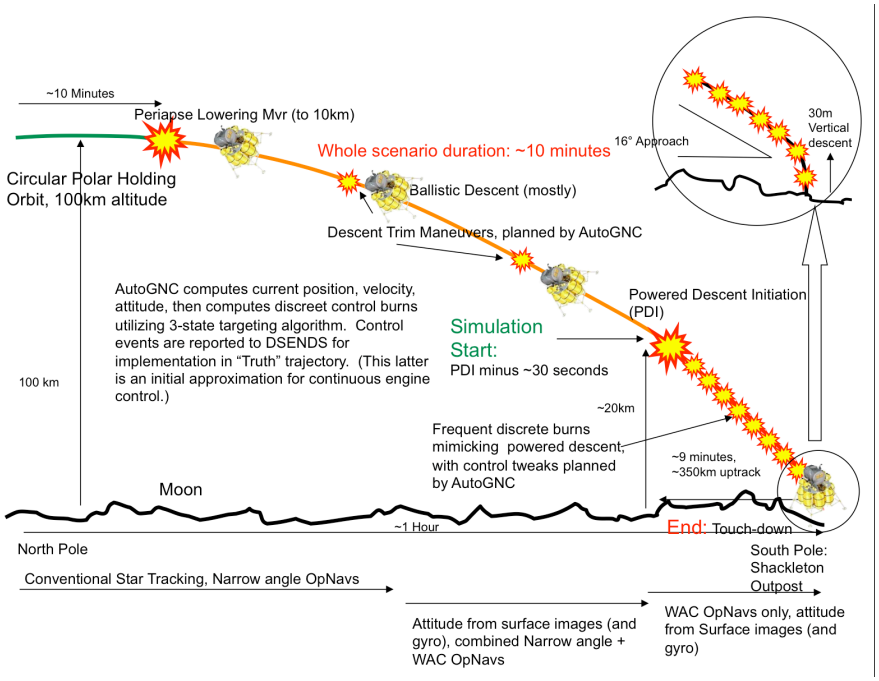

Figure 44. Schemmatic of Powered Descent Phase of Altair Lunar Landing Simulation*; Simulation is of the baseline 16deg glideslope Apollo-like trajectory*

Figure 43 shows the status of the contact sensor as a function of time. This plot indicates a single spike at contact, which is expected. The contact sensor is enabled on the terminal approach and should report "False" until contact with the surface is detected. The contact sensor then indicates "True" for contact as long as the sampler is on (or below) the surface, and then turns off again when the spacecraft departs. Thus, this plot is an indicator of contact time, contact duration, the basic operation of the contact sensor, and the presence of potential multiple contacts.

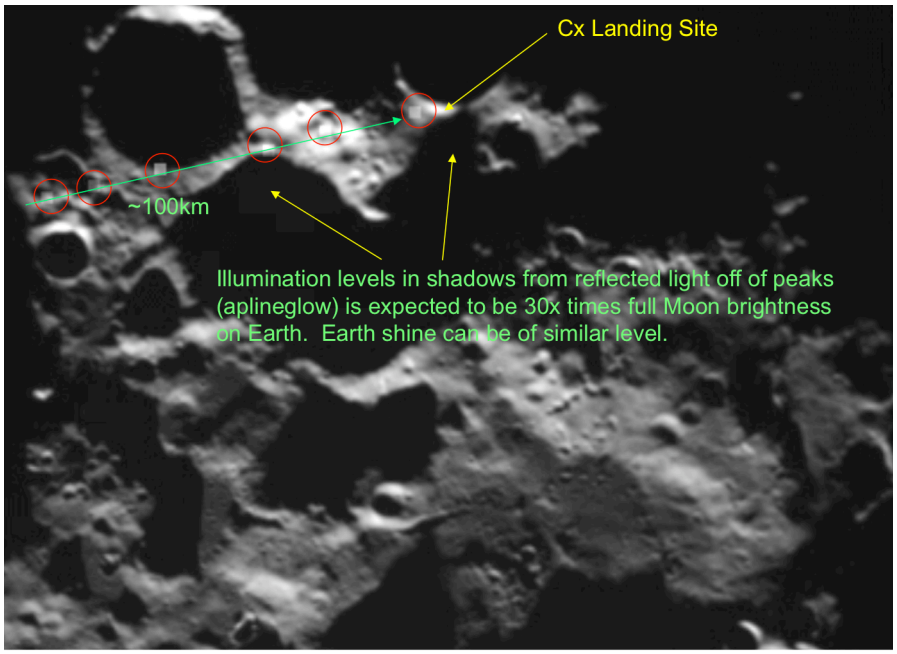

Figure 45. Clementine VIS image of the Shackleton Region, with an Overlay of the Simulated Landing Trajectory; *100km of the ~350km powered descent, showing several navigation landmarks and the landing site.*

B. Simulation of Lunar Landing using AutoGNC

1. Overview of the Lunar Landing Simulation

The thorough discussion of the small-body mission emulations using AutoGNC discussed above is highly relevant to the application of the system to an Altair landing, as most – if not all – of the Executive, navigation, trajectory and attitude control susbsystems applicable to small body operation also apply to the lunar case. Using the AutoGNC simulation system described above, an emulation of the landing of Altair at Shackleton has indeed been developed, and another landing emulation at a mid-latitude Constellation landing site is currently in work. This simulation follows the Altair landing through the powered descent, and is shown schematically in Fig 44. At the time of the development of this simulation, neither GSSR (Ref. 13) nor LRO data was available at the South Lunar Pole to develop maps of sufficient accuracy (up to 0.5cm) necessary for the simulation. Even at this writing, LRO data has not been assembled into a map, and the GSSR data presents substantial numbers of radar shadows in the South Pole region where some early Altair project landing is anticipated. The first step in the reference map-generation process process was to use the 1994-era Clementine data to create a map of the Shackleton region to the maximum achievable resolution allowed by the available data.

