(Preprint) AAS 16-113

ORION BURN MANAGEMENT, NOMINAL AND RESPONSE TO FAILURES

Ryan Odegard*, John L. Goodman†, Charles P. Barrett‡, Kara Pohlkamp§, and Shane Robinson¶

An approach for managing Orion on-orbit burn execution is described for nominal and failure response scenarios. The burn management strategy for Orion takes into account per-burn variations in targeting, timing, and execution; crew and ground operator intervention and overrides; defined burn failure triggers and responses; and corresponding on-board software sequencing functionality. Burn-to-burn variations are managed through the identification of specific parameters that may be updated for each progressive burn. Failure triggers and automatic responses during the burn timeframe are defined to provide safety for the crew in the case of vehicle failures, along with override capabilities to ensure operational control of the vehicle. On-board sequencing software provides the timeline coordination for performing the required activities related to targeting, burn execution, and responding to burn failures.

INTRODUCTION

The trend of increasing automation and autonomy of NASA's crewed spacecraft continues with Orion. The exploration mission flight profiles, communications time delays, and a requirement for vehicle return in the event of loss of communications drive a need for a greater level of on-board automation (no human interaction with the flight computer or spacecraft) and autonomy (no Mission Control interaction with the spacecraft or crew). Management of on-board systems for burns is a major factor in the design of the Orion automation and sequencing. The design must also accommodate crew or Mission Control override of automated burn sequencing.

Figure 1 shows the Exploration Mission 1 (EM-1) flight profile for Orion. [1] For the uncrewed EM-1 mission the Orion on-board targeted burns include four Outbound Trajectory Correction (OTC) burns, the Outbound Powered Flyby (OPF) burn at the Moon, two more OTC burns, the Distant Retrograde Orbit (DRO) Insertion (DRI) burn, three Orbit Maintenance (OM) burns while in the DRO, the DRO Departure (DRD) burn, three Return Trajectory Correction (RTC) burns, the Return Powered Flyby (RPF) burn at the Moon, and three more RTC burns executed before the vehicle reaches Entry Interface (EI).

*Senior Member Technical Staff, Spacecraft GN&C and Mission Operations Group, Charles Stark Draper Laboratory, 17629 El Camino Real, Houston, TX, 77578.

†Specialist Engineer, Odyssey Space Research, LLC., 1120 NASA Parkway, Suite 505 Houston, TX 77058.

‡Flight Dynamics Officer, Flight Dynamics Division, NASA Johnson Space Center, 2101 NASA Parkway, Houston TX 77058.

§GNC Operations, Flight Dynamics Division, NASA Johnson Space Center, 2101 NASA Parkway, Houston TX 77058.

¶Aerospace Engineer, Aeroscience and Flight Mechanics Division, NASA Johnson Space Center, 2101 NASA Parkway, Houston TX 77058.

