POLITECNICO DI TORINO
Repository ISTITUZIONALE

Enhancing the iDREAM Toolset for Human Lunar Lander Design: A Tightly Coupled Multidisciplinary Design Loop Integrating Mission Analysis, Cost Estimation, and Technology

Original
Enhancing the iDREAM Toolset for Human Lunar Lander Design: A Tightly Coupled Multidisciplinary Design Loop Integrating Mission Analysis, Cost Estimation, and Technology Roadmapping / Rimani, J., Luccisano, G., Fusaro, R., Viola, N., Steelant, J.. - ELETTRONICO. - (2026). (77th International Astronautical Congress (IAC 2026) Antalya (Türkiye) 5-9 October 2026).

Availability:
This version is available at: 11583/3017542 since: 2026-11-04T20:49:25Z

Publisher:
International Astronautical Federation (IAF)

Published
DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

21 November 2026

IAC-26,D1,6,11,x109390

Enhancing the iDREAM Toolset for Human Lunar Lander Design: A Tightly Coupled Multidisciplinary Design Loop Integrating Mission Analysis, Cost Estimation, and Technology Roadmapping

Jasmine Rimani[a,*], Giacomo Luccisano[b,*], Roberta Fusaro[b], Nicole Viola[a], Johan Steelant[c]

[a] *Department of Management and Production Engineering (DIGEP), Politecnico di Torino, 10129 Turin, Italy.*
[b] *Department of Mechanical and Aerospace Engineering (DIMEAS), Politecnico di Torino, 10129 Turin, Italy.*
[c] *ESA ESTEC, Noordwijk, The Netherlands.*
[*] Corresponding Authors: jasmine.rimani@polito.it (J.R.); giacomo.luccisano@polito.it (G.L.).

Abstract

The design of Human Landing Systems (HLS) for sustained lunar exploration requires concurrent evaluation of vehicle performance, mission feasibility, cost, and technological maturity from the earliest design phases. Existing conceptual design frameworks, however, address these dimensions sequentially: cost estimation and technology assessment are performed downstream of mission analysis and vehicle sizing, with no feedback loop to the design. This decoupling is particularly consequential for crew-critical subsystems such as the Environmental Control and Life Support System (ECLSS), whose feasibility depends strongly on technology maturity and which represents one of the principal bottlenecks to achieving a viable, closed-vehicle design. The iDREAM toolset, developed jointly by Politecnico di Torino and ESA, addresses this gap with a unified framework integrating conceptual vehicle design (ASTRID-H), life-cycle cost estimation (HyCost), and technology readiness assessment with roadmap generation (TRIS), supported by a common reference database. A first development phase validated the toolset against a reference micro-launcher at the system and subsystem levels. Furthermore, it explored the overall methodology through a pre-phase A sizing of a lunar lander, focusing on wet and dry mass closure and a simplified mission analysis routine without a detailed subsystem breakdown. This paper presents two methodological enhancements introduced in the second development phase. First, vehicle sizing and mission analysis are coupled in a closed iterative loop: reduced-order models for lunar orbit transfer, powered descent, and ascent compute the delta-v budget and propellant mass fractions, feeding back into ASTRID-H for resizing until mass convergence within 1%. Second, the sizing now iterates on subsystem-level mass, power, volume, and technology allocation. ECLSS is embedded directly within this loop as the subsystem that most constrains the overall design and differentiates crewed vehicles from unmanned ones, allowing the design to respond to ECLSS technology assumptions in real time during convergence. The converged design is then evaluated against the technology roadmapping and cost routines to assess whether available technologies can satisfy the subsystem-level requirements within programmatic constraints, enabling a sensitivity analysis linking technology-readiness-level assumptions, critical for regenerative life-support technologies, to programme cost and development schedule. Where no match exists between the design and the roadmap, an alternative technology selection and ECLSS closure is reinvestigated. Application to a reference HLS short-sortie mission to the lunar south pole demonstrates the framework, with focus on ECLSS mass sensitivity to mission duration and regeneration technology assumptions, and identification of the technology development priorities most strongly driving programme cost and schedule.

Keywords: lunar lander; conceptual design; trajectory modelling; vehicle sizing; multidisciplinary design; iDREAM; life support; cost estimation; technology roadmapping.

1 Introduction

Human lunar landers bring together requirements that are strongly connected from the earliest design activities. The vehicle must deliver the required payload to the lunar surface, support the planned operations and, when required by the mission architecture, return the crew and the ascent payload to orbit. Each of these functions contributes to the vehicle mass and changes the propulsion demand. A useful conceptual design method must therefore link the mission definition to the vehicle's physical characteristics through explicit, consistent interfaces.