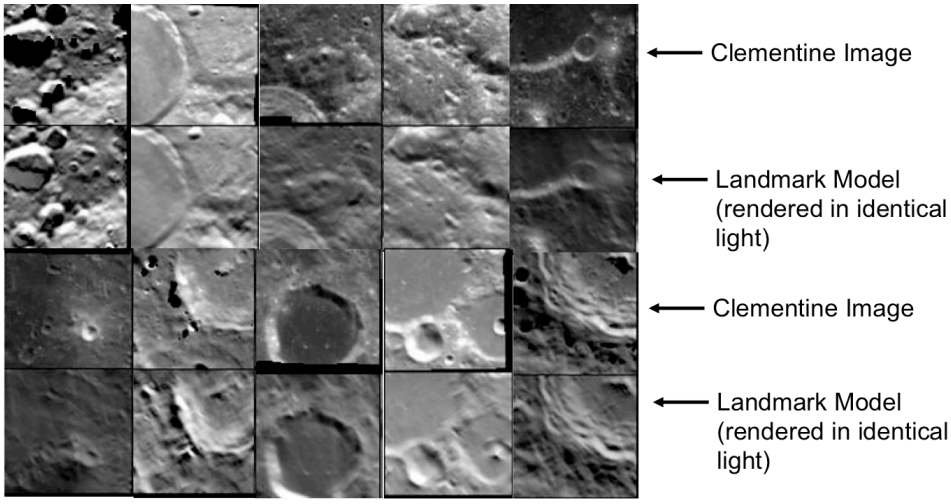

Figure 46: Stereophotoclinometrically-Generated Landmark Maps and Representative Clementine Images; *Clementine images were the basis for generating the landmark digital elevation maps used in the lunar landing simulation*

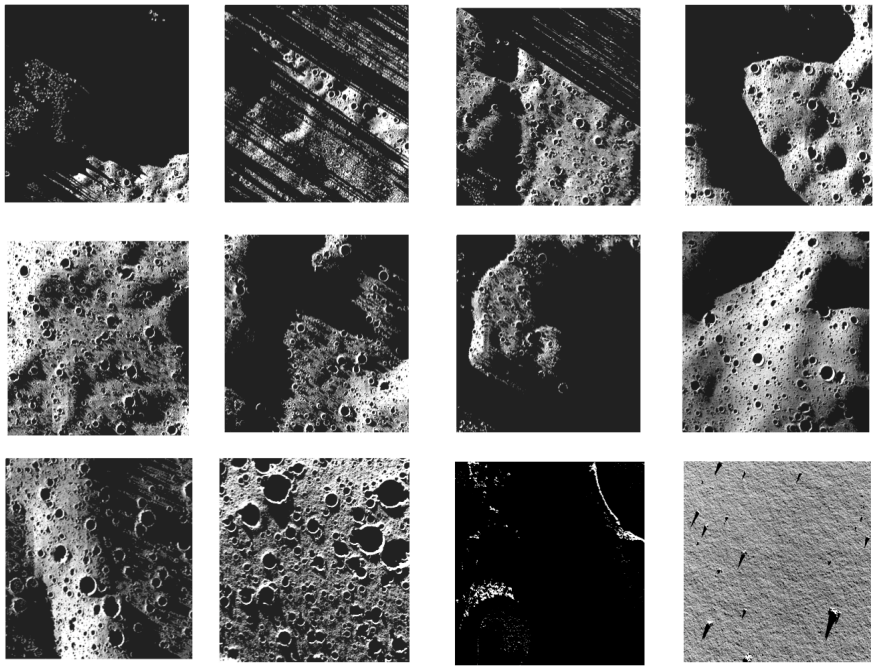

Figure 47. Representative Navigation Frames from the Simulated Power Descent, Starting at 21km Altitude, and Ending at the Initiation of the Vertical Descent at 30m Altitude; *The modeled terrain is based on Clementine Mission imaging data, and is at best of 200m resolution; necessarily much finer detail is present in these models, and was created using realistic sequential texturing processes, including cratering rock distribution and weathering processes.*

Though the Clementine survey mission was during the southern winter, and therefore had very poor visibility, it did survey sufficiently to reveal a "light-bridge" which would make it possible at a particular time of the lunar day to traverse to the Shackleton crater with directly illuminated terrain nadir to the spacecraft for the whole descent, even in that light-challenged season. This is a very conservative result and beneficial to the Constellation Program. It is important to note that nadir or nearly-nadir observations of the terrain are not necessary, as with a gimbaled camera it is possible to slew up toward a horizon to capture the many illuminated prominences that would be visible from almost any vantage and altitude.

Modelling the Shackleton Lunar Terrain

With the Clementine data in hand, stereophotoclinometry (Refs. 9,10) was used to develop the maps at the ~200m resolution that was the best obtainable from that data. Fig. 45 shows a Clementine picture of the Shackleton region, and an overlay of the tested landing trajectory. Several hundred individual maps were created as part of, and independent of, the landing simulation preparation. Figure 46 shows a few of the landmark maps that were developed along with an image from the Clementine mission of the same site. In particular, along the descent trajectory, specific submaps were created corresponding to each of the predicted descent picture times along that trajectory. Within those maps, a realistic texturing and detailing process was used to add fractal terrain variations, craters and rocks in power-law distributions consistent with studies of these distributions on the Moon. Finally, within the maps associated with the terrain for each picture time, landmark

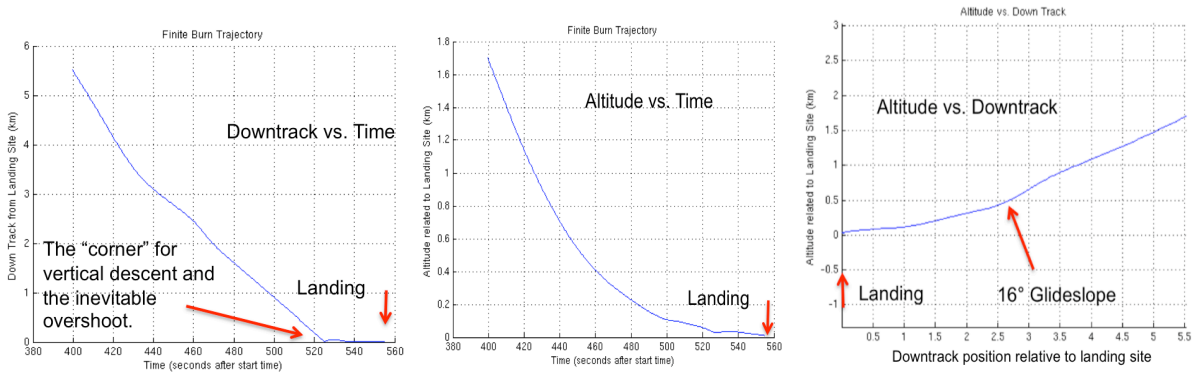

Figure 48: Aspects of the Nominal Altair Apollo-like 16° Glideslope Landing Trajectory: *GN&C stress is indicated by the discontinuities in acceleration, most notably the "corner" directly over the landing site where several 10's of m/s of lander velocity must nulled in a very short period of time*

features were selected. For each of these landmarks, a small terrain map was made for purposes of AutoGNC processing. A series of images created for the simulation is shown in Figure 47; these were taken during the powered descent. It is important to note that separate truth and nominal maps were developed and used. The truth maps were high resolution, matching the best resolution that the simulated camera would obtain. The nominal maps, on the other hand, had the maximum resolution expected to be obtained from a precursor sampling mission (e.g. Lunar Reconnaissance Orbiter, which will obtain 50 m maps globally, and 0.5m maps locally near landing sites), as it is the nominal maps that are used for landmark tracking in AutoGNC.