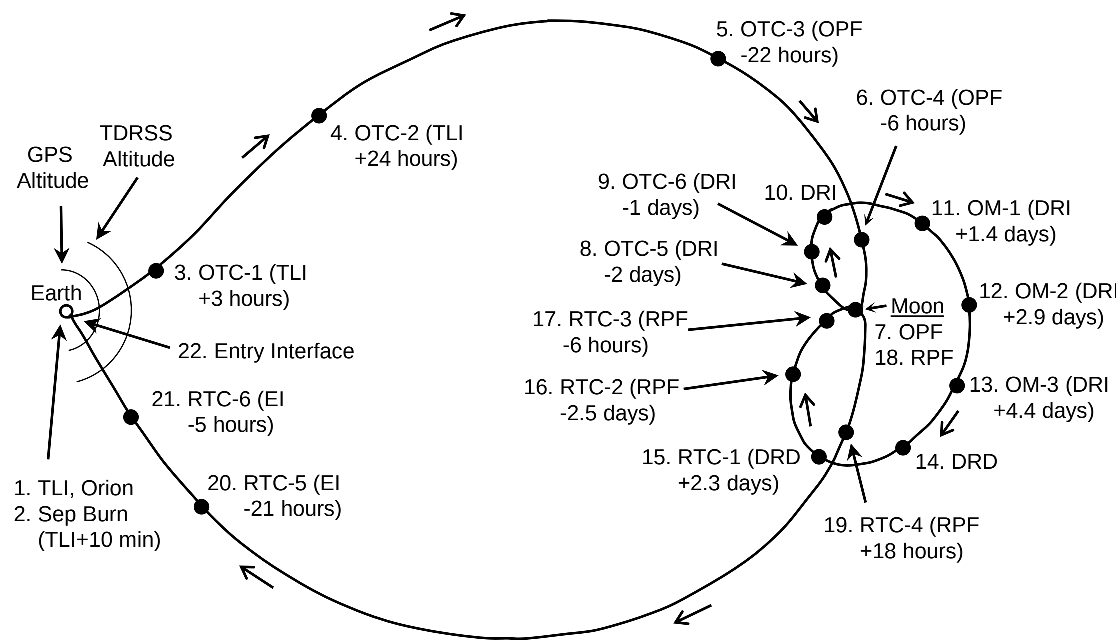

Figure 1. EM-1 mission profile.

The trajectory profile for the crewed EM-2 mission is a high lunar orbit $[100\,\mathrm{km} \times 10,000\,\mathrm{km}]$ $(328,084\,\mathrm{ft} \times 32,808,400\,\mathrm{ft})$. The on-board targeted burns include the Outbound Trajectory Adjust (OTA) burn, three OTC burns, Lunar Orbit Insertion (LOI) burn, OM burns while in the high lunar orbit, the Trans-Earth Injection (TEI) burn, and three RTC burns executed before EI.

The burns for both missions after Trans Lunar Injection (TLI) are targeted on-board the Orion vehicle with the Two Level Targeter (TLT). [2] Burn targeting by the Mission Control is available for comparison purposes and as a backup, as well as trajectory re-optimization using the Copernicus tool, if it is required. [3] With the exception of spacecraft separation burns, all Orion-controlled, on-orbit SM burns utilize the same burn management approach. Although there are differences between EM-1 and EM-2 (uncrewed and crewed, DRO versus high lunar orbit), the burn management approach is the similar for both missions.

On-orbit aborts can be performed after TLI to return the vehicle safely to Earth. For the crewed EM-2 mission the vehicle will be able to return the crew safely to Earth within 120 hours at any point during the flight. All abort burns, for both direct and lunar fly-by aborts, can be targeted on-board using the TLT.

Translational burns fall into one of three criticality classifications. A *critical* burn must be successfully executed at the upcoming burn opportunity to ensure the safety of the vehicle and crew. A *mandatory* burn must be successfully executed at the upcoming burn opportunity to ensure mission success, but failure to execute the burn does not threaten safety of flight. Execution of a *non-critical* burn can be delayed to a future opportunity without compromising safety or mission success.

The propulsion system used for the on-orbit burns resides on the European Service Module (SM). Three different types of engines on the SM may be used for translational burns. All three use hypergolic propellants. The first is one Space Shuttle-heritage Aerojet Rocketdyne Orbital Maneuvering

System Engine (OMS-E) of 6,000 lbf thrust with Thrust Vector Control (TVC) for steering during the burn. Second, eight Aerojet Rocketdyne R-4D-11 auxiliary engines provide 110 lbf thrust each, and are called AUX (+X) engines. They provide a back-up in the event of an OMS-E failure. The third propulsion system consists of 24 Automated Transfer Vehicle Reaction Control System (RCS) thrusters rated at 50 lbf of thrust each. [4] The RCS thruster layout allows translational burns to be performed without requiring a maneuver to the burn attitude.

This paper begins with an overview of the on-orbit concept of operations involving targeting and burn execution. Next is an overview of the mission sequencing architecture that coordinates the Guidance, Navigation, and Control (GN&C) activities. Third, the burn plan management approach is described for handling burn-to-burn variations, alternate burn plans, targeting and guidance, and control of automated sequencing. Finally, an overview of burn sequencing and automated responses to burn failures is presented.