The life support system makes this coupling stronger for a crewed lander. The ECLSS mass grows with the number of crew members and the time they spend on board, and the time on board depends on the trajectory: transfers, orbital loiters, and waiting for the return window add days to the surface stay. The technology selected for each ECLSS function then determines whether mass grows with duration, as with expendable consumables, or stays close to a fixed hardware value, as with regenerative equipment. Regenerative technologies, however, are often less mature, so the lightest ECLSS on paper may be the one that the programme cannot develop in time or within budget.

The iDREAM methodology was developed to support the preliminary design of aerospace vehicles and assess their economic and technological viability. The previous paper introduced the integrated toolset and described its application to a microlauncher reference case [1]. Vehicle design and mission analysis form the methodology's engineering core, while HyCost and TRIS support cost estimation and technology roadmapping. The first iDREAM contract also applied

Nomenclature

a, e, p	Semi major axis, eccentricity and semilatus rectum
C	Development and production cost
$c_k,\ \tau_k$	Cost and time fractions of the transit to TRL k
g_0	Standard terrestrial gravitational acceleration
I_{sp}	Specific impulse
L_*, T_*	CR3BP length and time scales
m, T	Vehicle mass and maximum thrust
N_c	Number of crew members
r, θ	Lunar radius coordinate and in plane angle
$R_{\mathrm{M}}, \mu_{\mathrm{M}}$	Lunar radius and gravitational parameter
$t_c,\ t_s$	Crewed duration and surface stay
v_r, v_t	Radial and tangential velocity
α, η	Thrust angle from the local vertical and throttle
Δv	Propulsive velocity increment
ϵ	Relative mass change between iterations
μ	Dimensionless lunar mass parameter
τ	Dimensionless CR3BP time
AMCM	Advanced Missions Cost Model
CaC	Cost at Completion
CR3BP	Circular restricted three body problem
DOI, PDI	Descent orbit insertion, powered descent initiation
ECLSS	Environmental Control and Life Support System
EVA	Extravehicular activity
HLS	Human Landing System
LLO	Low lunar orbit
MER	Mass estimating relationship
NRHO	Near rectilinear halo orbit
OCP	Optimal control problem
TRL	Technology Readiness Level

the methodology to a pre-phase A lunar lander [3, 4]. That study closed the vehicle's wet and dry mass, but the mission analysis was a simplified routine; the loop between sizing and trajectory required the user to restart the simulation by hand; and the cost and roadmap routines ran after the design, with no path back to it. This structure provides the starting point for extending the engineering routines to human lunar landers. For a lander, the trajectory cannot be reduced to a single prescribed velocity increment without losing information relevant to sizing. The parking orbit defines the conditions at descent initiation. Engine thrust and throttle capability influence the duration of powered flight and the propellant consumed while counteracting gravity. The ascent payload and the mass retained on the surface determine the vehicle configuration at liftoff. When the mission includes an NRHO, the transfer geometry and available return window also link mission duration to propulsion requirements.

This paper presents the two enhancements introduced in the second development phase, iDREAM 2.0, and applies them to a human lander. First, the ASTRID-H sizing routine and a new set of reduced-order mission models run in one closed loop: optimal transfers between the NRHO and LLO in the Earth-Moon circular restricted three-body problem, and optimal powered descent and ascent in a planar lunar model. The loop exchanges velocity increments, propellant, installed thrust and crewed duration until the vehicle mass changes by less than 1% between iterations. Second, the sizing works at the subsystem level: an ECLSS module selects one technology per life support function and returns its mass, power, and volume inside the loop, while propulsion, envelope, communication, and avionics modules detail the rest of the vehicle. The converged design then feeds the cost and roadmapping routines of the first iDREAM version. When the roadmap cannot mature a selected technology within the programme window, the process selects an alternative technology and closes the vehicle again.

The paper applies the process to a reference short sortie to the lunar south pole: four crew members, a 72 h surface stay, and a two-stage lander that departs from and returns to the Gateway in NRHO. The results show how the surface stay drives lander mass more than ECLSS consumables when regeneration pays off, and which technology developments drive the programme schedule.

Section 2 recalls the iDREAM background. Section 3 describes the closed design loop, the mission models, the sizing modules and the cost and roadmapping routines. Section 4 verifies the modules, presents the new mission analysis for the NASA HLS sortie, and presents the reference design and its sensitivity to mission duration, ECLSS technology, and technology readiness. Section 6 summarises the findings and the next steps.