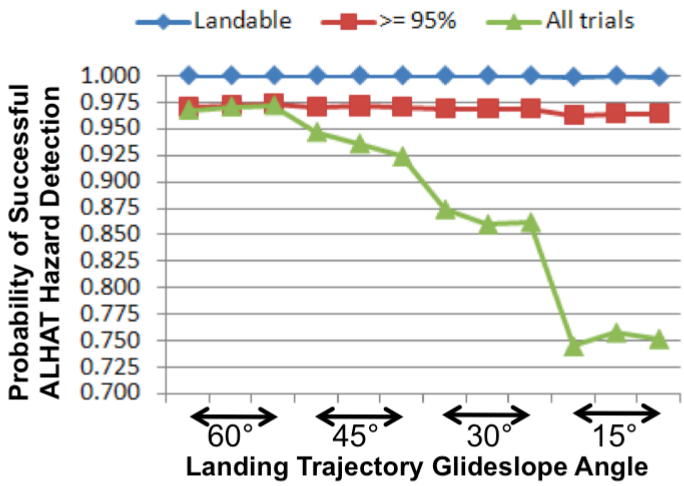

Figure 49. Probability of Successful Hazard Detection as a Function of Landing Site Approach Glideslope (Green Curve): *Using a shallow slantrange LIDAR observation of the landing site impedes the ability of the ALHAT system to successfuly identify hazards and avoid them (Courtesey Andrew Johnson, Alan Strahan, ALHAT Project)*

2. Results from the Lunar Landing Simulation

The discussion and charts that follow describe the ability of the AutoGNC system to ascertain the spacecraft position relative to the terrain throughout the approximately 9 minute powered descent to the lunar surface on the Shackleton rim, and to maintain the desired course with periodic trims to the engine thrust vector.

One of the key aspects of the lunar landing demonstration so far is a completely successful operation of the OBIRON components, successfully acquiring and tracking hundreds of landmarks on the ground track to Shackleton. This approach occurred over a very challenging region of image processing, with deep shadows induced by extremely rugged terrain illuminated by a shallow, grazing incidence sun angle. Because the Clementine images of the southern polar regions were taken in mid-southern winter, the terrain that could be effectively deduced by the stereo-photoclinometry technique used was limited to that seen by Clementine. It should be noted that stereophotoclinometry is far more effective at determining topography from images than conventional stereo-photogrammetry. As a result, through the landing was limited to a very stressful winter landing, it is nonetheless a successfully operating system. Thus this demonstration represents a possible worst-case scenario to test the AutoGNC techniques and software and also represents an important worst case technology feasibility demonstration for the Constellation Program.

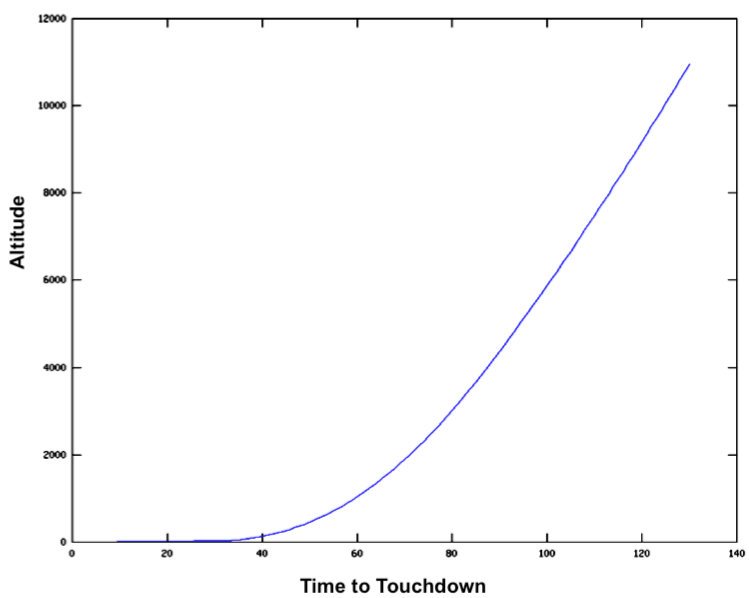

Figure 50. The Last 130 Seconds of a GN&C-, and Especially OpNav-Optimum Trajectory for a Lunar landing; *By making the landing profile less stressful, and therefore more balistically optimum, the landing trajectory becomes much more commodious for OpNav observations as well as for the ALHAT system.*

Performance of the system was very good, as shown in Fig. 3. This plot shows the difference between the targeted trajectory and the delivered trajectory. These errors represent a roll-up of all errors including orbit determination, attitude and trajectory control errors. The errors never exceed 50m (except for a brief inialization transient) and reach a minimum level of under 10m at about 60seconds before landing, and at this time the OBIRON loses lock on the surface, and errors begin to subsequently grow due to integrated IMU errors. The trajectory used in this simulation is a version of the nominal Altair trajectory which is based on Apollo "low-glide-slope" practice. This form of trajectory is a major contributior to the loss of visual lock, due to inability to fix the OpNav system on a single region of landmarks – especially the landing site. As a result, a rapidly streaming set of landmarks is passing through the filter, leaving it susceptible to a difficult region of terrain disrupting the filter's hold on the terrain.