ON-ORBIT CONCEPT OF OPERATIONS

Trajectory translational burns that require lunar gravity assist (OPF, RPF) most likely will be performed while line-of-sight communications between the spacecraft and Mission Control is not available. Since the first mission is uncrewed, burn execution must be automated to complete mission objectives. In addition, in the event that Orion is already on its way back to the Earth and communication with Orion is lost, program requirements dictate that Orion have the ability to autonomously perform the burns required for it to meet its EI conditions. These two scenarios make it necessary for Orion to be able to target and execute burns automatically. The following operations concept has been formulated to address these scenarios.

While there is communications between Orion and Mission Control, the flight control team in the Mission Control Center at the NASA's Johnson Space Center keeps Orion's navigation state vector (time, position, velocity) up to date using Deep Space Network tracking. Mission Control provides Orion with an optimal trajectory to follow. Before each burn, Mission Control can provide updated burn targeting and state vector data to Orion. At the requested time, the on-board TLT computes burn guidance targets for use by Orion on-board guidance, with monitoring by Mission Control. Mission Control also computes burn guidance targets for comparison with the on-board targets and for possible use during the burn instead of the TLT solution. Once acceptable targets are on-board, Orion prepares for the upcoming burn. Mission Control has the ability to intervene in the automated activities as necessary.

Burn preparations include events such as powering up systems required to perform the burn, aligning the Inertial Measurement Units (IMUs), maneuvering to burn attitude, positioning solar arrays, and configuring the propulsion systems. Burn execution nominally begins at the time of ignition. If the remaining velocity residuals after engine cut-off are above acceptable limits, Orion trims the burn to within acceptable tolerances using the SM RCS thrusters. Once the burn is complete, Orion performs post-burn checks and system reconfiguration. As the next burn in the timeline approaches, this process is repeated.

The burn concept of operations must also account for times when Orion has lost communications with Mission Control. This loss of communications failure prevents Mission Control from updating the on-board navigation state vector. To accommodate this, Orion is equipped with an optical navigation system that keeps its navigation state vector within accuracy limits acceptable to achieve the required EI conditions for return to the Earth. [5] In this scenario burn targeting and burn execution operations are performed autonomously.

MISSION SEQUENCING ARCHITECTURE

The responsibility for on-board sequencing is distributed among the Timeline Management (TMG) software and the vehicle subsystems. The TMG software is responsible for the overall mission timeline and coordination of the Orion subsystems. This knowledge of the overall timeline is based on a sequence of mission phases and mission segments. Within the GN&C subsystem, a series of activities are sequenced in order to meet the objectives of the mission segment. [6]

Figure 2 shows a portion of an example mission timeline of an Orion flight, and is represented in a hierarchy of Phases, Segments, Activities, and Modes (PSAM). The mission phases represent the major operational portions of the timeline, such as ascent, lunar transit, and entry. Each mission phase contains mission segments, which correspond to the major vehicle events that will occur during that mission phase. During flight, the TMG software communicates the current phase and segment to the Orion subsystems to coordinate the vehicle events. For example, during the orbit phase, the subsystem functionality will change in response to changing segments from "coast" to "burn." In some cases the subsystem response to the current phase and segment will require internal sequencing, as is the case for GN&C. Figure 2 depicts a Burn segment containing multiple GN&C activities.