2 Methodological background

Conceptual vehicle design benefits from an iterative treatment of requirements, architecture and engineering analysis. The NASA Systems Engineering Handbook describes design development through successive definition, assessment, and refinement activities, with traceability between requirements and the selected solution [2]. For this application, this principle is expressed through the exchange of vehicle mass, propulsion capability, mission duration, and trajectory demand.

iDREAM combines three capabilities around a common database [1, 3]. ASTRID-H performs the conceptual design of the vehicle and of its mission. HyCost estimates life cycle cost; for the lander, it applies an analogy method [3, 5, 6]. TRIS generates technology roadmaps in five steps: stakeholder analysis, element definition, prioritisation, planning and results evaluation, and it plans the TRL transitions of each technology with the budget and time they need [7]. The HyDat and TREX databases store the technologies, their TRL milestones and their Cost at Completion (CaC) [8].

For crewed landers, Isaji et al. derived mass-estimating relationships from Apollo and Constellation-era designs and used them to size a lander based on crew, surface stay, payloads, and velocity increments [9]. NASA studies of the NRHO-staged

Algorithm 1 Closed design loop of the iDREAM 2.0 lander routine

Require: $\boldsymbol{q}$: crew N_c, surface stay t_s, site, payloads down and up, propellant and T/W of each stage, mission profile, one technology per ECLSS function; $\epsilon_{\mathrm{req}} = 1\%$

1: $\Delta v_d, \Delta v_a \leftarrow$ reference values of the profile
2: $t_c \leftarrow t_s + 2\,t_{\max},\ k \leftarrow 0$
3: **repeat**
4: $\quad m_{\mathrm{ECLSS}}, P, V \leftarrow \textsc{Eclss}(N_c, t_c, t_s)$ (Eqs. 25–27)
5: $\quad \boldsymbol{x}^{(k)} \leftarrow \textsc{Mass}(\Delta v_d, \Delta v_a, m_{\mathrm{ECLSS}})$ (Eqs. 21–23)
6: $\quad$ **for** stage $\in$ {ascent, descent} **do**
7: $\quad\quad$ engines, $T_{\mathrm{cl}} \leftarrow \textsc{Propulsion}(m_0, \dot{m}_{\max})$ (Eq. 29)
8: $\quad$ **end for**
9: $\quad \boldsymbol{a}^{(k)} \leftarrow \textsc{MissionAnalysis}(\boldsymbol{x}^{(k)}, T_{\mathrm{cl}})$ (Eqs. 6, 14)
10: $\quad$ update $\Delta v_d, \Delta v_a, t_c, \dot{m}_{\max}$ from $\boldsymbol{a}^{(k)}$ (Eqs. 24, 20)
11: $\quad k \leftarrow k + 1$
12: **until** $\epsilon_m^{(k)} \le \epsilon_{\mathrm{req}}$ and the mission is feasible (Eq. 2)
13: size tanks, envelope and fairing, communication, avionics
14: $C \leftarrow \textsc{HyCost}(m_{\mathrm{dry}}, \ldots)$ (Eq. 32)
15: **for** each technology j of the design **do**
16: $\quad C_{\mathrm{rem},j}, t_{\mathrm{rem},j} \leftarrow \textsc{Tris}(\mathrm{TRL}_j, \mathrm{CaC}_j)$ (Eq. 33)
17: $\quad$ **if** $t_{\mathrm{rem},j} > W$ **then**
18: $\quad\quad$ replace j with its most mature alternative; go to line 3
19: $\quad$ **end if**
20: **end for**
21: **return** design, cost C and roadmap

HLS optimise transfers, powered descent, and ascent as a single multiple-shooting problem [10], and the Altair reference trajectories define the terminal descent profile a crewed lander flies [11]. The present work brings these elements into the iDREAM loop at a fidelity suitable for pre-phase A trade studies.

Lunar mission experience also supports explicitly separating orbital preparation from powered flight. The NASA report on mission planning for lunar module descent and ascent describes the sequence of descent orbit preparation, powered descent and ascent, together with the operational considerations that shaped these phases [12]. This provides a physical reference for the mission decomposition used here. The present model uses its own numerical controls and boundary conditions and does not reproduce the Apollo flight guidance system.