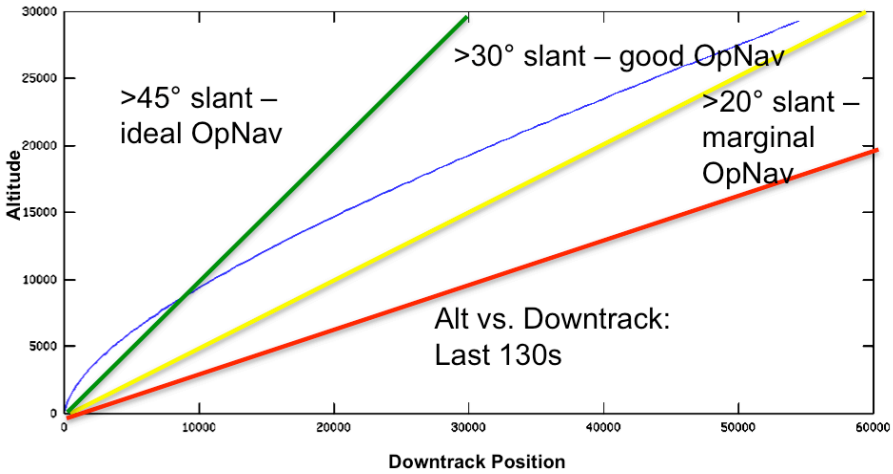

Figure 51. Alternative Altair Landing Trajectory Geometry in the Final 130s of Descent; *By utilizing a more dynamically optimum approach to the landing site, the ability of OpNav data to guide the vehicle is dramatically improved over the 16°Apollo-like glide-slope trajectory*

3. Seeing Where You're Landing and Landing Where You're Seeing

Fig. 48 shows the nominal design of the Altair landing trajectory, displaying the principal characteristic of the trajectory, namely the more than one minute of shallow (16°) glide-slope. This design was adapted from Apollo largely for visibility by the astronauts of the terrain. The Apollo LEM (Lunar Excursion Module) had no electronic cameras aboard, so the critical seeing of the landing site had to be "biological" as opposed to electronic. The visibility of the landing site from Altair is far worse than from Apollo (see Fig. 49), meaning that even with such a shallow glide-slope the landing site will be invisible for many seconds before reaching the terminal vertical descent. This form of trajectory also puts many stresses on the GN&C system. As discussed above, the rapid traverse at low altitude is stressful for the optical navigation system, with multiple different landmarks and "frame-tie" regions passing through the navigator, and no view of the landing site. Dynamically, the long, slowly-descending 30% throttle hover is highly fuel consumptive, the shallow approach is very difficult for the ALHAT landing LIDAR and hazard avoidance system to contend with, and the abrupt accelertation changes, particulary at "tip-up" when the thrust drops from nearly 100% to about 33% and the orientation of Altair changes from pointing largely in the velocity direction to mostly vertical.

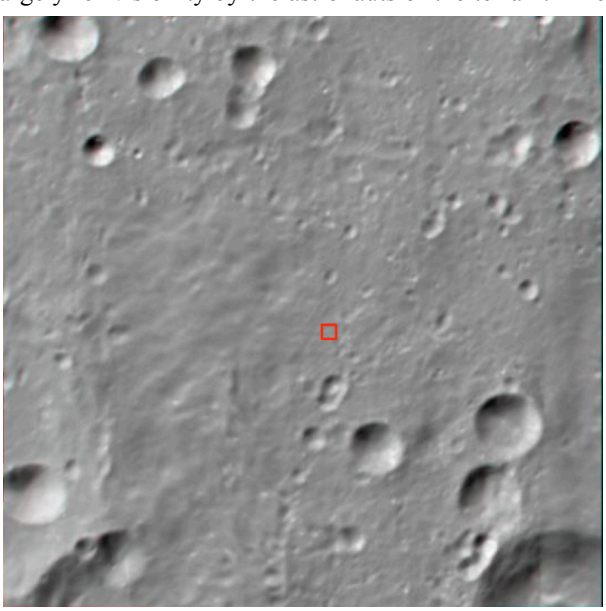

Figure 52. Composite Clementine-Derived Topographic Map at -59.7S, -160.7E; *This map is of approximately 220m resolution, and shows a small region which is shown in high detail in the next figure*

Figure 7: Nearly OpNav-Optimum™ and Ballistically Optimum Trajectory Characteristics

Time Relative to Landing (s)	Throttle Level	Characteristic
0 - 30	~30%	Vertical Descent
30 - 60	Linear throttle down from 100% to ~30%	Rotate engine bell from vertical to 30° from vertical at 1°/s
60 - 90	~100%	Rotate engine bell from 30° to 60° at 1°/s
90 - 405s	~100%	Rotate engine bell from 60° to ~90° at 0.1°/s

If "electronic vision" is adopted as the principal method of visualization of the landing site by the crew, a number of advantages immediately come to the fore. The first, most obvious, and most germane to the current discussion is that the trajectory need no longer be designed for visibility of the landing site out the view-ports. Freed of this constraint, the trajectory can be configured with other conflicting demands optimized, including propellant use, and visibility of the landing site by electronic instruments, principally the ONSS and the ALHAT hazard detection and avoidance system. Figure 50 shows aspects of an example of an alternate Altair landing trajectory, that saves over 1000kg in propellant relative to the previously shown trajectory. This approach to the landing site is much more conducive to OpNav relative to the landing site, as shown in Fig. 51, which shows the relative quality ranges of view angle to an OpNav target when landmarks are being used as data. A summary of the phases of this trajectory is shown in Table 7, where two important attitude and thrust challenges to the landing sequence are eliminated, the discontinuous "pitch up" and the arresting of horizontal velocity above the landing site. Both of these events present substantial levels of jerk, which increased noise in the navigation system due to control errors, and their elimination will make the descent much more accurate and comfortable for both the crew and the GN&C system.

4. Landing in the Aitken Basin

Besides the Shackleton site, Constellation has a list continaing several dozen potential landing sites, including several in the Aitken Basin. As another test of the lunar landing applications for AutoGNC, a landing demonstration is being prepared for a Cx landing site at -60 lat., -200 lon. Figure 52 shows an image of the initial coarse model (200m resolution) based on the Clementine visual survey. In the new simulation the "electronic vision"-tuned trajectory will be utilized. LRO images are being utilized at the time of this writing to develop topography maps using stereophotoclinometry, and Fig. 53. shows one of these LRO images.