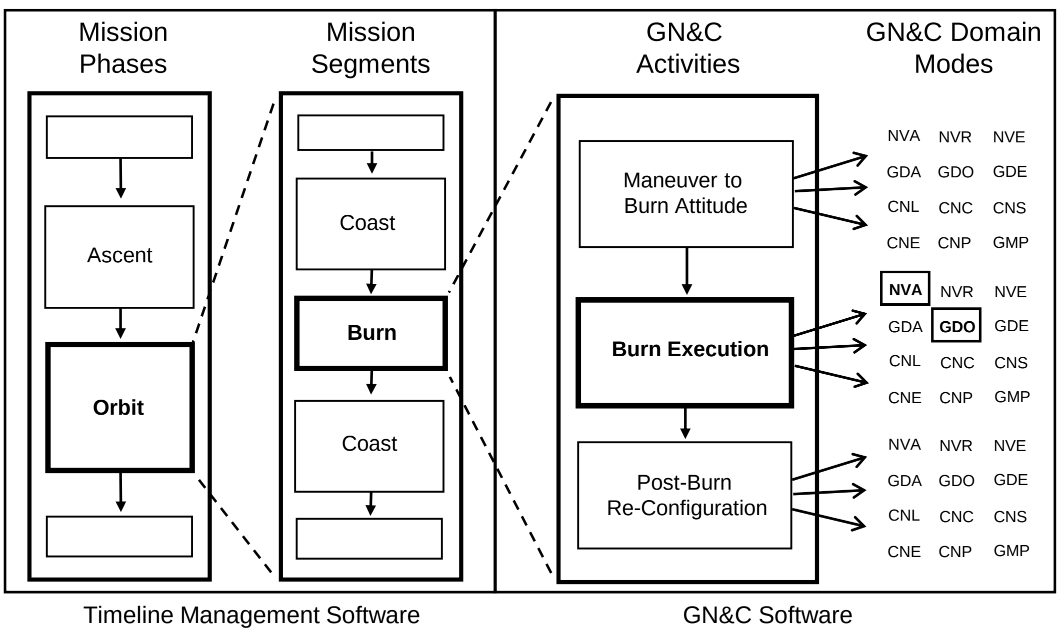

Figure 2. Phases, Segments, Activities, and Modes (PSAM) sequencing hierarchy.

To support operational flexibility and off-nominal sequencing capabilities, GN&C sequencing of activities includes branching functionality. While Figure 2 represents a linear sequence of activities that execute in GN&C within a mission segment, in some cases non-linear sequences are needed. The design for branching activities consists of transitions from an originating activity to one or more destination activities. This is particularly relevant for on-orbit burn management in two ways. The first is to support sequencing off-ramps to support failure or contingency scenarios, and the second is to provide Mission Control and the crew with control over what automated activities can execute by precluding or enabling certain pathways through the sequence. An example of the first is the ability to stop a burn before it has finished nominally, which would have a different resulting sequence than if it had completed as planned. An example of the second is if Mission Control needs to skip planned activities such as IMU alignment maneuvers or TVC gimbal checks, or to dictate the engine(s) that will be used for a particular burn.

BURN MANAGEMENT

The concept of a "burn plan" provides the method for parameterizing the on-board sequencing constructs (i.e., TMG segments and GN&C activities) in order to sequence through various burns during the mission. Figure 3 shows how the burn plan parameters are issued for each instance of a burn that is executed throughout the mission. Since the execution steps for performing a burn are generally repeatable, a subset of items that vary from burn to burn comprise the parameters of the burn plan. This allows the sequences to be developed and verified prior to flight while providing flexibility to parameterize each burn during a mission as needed. In addition, the total set of data describing the on-orbit segments and activities is reduced by being able to reuse a single burn segment; the parameters from the burn plan contain the specific values to be used for each individual burn.

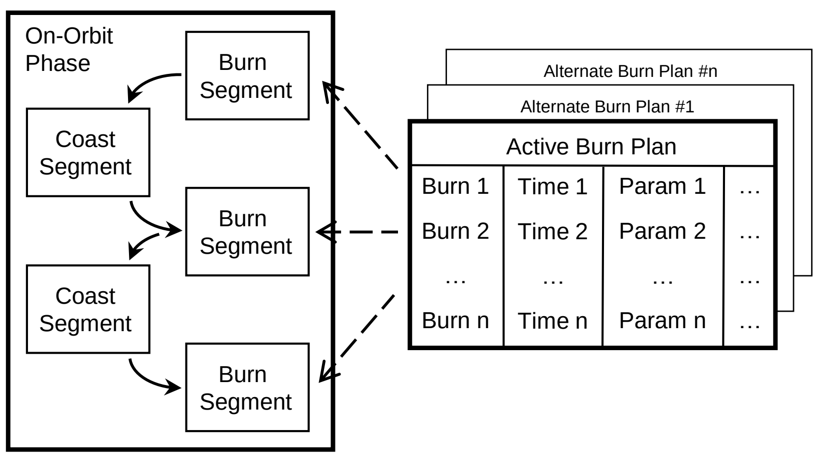

Figure 3. Mission sequencing with burn plan parameterization.

The burn plan contains the information needed to target and execute the specified burns for the mission. The types of parameters in the burn plan include target inputs, which engine to use, the guidance option during the burn, attitude maneuver information, solar array positions and timing, automated sequencing control, and others. The burn plan can be updated by Mission Control and uplinked to Orion if mission planning and systems performance evaluation necessitates a change to any of the currently planned burns.