For missions involving an NRHO, the Earth and Moon both contribute to the orbital environment. NASA investigations of NRHO operations address orbit maintenance, navigation and attitude control, illustrating the additional considerations associated with this mission domain [13]. The lander routine represents the transfer geometry at a level suitable for preliminary sizing. It retains a simpler description close to the lunar surface, where a central lunar gravity model provides the basis for the powered flight equations.

3 Integrated lander design methodology

3.1 Closed design loop

Algorithm 1 gives the loop in pseudo-code. The user defines the crew, the surface stay, the landing site, the payload delivered to the surface and the payload returned, the propellant and the thrust-to-weight ratio of each stage, a mission profile and one technology for each ECLSS function. From these inputs, each iteration runs four modules in sequence. The ECLSS module sizes life support for the crewed duration t_c, the time between departure from the NRHO and return to it. The mass estimation module closes the vehicle dry mass and computes the propellant for each stage from the velocity increments. The propulsion allocation selects the engines of each stage and returns the thrust they install. The mission analysis flies the resulting vehicle through the entire mission and returns the velocity increment and propellant for each leg, the peak mass flow of the powered legs, and the crewed duration.

The intended coupling can be represented by a trajectory operator $\mathcal{A}$ and a sizing operator $\mathcal{S}$:

$$\begin{aligned} \boldsymbol{a}^{(k)} &= \mathcal{A}(\boldsymbol{x}^{(k)}, \boldsymbol{q}), \\ \boldsymbol{x}^{(k+1)} &= \mathcal{S}(\boldsymbol{a}^{(k)}, \boldsymbol{q}). \end{aligned} \tag{1}$$

The design vector $\boldsymbol{x}$ includes stage dry masses, propellant allocations, thrust and subsystem selections. The analysis output $\boldsymbol{a}$ includes the velocity increments, consumed propellant, durations, terminal states and feasibility diagnostics. The mission requirements $\boldsymbol{q}$ remain fixed within a given design case.

In the loop, $\boldsymbol{x}$ holds the wet mass of the vehicle and of the ascent stage, the usable propellant of each stage, the installed thrust and the ECLSS mass, and $\boldsymbol{a}$ holds the velocity increment assigned to each stage, the peak mass flow and t_c. At each iteration, the vehicle definition supplies the trajectory model with consistent initial masses and propulsion characteristics. The resulting demand then updates the propellant, propulsion and subsystem budgets. When a discrete technology or engine configuration changes, the trajectory must be reevaluated for that configuration.

The loop stops when

$$\epsilon_m^{(k)} = \max_{i \in \mathcal{I}} \frac{|m_i^{(k+1)} - m_i^{(k)}|}{\max(|m_i^{(k)}|, m_{\mathrm{floor}})} \le \epsilon_{\mathrm{req}}, \tag{2}$$

where $\mathcal{I}$ includes the principal stage and vehicle mass quantities. A small $m_{\mathrm{floor}} > 0$ prevents division by zero for inactive components. In this work $\mathcal{I}$ contains the vehicle and stage wet masses, the dry mass and the usable propellant of each stage, $m_{\mathrm{floor}} = 1$ kg and $\epsilon_{\mathrm{req}} = 1\%$. Acceptance additionally requires a feasible trajectory, nonnegative propellant balances, consistent stage boundaries and satisfaction of the selected subsystem requirements.

After convergence, the detailed modules size the propellant tanks, the lander envelope inside a launcher fairing, the communication link and the avionics. HyCost and TRIS then evaluate the vehicle. If the roadmap shows that a selected technology cannot reach TRL 9 within the programme window, the process replaces it with a more mature alternative and runs the loop again (lines 13 to 20 of Algorithm 1).

3.2 Mission analysis

The mission analysis follows the NRHO-staged lunar sortie of NASA studies [10]. The lander leaves the Gateway in the 9:2 NRHO, transfers to a 100 km circular polar LLO, loiters for two navigation revolutions and one phasing revolution, lowers its perilune to 15 km with the descent orbit insertion (DOI) burn, coasts for half an orbit and starts the powered descent at perilune. After the surface stay, the ascent stage lifts off, reaches the perilune of a 15 by 100 km insertion ellipse, circularises at apolune, loiters for at least one revolution and transfers back to the Gateway. The routine also runs the one-way mission and the LLO-staged missions of Apollo and LSAM. Mission profiles store standard combinations of surface stay, transfer time cap and return loiter window.

The orbital transfers use the CR3BP in an Earth- and Moon-rotating frame, and the powered legs use dimensional polar coordinates centred on the Moon, as in the previous formulation. Analytical relations connect the circular LLO, the descent ellipse and the ascent insertion ellipse.