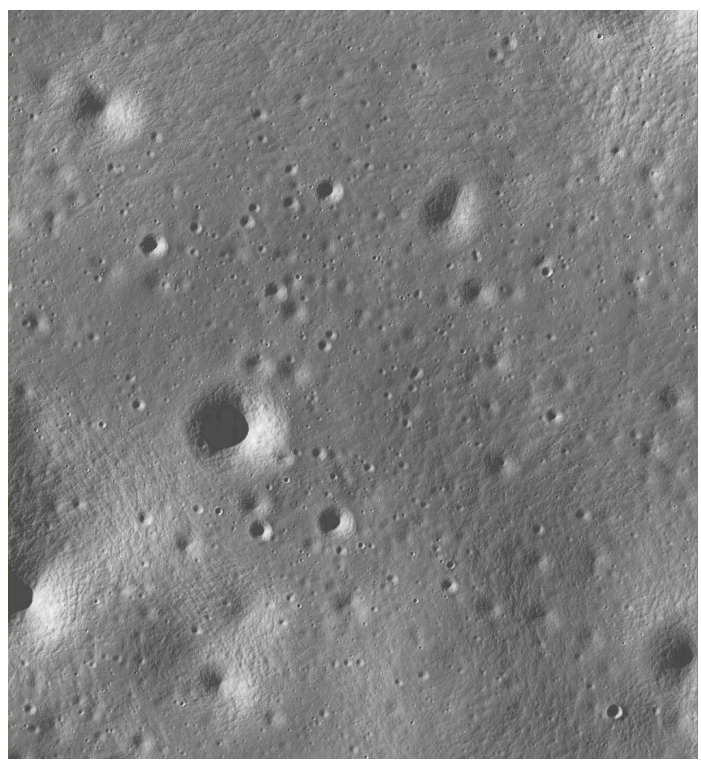

Figure 53: LRO High-Resolution Image at -59.69S, -160.69E, (M103189611R); *This image has 0.5m resoluiton and is roughly oriented in the central square of the previous figure*

VII. First NASA Use of Landmark-based Optical Navigation at the Moon: LCROSS

The recent lunar-surface-impacting-mission LCROSS (Lunar CRater Observation and Sensing Satellite) utilized OpNav to reconstruct the landing site of the "Nurse Ship" that guided the Centaur upper stage to its impact site, and then followed the impactor closely along the same trajectory. Shown in Figs. 54 and 55 are two examples of the OpNav imaging used for LCROSS, showing the sets of SPC-generated landmarks and regions in the OpNav frame against which they were matched, these are the first and the last pictures taken and used in this orbit determination sequence. This data was not processed onboard the spacecraft, but was performed as a post-flight reconstruction. The maps used for the optical navigation were based on Clementine, Apollo, LO, and LRO imaging and developed by Robert Gaskell (Ref. 14). The horizontal resolution of the map varies from 100 to 200m and is about 150m in the Cabeus region where the LCROSS vehicles impacted. However, the ensemble registration accuracy of the maps, or "frame-tie accuracy" is belived to be at the 50m level, in other words, groups of landmarks in different portions of the maps might be shifted as an ensemble by as much as 50m from another group. This analysis processed 54 pictures taken during the final hour of descent of the LCROSS nurse ship, from an altitude of 5500 to 175km. The LCROSS VIS camera had a 30° x 23° field of view, which is sufficiently wide to extract altitude information. The VIS imager had a rectangular detector of 720 x 480 pixels, for an approximate pixel anglular size of 0.88mrad.

The 54 LCROSS pictures were processed using the OBIRON system and several dozen landmarks extracted from the SPC-generated topography map, and combined with the powerful radiometric data. The result of this solution was a shift of 50m as compared to the radio-only solution. The magnitude of this 50m shift is consistent with the believed frame-tie accuracy of the composite map. Figure 56 shows a view of the nominally planned impact target, and the final combined solution, showing agreement within 200m. To put this 50m shift in final solution in further context of the LCROSS optical data; the final picture had a resolution of 150m/pixel, the post-fit RMS residuals are about 75m (~1/2 pixel) across all 54 pictures, indicating that the 50m offset is of bare statistical significance. The ensemble residuals are presented in Fig. 57, where the coordinates are the VIS picture frame. Individual landmarks of the two-dozen or so used can be traced in the sequence of pictures as they are approached by the vehicle. A gross residual bias between the optical and radio cannot be perceived, although individual landmarks show small biases relative to others – indications of landmark location errors as a result of presumed frame-tie errors, or possibly errors in the topography map itself, due to altitude errors arising from the fact that the global topographic map is not yet complete, and therefore is subject to some possible residual discontinuities.

VIII. Photometry Challenges for Optical Navigation at the Moon

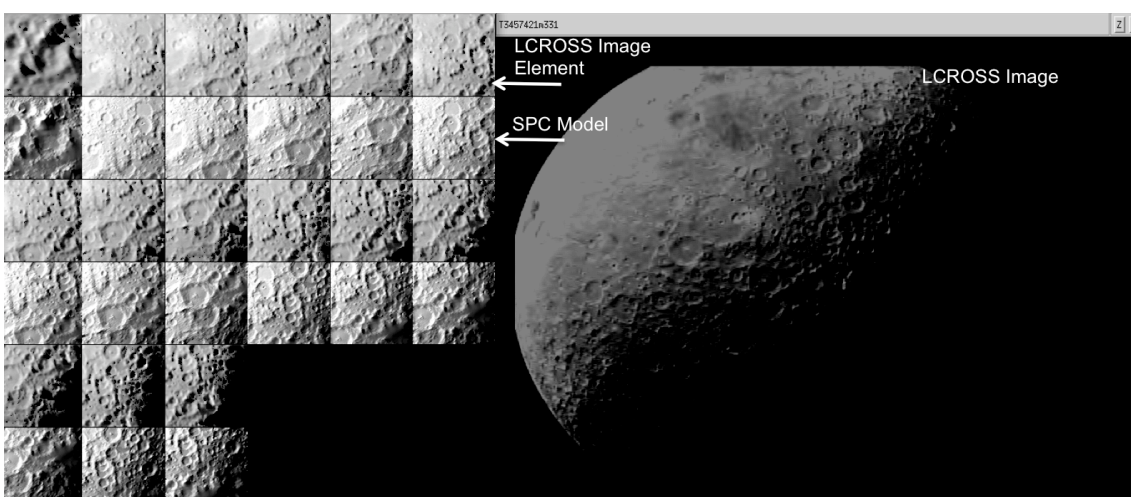

Figure 54. First Image in LCROSS Terminal Descent OpNav Sequence; *On the left are the landmarks used in the processing of this picture, with alternating rows of pictrue-elements and Stereophotoclinometrically (SPC) generated and rendered models*