The burn plan currently being used by the vehicle is called the active burn plan. In addition to the active burn plan, other burn plans covering other scenarios are on-board and maintained by Mission Control. Before launch, the planned active burn plan and burn plans covering possible alternate mission scenarios that could occur on that flight day and the next are uploaded on-board. Burn plans for alternate scenarios are uplinked before major flight activities, such as burns. This permits the vehicle to have an up-to-date burn plan at all times in the event of a loss of communications.

The parameter data that configure the flight software for executing a burn falls into three categories: 1) mission constant, 2) plan specific , and 3) burn specific. The mission constant data include parameters such as physical constants, ephemeris data, threshold values, maximum number of algorithm iterations, etc. It is possible to update the constant parameters in flight, but this would only occur for rare, unanticipated contingency scenarios. The plan-specific parameters are applicable to a burn plan as a whole, and include information such as the trajectory constraints (e.g. minimum flyby altitude, EI target line, maximum ΔV) used by the TLT. The burn-specific parameters are unique to a burn. Examples of these include burn ignition time, guidance option and targets, and associated patch point data for the TLT. These parameters vary from burn to burn in the burn plan.

The burn plan provides the mechanism for managing burns throughout the mission in four main areas. The first is the ability to manage individual burns as they vary throughout the mission and need to be updated as the mission progresses, including adding unscheduled burns or removing burns from the mission timeline. The second is the management of alternate burn plans, in the case where aborts or loss of communications require executing a different series of burns than originally planned. The third area includes the specification of targeting and guidance configurations to be executed for each burn. The fourth area is the ability for crew and Mission Control to control certain aspects of the sequencing throughout the burn timeframe.

Burn-To-Burn Variations

The burn plan provides data describing all of the burns needed to complete the mission. Burns, however, differ depending on their purpose, expected size, importance to the mission (criticality), accuracy requirements and other factors. The burn plan allows Mission Control to specify how the Orion flight software is to prepare for and execute each burn.

Large burns require different engine selection than smaller burns and are typically executed using the OMS engine. However, due to system testing or flight test objectives, Mission Control could elect to use the AUX (+X) engines for the burn instead. Smaller burns may be executed using the SM RCS thrusters in the current attitude. This avoids a maneuver to burn attitude and therefore reduces trajectory perturbations. These variations will be specified in the burn plan for each burn.

Critical burns must be executed to ensure safe return of the crew, and mandatory burns must be executed for mission objectives to be accomplished. Execution of non-critical burns can be delayed if issues arise. To accommodate this, the burn plan includes burn criticality to ensure the important burns are executed. The permissible burn cut-off ΔV residuals will also vary from burn to burn. Some burns have large permissible ΔV residuals while others have smaller permissible residuals. Use of the same ΔV residual limits for each burn could waste propellant through unnecessary burn trims with the SM RCS thrusters.

The variations in burns or the real-time discovery of other constraints such as collision avoidance may result in the need to add or remove burns from the current mission timeline. This can be accomplished via the burn plan by inserting or removing burns from the plan, assuming the update is performed at a time when there is no conflict with current burn activities. The mission sequencing architecture would adapt to the updated plan and execute the appropriate burns.

Alternate Burn Plans

Alternate burn plans can include burn fail cases, alternate loss-of-communications trajectories, post TLI abort trajectories, minimum time to return, minimum propellant to return, deorbit from low Earth orbit, one-off-burns such as collision avoidance, and trajectories flown in response to certain types of failures, such as propulsion system failures.

For example, before the TLI burn, the active burn plan will assume that TLI execution will be successful. If TLI fails, however, the mission plan as defined may no longer be achievable. To account for this, alternate burn plans cover these scenarios. Once the failure is detected, Mission Control will then select the appropriate alternate plan for Orion to execute.

Another scenario that may cause a change to the mission plan is loss of communications. System failures on-board may dictate the need to perform different burns than nominally planned, including protections setup in advance by Mission Control to avoid Orion impacting the Earth or the Moon. Without communication with Mission Control, the EM-1 on-board flight software must automatically select and execute a loss-of-communications burn plan.