3.2.1 Optimal transfers between NRHO and LLO

Let L_* denote the reference separation of the primaries and define

$$T_* = \sqrt{\frac{L_*^3}{\mu_{\mathrm{E}} + \mu_{\mathrm{M}}}}, \quad V_* = \frac{L_*}{T_*}, \quad \mu = \frac{\mu_{\mathrm{M}}}{\mu_{\mathrm{E}} + \mu_{\mathrm{M}}}. \tag{3}$$

The dimensionless rotating frame places Earth at $(-\mu, 0, 0)$ and the Moon at $(1-\mu, 0, 0)$. With prime denoting differentiation with respect to $\tau = t/T_*$, the potential is

$$\Omega = \frac{x^2 + y^2}{2} + \frac{1-\mu}{r_1} + \frac{\mu}{r_2}, \tag{4}$$

where r_1 and r_2 are the distances to Earth and Moon in units of L_*. The equations of motion are

$$\begin{aligned} x'' - 2y' &= \frac{\partial \Omega}{\partial x}, \\ y'' + 2x' &= \frac{\partial \Omega}{\partial y}, \\ z'' &= \frac{\partial \Omega}{\partial z}. \end{aligned} \tag{5}$$

Propagate an NRHO reference state to the selected departure epoch. The epoch is a time coordinate along the reference orbit. It is not a Keplerian true anomaly, since the NRHO is a periodic solution in the CR3BP.

Each transfer is a two-burn optimal control problem (OCP). For the outbound leg,

$$\begin{aligned} \min\ & \|\Delta \boldsymbol{v}_1\| + \|\Delta \boldsymbol{v}_2\| \\ \text{s.t. } & \boldsymbol{s}' = \boldsymbol{f}_{\mathrm{CR3BP}}(\boldsymbol{s}), \quad \tau \in [\tau_d, \tau_a], \\ & \boldsymbol{r}(\tau_d) = \boldsymbol{r}_{\mathrm{G}}(\tau_d), \\ & \boldsymbol{v}(\tau_d^+) = \boldsymbol{v}_{\mathrm{G}}(\tau_d) + \Delta \boldsymbol{v}_1, \\ & \boldsymbol{r}(\tau_a) = \boldsymbol{r}_{\mathrm{M}} + r_{\mathrm{L}}\, \hat{\boldsymbol{r}}, \\ & \boldsymbol{v}(\tau_a^-) + \Delta \boldsymbol{v}_2 = \boldsymbol{v}_{\mathrm{L}}(\hat{\boldsymbol{r}}, \hat{\boldsymbol{t}}), \\ & \hat{\boldsymbol{r}} \cdot \hat{\boldsymbol{t}} = 0, \quad (\hat{\boldsymbol{r}} \times \hat{\boldsymbol{t}}) \cdot \hat{\boldsymbol{z}} = 0, \\ & \tau_d \in W_d, \quad t_a - t_d \in [t_{\min}, t_{\max}], \\ & \|\boldsymbol{r} - \boldsymbol{r}_{\mathrm{M}}\| \geq R_{\mathrm{M}} + h_{\min}, \end{aligned} \tag{6}$$

where $(\boldsymbol{r}_{\mathrm{G}}, \boldsymbol{v}_{\mathrm{G}})$ is the Gateway state on the reference NRHO at epoch τ and $\boldsymbol{v}_{\mathrm{L}}$ is the circular velocity of the LLO at the arrival point, along the unit vector $\hat{\boldsymbol{t}}$, expressed in the rotating frame. The two orthogonality conditions define a polar plane: its normal is perpendicular to the Earth and Moon rotation axis, so the orbit overflies both poles and serves a south pole site. The departure epoch, arrival point, plane orientation, and time of flight are free within the departure window W_d, the time-of-flight bounds, and an altitude floor $h_{\min} = 30$ km. The return leg solves the same problem backwards: the LLO departure is free inside a window of extra loiter after the ascent, and the arrival state must equal the Gateway state at the arrival epoch. The lander therefore meets the Gateway by construction, whatever the surface stay.