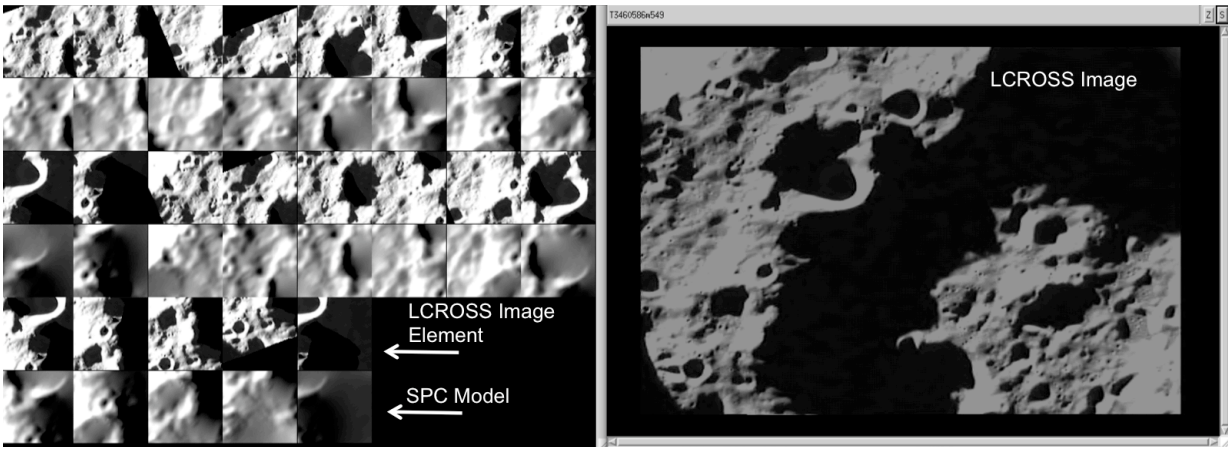

Figure 55. Final Image in LCROSS Terminal Descent OpNav Sequence; *On the left are landmark images compared with modeled and rendered versions, showing the limitations of the resolution models*

Optical navigation at the Moon will encounter two different lighting condition stresses: too bright and too dim. But except for missions that call for night-time landing in mid-latitudes, these stresses can be overcome with good cameras such as those envisioned as components of the ONSS. The former problem is chiefly a challenge of Cis-Lunar navigation where approach or departing images are taken of the limb of the moon with the desire to also capture stars in order to eliminate pointing errors. Despite an average albedo of about 0.2, the sub-solar limb of the Moon is very bright, and even the wide dynamic range of the ONC will necessitate bright stars of 5-6th M when imaging the sub-solar limb. Such stars will occur typically at a spatial density of less then one per ONC frame on average – thereby reducing OpNav opportunities. Fortunately, there are several ameliorations to this challenge. With a target as large as the Moon (even in mid Cis-Lunar transit the Moon subtends nearly 700 pixels) it is seldom the case that the only pointing choice is the sub-solar limb. The limb of a near-0-phase Moon or the cusps of a crescent Moon present image opportunities several magnitudes dimmer than the sub-solar illumination point, which will allow imaging of correspondingly dimmers stars with the limb, thereby guaranteeing star-image opportunites. Even when limited to sub-solar limb observations, such as close to LOI, the rapidly changing geometry will provide for change of viewing angle, causing other (and eventually brighter) stars to enter the field of view. A less desireable but still reasonably effective method to overcome dynamic range limitations is to take in quick succession images of (relatively) long and short exposure. For typical Moon OpNavs where stars are desired against the sub-solar limb, these might be in the range of 0.075 and 0.015 seconds respectively. In the longer exposure the lunar suface will be over-exposed and featureless, but stars as dim as 9th magnitude will be visible in the ONC. In the short exposure, the short exposure, landmarks on the Moon will be trackable, but only stars as bright as 5^{th-} m will be visible. This last method is a less favorable method as pointing knowledge loss between the pictures can represent a substantial reduction in accuracy. Finally, when nearly LOI and in orbit, the WAC is probably going to be the instrument of choice, in which case, obtaining stellar images is unnecessary as the wide field of at least 20° provides a means of differentiating between position and pointing error.

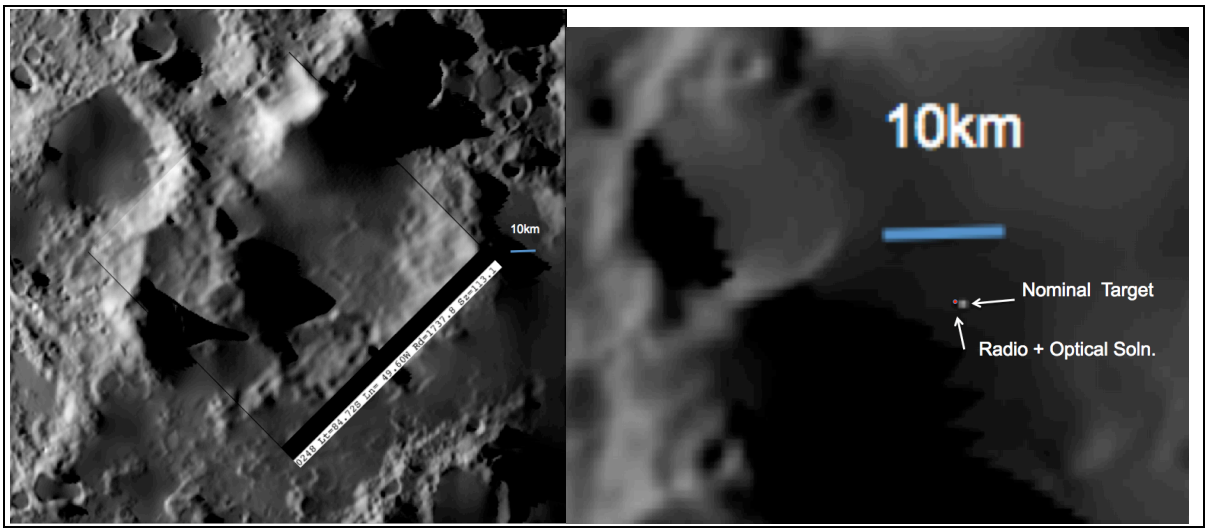

Figure 56. LCROSS Local Terrain Map of Target, and Combined Radio Plus Optical Navigation Solution: *On the left is an LCROSS VIS image, overlayed in the Cabeus target area with the Clementine/LO/Apollo-generated topographic map. On the right is a close-up of the nominal target and the combined solution, showing a less than 250m offset from the nominal. The radio-only and optical solutions differed by approximately 50m.*