Targeting and Guidance

The burn plan supplies the parameters required to configure and operate the on-orbit burn targeting and guidance software. In addition, the burn plan contains data to be used downstream if any element of the targeting and guidance software fails. For non-critical burns a failure of targeting and guidance cancels the burn. Critical burns must be executed even if a particular element of the targeting and guidance system fails. Figure 4 shows the data flow from the burn plan to the targeting and the Orbital Guidance (OrbGuid) algorithms. [7] The TLT simultaneously targets all of the burns yet to be executed in the active plan. OrbGuid produces the active guidance (i.e. closed-loop steering) commands for the current burn.

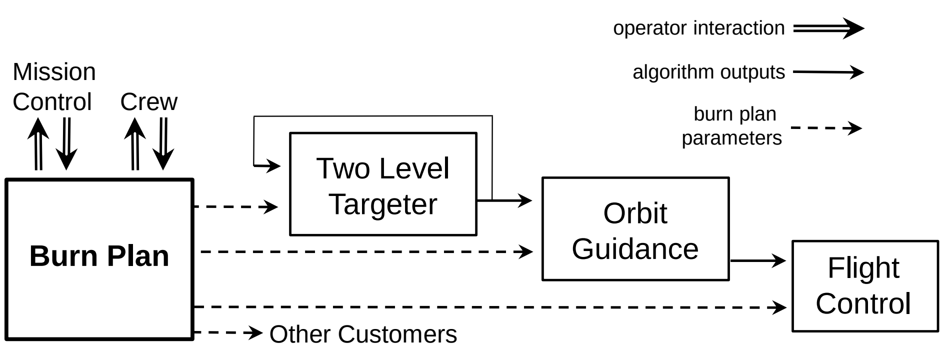

Figure 4. Burn plan data flow.

The TLT can begin its iterative convergence process using initial data provided directly from the burn plan, or data computed during the last successful execution of the TLT. The TLT output is not used to overwrite burn plan data. The feedback of the TLT outputs is particularly applicable during loss of communications scenarios to help ensure convergence, but may be used nominally as well to reduce the number of iterations required for convergence.

Similarly, OrbGuid can execute a burn using the guidance targets provided by the TLT, or the guidance targets from the burn plan. This allows for guidance targets to be sent directly to OrbGuid while bypassing the TLT. OrbGuid has seven guidance target options and the TLT computes primary guidance targets of the same type as the targets specified in the burn plan. [7] Guidance target redundancy has two purposes. First, it allows for one-off burns, targeted by Mission Control, to be inserted into the burn plan. These may be needed for collision avoidance or to accommodate other changes in the mission plan. Second, it allows for critical burns to be executed even in the event of a TLT failure.

OrbGuid always provides the last valid steering command to the flight control system. The steering commands output from OrbGuid are valid for the timeframe of the entire burn arc. Therefore a stale command may be used to successfully complete the burn in an open loop fashion. Using the most recent command ensures that the steering parameters reflect the most recent navigation state. Clearly, steering parameters from the burn plan would become stale and should only be used in the unlikely event that guidance never successfully converges for a critical burn, or in the case of an IMU velocity dilemma.

Control of Sequencing

The burn plan contains not only burn attributes such as guidance targets and engine selection, but also control over specific sequencing throughout the burn timeframe. The control over automated sequencing is an important aspect of Orion being a human-rated vehicle such that the crew or Mission Control can intervene as required where automation is limited or during failure cases that the on-board automation is not designed to handle.

The burn plan allows specific sequencing items to be managed for each burn. The types of sequencing control for burns are shown in Table 1.

Other control of sequencing is supported via specific commands that could be issued during the mission. Two examples are commands to enable or inhibit the TVC gimbal check before the burn and the gimbal check just after the burn. These are not part of the burn plan itself, but are commands available to the crew or Mission Control during the burn time frame. In this case, the crew or Mission Control determines whether or not the gimbal checks occur as part of the normal burn sequencing.