The challenge of insufficient light is greater, but not by any means insurmountable. Already discussed is the plausibility of a wintertime landing in the Shackleton region using OpNav as demonstrated by the AutoGNC-based simulation? In a more illuminated season, with the sun 1° above the horizon of the South polar region, the surface illumination will be (to put it in a familiar frame) 0.5% of a tropical, cloudless Earth-day at noon, or 570 times brighter than full-Moon illumination on the Earth. This is sufficient light for the ONC to take very good OpNav frames with 0.1sec exposures. But OpNav is still possible where the Sun doesn't shine – ever. Shackleton crater is a region of great interest for well-known reasons, and a place where eventual landings will occur – if not perhaps by the Altair project. A common phenomenon well known to hikers is bright illumination by nearby peaks even when the local terrain is in shadow. On the Moon this will present important opportunities for GN&C operations. The bottom of the large Shackleton crater is surrounded by tall rugged peaks of the rim (as are most of the craters in the region). Again, assuming a 0.2 albedo lunar surface, a 10km^2 illuminated region of rim 25km distant, and a reflection/incidence loss of 60%, the net surface illumination in the bottom of the crater due to alpenglow is

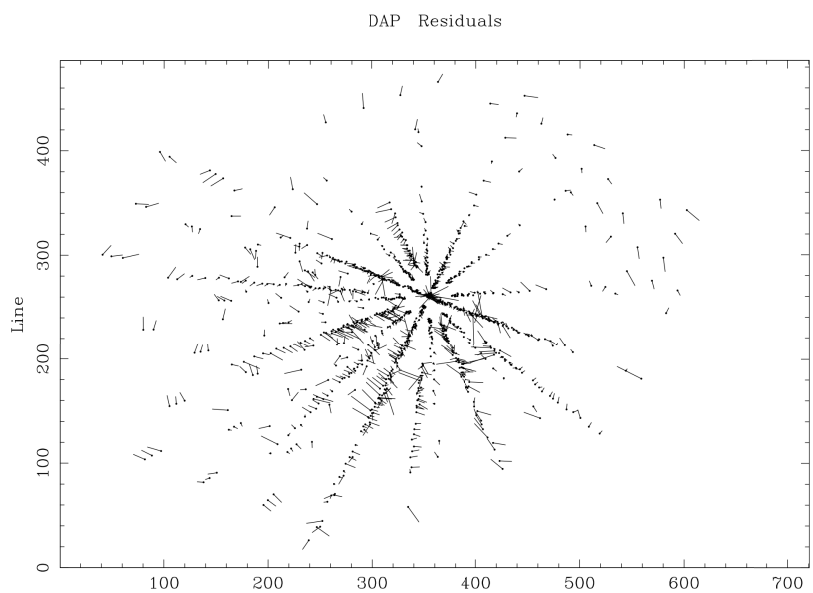

Figure 57. LCROSS OpNav Residuals: *Pixel and Line residuals of 54 LCROSS terminal descent OpNav images shown as a composite in VIS imager coordinates – the residuals are scaled by x10*

~3e-4 of direct sunlight, or 30x surface-of-Earth "Full Moon" illumination. Such a level of brightness (based on conservative assumptions) is adequate for OpNav producing a SNR of at least 10 in a 0.5 second exposure in the ONC or equivalently sensitive camera. As an example of the usability of alpine-glow, Fig 58. shows an image from the NEAR mission, showing a sun-occulted region of the surface of Eros, illuminated only by alpine-glow of nearby prominences. In this case, the brightness of this reflected-light region is 300 times the Earth-surface luminance due to a full Moon. Terrain is easily visible in this light, ten times brighter than the calculation above, probably attributable to the nearness (less than 1 km) to the lit terrain. This case has direct relevance to the lunar polar regions, as many such permanently shadowed regions exist in nearby proximity to tall, steep topography. Thus OpNav can effectively navigate a vehicle such as Altair even to a place "where the sun don't shine."

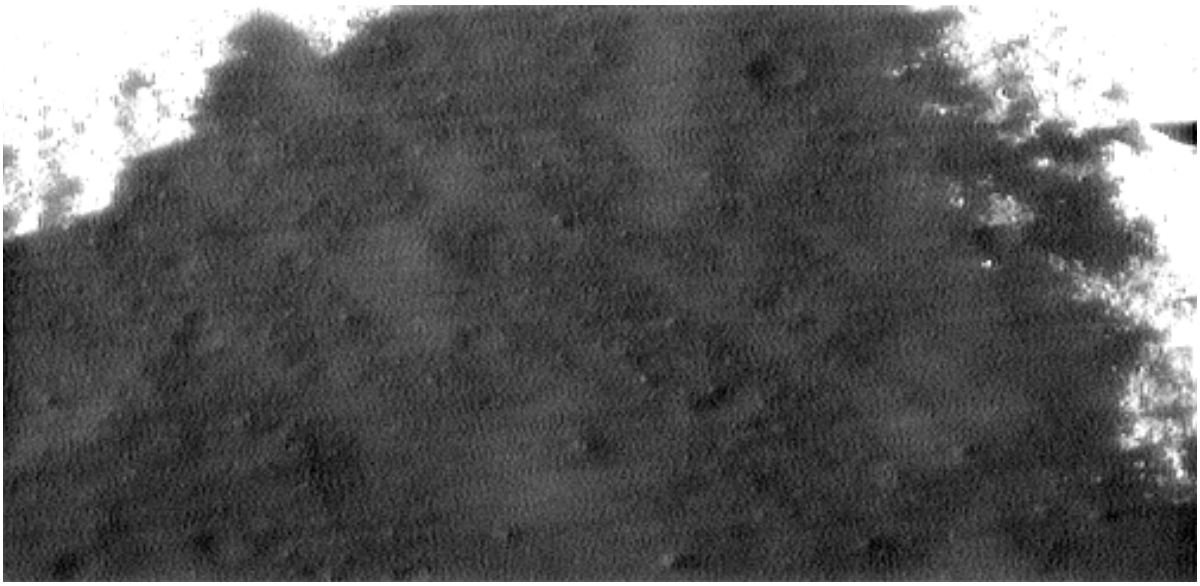

Figure 58: Alpine Glow as Seen in a NEAR Mission Image of Eros; *the terrain shown in the shadow is illuminated only by reflected light from other portions of the Eros asteroid surface, and is 300x brighter than Earth's surface on a "full Moon" night*

IX. Conclusions

Optical navigation is a powerful technique for deep space navigation with a heritage of over forty years. The Altair missions, and virtually all lunar landing and virtually all exo-Earth missions that require proximity operations, precision landing and/or on-orbit rendezvous will require some form of OpNav. In addition, the Constellation program wisely required that both the Orion and Altair projects provide for safe return of the crew when the radiometric navigation link is compromised; and this requires the use of OpNav by both vehicles. For Altair, OpNav means use of landmarks as navigational targets, necessitating the generation of and utilization onboard of accurate topography maps, a clear departure from methods used by Apollo.

Besides providing necessary backup navigation capability, OpNav will dramatically improve the Altair performance during orbit and landing, and provide important back-up functions for RPOD. Though still limited by non-gravitational acceleration issues, raw data power of OpNav landmark tracking is as good as 10m from orbit. The capability of OpNav to provide precise landing has been shown and in the regime of 10m or better.

Comprehensive simulations of software systems necessary to perfom optical navigation and trajectory control, indeed, fully comprehensive GN&C, have been demonstrated for a related mission scenario, rendezvous and contact with a small asteroid, with great success. This same system has been used to emulate the powered descent and landing on the Shackleton rim, in Winter, thereby demonstrating the robustness of the OpNav techniques for the Alatiar application.

The LCROSS mission has utilized OpNav for post-flight reconstruction, becoming the first NASA mission to use a network of lunar landmarks for optical navigation. The agreement between the radiometric data and the optical was within the expected accuracy of the maps (about 50m) and bodes well for safe and precise landing for Altiar, for which missions the LRO mission imaging data will have enabled the development of maps of 5m accuracy or better in the Constellation landing regions. Finally, photometric challenges will be presented to the use of lunar OpNav, but with good cameras, judicious picture planning, and sensible mission design these can always be overcome, allowing optical navigation to be a key component of the Altair GN&C suite of capabilities throughout the mission.

Acknowledgement

The research described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. Acknowledgement is due J. Balaram and C. Lim of JPL for extensive and ongoing assistance in configuring the DSENDS system for the wide-ranging integration and demonstration efforts described in this paper.

References

[1] Morabito, David D., Riedel, J.E., Han, D., "The Effects of Earthward Directed Coronal Mass Ejections on Near-Earth S-band Signal Links", submitted to Radio Science, August, 2010

[2] Ely, Todd, Heyne, Martin, Riedel, Joseph E., "Altair Navigation during Trans-Lunar Cruise, Lunar Orbit, Descent and Landing", AIAA GN&C Specialty Conference, Toronto, August 2010

[3] Riedel, R., Bhaskaran, S., Eldred, D., Gaskell, R., Grasso, C., Kennedy, B., Kubitschek, D., Mastrodemos, N., Synnott, S., Vaughan, A., Werner, R., "AutoNav Mark3: Engineering the Next Generation of Autonomous Onboard Navigation and Guidance," AIAA Paper AIAA 2006-6708, AIAA GN&C Conference, 21-24 August 2006, Keystone, CO.

[4] Brady, T., Tillier, C., Brown, R., Jimenez, A., Kourepenis, A., "The Inertial Stellar Compass: A New Direction in Spacecraft Attitude Determination," 16th Annual AIAA, USU Conference on Small Satellites, Logan, Utah, August 12-15,2002.

[5] http://www.moog.com/Media/1/Biax%20Intro.pdf

[6] Riedel, J.E., Bhaskaran, S., Synnott, S.P., Desai, S.D., Bollman, W.E., Dumont, P.J., Halsell, C.A., Han, D., Kennedy, B.M., Null, G.W., Owen Jr., W.M., Werner, R.A., and Williams, B.G. "Navigation for the New Millenium: Autonomous Navigation for Deep Space-1," Proceedings of the 12th International Symposium on Flight Dynamics, Darmstadt, Germany, Jun. 1997.

[7] Emil Schluser, Apollo Project Navigation Team, NASA, Private Communication

[8] R. W. Gaskell And N. Mastrodemos, **"**Lunar Topography From Stereophotoclinometry,**"** Paper #1152, XXXIX LPSC Conference, Houston, TX, March 2008

[9] Gaskell, Robert A., "Landmark Navigation and Target Characterization in a Simulated Itokawa Encounter," 2005 AAS/AIAA Astrodynamics Specialist Conference, Lake Tahoe, California.

[10] Gaskell, R, O. Barnouin-Jha, D. Scheeres, T. Mukai, N. Hirata, S. Abe, J. Saito, M. Ishiguro, T. Kubota, T. Hashimoto, J. Kawaguchi, M. Yoshikawa, K. Shirakawa, T. Kominato, Landmark Navigation Studies and Target Characterization in the Hayabusa Encounter with Itokawa, AIAA paper 2006_6660, AAS/AIAA Astrodynamics Specialists Conf., Keystone, CO, Aug 2006

[11] Bhaskaran, S., J. E. Riedel, B. Kennedy, T. C. Wang, "Navigation of the Deep Space 1 Spacecraft at Borrelly," AIAA paper 2002-4815, AIAA/AAS Astrodynamics Specialist Conference and Exhibit, Monterey, CA, August5-8, 2002.

[12] Kubitschek, D., Mastrodemos, N., Werner, R., Kennedy, B., Synnott, S., Null, G, Bhaskaran, S., Riedel, J. , Vaughan, A., "Deep Impact Autonomous Navigation: The Trials of Targeting the Unknown," AAS 06-081, 29th Annual AAS Guidance and Control Conference, Breckenridge, Co., Feb. 4-8, 2006.

[13] GSSR Reference

[14] Robert Gaskell, Planetary Sciences Institute, Private Communicaiton

[15] R. W. Gaskell And N. Mastrodemos, **"**Lunar Topography From Stereophotoclinometry,**"** Paper #1152, XXXIX LPSC Conference, Houston, TX, March 2008

Copyright © 2010 California Institute of Technology. Government sponsorship acknowledged.