BURN FAILURE AND AUTOMATED RESPONSES

Orion is being designed to autonomously handle a first failure that could result in a catastrophic hazard. Many software and hardware failures have an increased hazardous impact if they occur during burn operations. As a result, specific automation has been included to handle failures from burn targeting through the end of burn reconfiguration. This automation is applicable to the orbit phase only and not for SM abort burns in the ascent phase where Orion may be required to achieve orbit under failure scenarios.

To determine the scope of the automation, failures impacting burn operations were determined along with the optimal software response. The failure list was then reduced through a decision tree process. Some failures take so much capability from the vehicle they were deemed out of scope for automation in the primary flight software, and in many cases, backup flight software and operation modes exist to handle these situations. Examples of these failures include if the TLT converges to bad guidance targets or if a cabin or other consumable leak occurs that impacts vehicle control.

Table 1. Types of control of automated burn sequencing.

Sequencing Control	Description
Criticality	Control burn failure responses based on whether the burn is critical, mandatory, or non-critical.
Authorization to Maneuver to the Burn Attitude	Provide authorization to perform the maneuver to the burn attitude.
Authorization to Execute the Burn	Provide authorization to execute the burn.
Enabling the IMU Alignment Maneuver	Enable or inhibit the maneuver associated with an IMU alignment before the burn.
Loading Mission Control Guidance Targets In Lieu of On-board-Computed Guidance Targets	Loading alternate guidance targets from Mission Control for the burn, rather than executing the on-board TLT or using the guidance targets computed by the on-board TLT.
Allowed Engine Downmodes	Control which engine downmodes are allowed for certain burns.
Solar Array Park Time	Change the time of solar array positioning and parking relative to the time-of-ignition.

Some single-fault failures may impact the desire to do a burn but do not impact the vehicle's ability to perform the burn. For example, a gaseous nitrogen tank failure means there is only one remaining start for the OMS-E engine. The subsequent OMS-E burn could be performed but it may be more prudent to save the last OMS-E start for a future burn. Therefore, there exists a class of failures where a human-in-the-loop is more capable to respond than software. Mission Control, and crew for EM-2, will have the ability to cancel a burn before it begins or stop a burn in progress in response to these types of situations and automation is not required.

Since Orion requirements state the vehicle must handle the first failure, scenarios of multiple failures were deemed out of scope for automation except for IMU dilemmas. Orion has three IMUs that provide $\Delta\theta$ and ΔV measurements. All three IMUs are planned to be operational during burns with one IMU powered down during quiescent flight for power savings. Therefore an IMU dilemma is expected to be a two-failure scenario during burn operations. If it is decided to perform non-critical burns with the third IMU powered down as with the Space Shuttle, then this is a one-failure scenario during burn operations. Even as a two-failure scenario, the impact of an IMU dilemma upon vehicle performance could be so severe that an automated response is provided for an IMU dilemma.

Due to the one-failure requirement, and for flight software simplicity, it is assumed failure triggers do not occur simultaneously. Table 2 identifies the fifteen failure triggers with automated software response during burn operations.

Each failure trigger is checked before each burn (each burn is a clean slate). Mission Control and the crew have the ability to inhibit a trigger should a trigger be a false-positive alert or the decision is made to proceed regardless.

Table 2. Burn failure triggers with automated responses.

Solar Arrays not in the proper position to minimize structural loading during the burn.
Solar Array brake not engaged.
TLT fails to converge pre-burn.
Orbit Guidance fails to converge.
OMS-E TVC does not power up or achieve commanded position before the burn.
IMU velocity dilemma.
IMU attitude dilemma.
Attitude or rate errors exceed allowable values during a burn.
The vehicle does not achieve the burn attitude within tolerance before the expected time of ignition.
OMS-E failure.
High oxidizer or fuel tank pressure beyond engine start conditions.
Pressure Regulation Unit failure.
Failure of any of the eight AUX (+X) engines.
A non-crucial burn preparation step is not complete before Time of Ignition (TIG). Non-crucial is defined to mean any burn preparation step that is not required to physically light the engine.
Failure to receive the Authority to Proceed command before TIG from either Mission Control or the crew. This command provides the software with final permission to execute the burn.

The responses to the failure triggers are simplified to six software capabilities. Thus the response of the vehicle to each failure may not be the best thing to do in each situation but consolidation reduces automation complexity. Table 3 captures each of these automated responses, which can also be invoked via Mission Control or crew command.

The mapping of failure trigger to software response is dependent upon whether the failure occurs pre- or post-ignition and the criticality of the burn. For example, if the TVC system (the OMS-E gimbals) do not achieve the pre-burn position for a non-critical burn, the best thing to do is cancel the burn to provide time for troubleshooting since burn execution can be delayed. However, for a mandatory or critical burn, the software response is to downmode from the OMS-E to AUX (+X) where the TVC system is not required. Once the burn has begun, feedback from the TVC system is not the best indicator of burn performance. After TIG this failure trigger is ignored entirely and if there is a problem in the TVC system an attitude and rate error trigger will trip. This helps limit the effect of false-positive system indications or TVC failures that are too small to affect vehicle control.

CONCLUSION

The exploration mission flight profiles, communications time delays, and a requirement for vehicle return in the event of loss of communications drive a need for on-board automation and au-

Table 3. Burn failure responses.

Failure Response	Description
Cancel Burn	Do not perform a burn that has not yet begun.
Stop Burn	Cutoff a burn that has started and do not trim remaining ΔV.
Downmode	Continue burn with a lower thrust engine if a downmode is permitted per the burn plan. If not, cancel or stop the burn.
TIG Slip	Delay the burn up to the amount of time permitted by the burn plan. Execute the burn if the failure trigger clears within this window.
Burn Previous Steering Commands and Cut-off Burn via Timer	Orbit Guidance outputs a solution for the entire burn each software cycle. Thus, if there is a problem in a subsequent cycle (e.g. IMU dilemma) the solution from the previous cycle can be used to steer the remaining burn arc. The burn is then ended via a timer since the real-time ΔV feedback to the software may be in error.
Ground Guidance Targets	Execute burn with guidance targets uplinked by Mission Control ahead of the burn instead of using the guidance targets computed by the TLT.

tonomy. The on-board sequencing and management of burns enables autonomous vehicle return to Earth in the event of a loss of communications. Furthermore, crew or Mission Control override of automated burn sequencing is needed to protect the crew in the case of failures. The approach for managing on-orbit burn execution on Orion handles both nominal and failure response scenarios. The burn management strategy takes into account per-burn variations in targeting, timing, and execution; crew and Mission Control intervention and overrides; defined burn failure triggers and responses; and corresponding on-board software sequencing functionality.

REFERENCES

[1] J. P. Gutkowski, T. F. Dawn, and R. M. Jedrey, "Evolution of Orion Mission Design for Exploration Mission 1 and 2," *AAS Guidance and Control Conference*, No. AAS 16-111, Breckenridge, CO, February 2016.

[2] S. K. Scarritt, B. G. Marchand, A. J. Brown, W. H. Tracy, and M. W. Weeks, "Finite-Burn Linear Targeting Algorithm for Autonomous Path Planning and Guidance," *Journal of Guidance, Control, and Dynamics*, Vol. 35, September-October 2012, pp. 1605–1615, 10.2514/1.54249.

[3] C. Ocampo and J. Senent, "The Design and Development of Copernicus: A Comprehensive Trajectory Design and Optimization System," *International Astronautical Congress*, No. IAC-06-C1.4.04, 2006.

[4] H. Hickman, "Evolution of MPCV Service Module Propulsion and GN&C Interface Requirements Between Constellation and European Service Module," *AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, Cleveland, OH, July 2014.

[5] C. D'Souza, G. Holt, R. Gay, and R. Zanetti, "Navigation Design and Analysis for the Orion Cislunar Exploration Missions," *AAS Rocky Mountain Guidance and Control Conference*, February 2014.

[6] R. G. Odegard, T. K. Sliwinski, E. T. King, and J. J. Hart, "Configuring the Orion Guidance, Navigation, and Control Flight Software for Automated Sequencing," *IEEE Aerospace Conference*, Big Sky, Montana, March 2011, pp. 1–13, 10.1109/AERO.2011.5747472.

[7] S. K. Scarritt, T. J. Fill, and S. B. Robinson, "Advances in Orion's On-Orbit Guidance and Targeting System Architecture," *AAS Guidance and Control Conference*, No. AAS 15-106, Breckenridge, CO, January-February 2015.