A multiple shooting scheme transcribes the problem on a regularised time s,

$$\frac{\mathrm{d}t}{\mathrm{d}s} = \rho^{3/2}, \qquad \rho = \frac{\|\boldsymbol{r} - \boldsymbol{r}_{\mathrm{M}}\|}{r_{\mathrm{L}}}, \tag{7}$$

which concentrates the 80 shooting intervals near the Moon, where the dynamics are fast. A B-spline of the reference NRHO gives $\boldsymbol{s}_{\mathrm{G}}$ as a smooth function of the epoch. The initial guesses come from a family of transfer arcs computed once over one NRHO period; the return leg uses the same arcs mirrored through the $y = 0$ symmetry of the CR3BP. CasADi [14] provides the exact derivatives and IPOPT [15] solves the problem. An independent propagation of each solution with an eighth-order Runge-Kutta integrator accepts it only when the arc closes within 2 km. Solving the problem for several caps on the time of flight gives the trade between velocity increment and transfer duration.

3.2.2 Orbital preparation for descent

The descent interface is defined through a circular parking orbit of radius r_a and a descent ellipse with perilune radius r_p. Their geometric relations are

$$a = \frac{r_a + r_p}{2}, \qquad e = \frac{r_a - r_p}{r_a + r_p}, \qquad p = a(1 - e^2). \tag{8}$$

The circular and elliptical velocities at apolune give the descent orbit insertion burn,

$$\Delta v_{\mathrm{DOI}} = \sqrt{\frac{\mu_{\mathrm{M}}}{r_a}} - \sqrt{\mu_{\mathrm{M}}\left(\frac{2}{r_a} - \frac{1}{a}\right)}. \tag{9}$$

With no intermediate thrust, the time from apolune to perilune is

$$t_{\mathrm{coast}} = \pi\sqrt{\frac{a^3}{\mu_{\mathrm{M}}}}. \tag{10}$$

Powered descent starts at perilune with zero radial velocity and the ellipse's tangential velocity, $v_{t,0} = \sqrt{\mu_{\mathrm{M}} p}/r_p$, which is higher than the circular velocity at the same altitude.

3.2.3 Powered flight equations

Powered descent and ascent share the state vector

$$\boldsymbol{y} = (r, \theta, v_r, v_t, m)^{\mathsf{T}}. \tag{11}$$

The angle α is measured from the local vertical. Thus, $\alpha = 0$ produces radial thrust; positive α gives a prograde component, and negative α gives a retrograde component. The throttle η scales the maximum thrust T. With a constant specific impulse during the leg, the equations are

$$\begin{aligned} \dot{r} &= v_r, & \dot{\theta} &= \frac{v_t}{r}, \\ \dot{v}_r &= \frac{v_t^2}{r} - \frac{\mu_{\mathrm{M}}}{r^2} + \frac{\eta T}{m}\cos\alpha, \\ \dot{v}_t &= -\frac{v_r v_t}{r} + \frac{\eta T}{m}\sin\alpha, \\ \dot{m} &= -\frac{\eta T}{I_{\mathrm{sp}} g_0}. \end{aligned} \tag{12}$$

The kinematic terms arise from the polar coordinate description. The model assumes motion within a fixed plane, a spherical Moon and central lunar gravity. It neglects lunar rotation, gravity harmonics, terrain, attitude dynamics and external perturbations during the powered leg.

The propulsive velocity increment of a leg follows from the propagated mass,

$$\Delta v = \int_{t_0}^{t_f} \frac{\eta T}{m}\,\mathrm{d}t = I_{\mathrm{sp}} g_0 \ln\frac{m_0}{m_f}. \tag{13}$$

Gravity acts through the trajectory equations and changes the required thrust history and burn duration. Therefore, do not add generic gravity loss to the velocity increment from Eq. (13).

Both powered legs are fuel optimal, multi-phase OCPs:

$$\begin{aligned} \max_{\eta,\alpha,t_p}\ & m(t_f) - w\sum_j \left[(\Delta\alpha_j)^2 + 0.1(\Delta\eta_j)^2\right] \\ \text{s.t. Eq. (12)},\quad & \eta_{\min,p} \le \eta \le \eta_{\max,p}, \\ & \alpha_{\min,p} \le \alpha \le \alpha_{\max,p}, \quad |\dot{\alpha}| \le \dot{\alpha}_{\max}, \\ & r \ge R_{\mathrm{M}} + h_{\min}, \\ & m \ge m_0 - m_{p,\mathrm{avail}}, \end{aligned} \tag{14}$$

together with equality conditions at the end of each phase p. The phase durations t_p are free. The small weight $w = 10^{-6}$ only regularises the control history. Piecewise constant controls on equal intervals, three Runge-Kutta substeps per interval, scaling by R_{M}, $\sqrt{\mu_{\mathrm{M}}/R_{\mathrm{M}}}$, and the initial mass, and IPOPT with exact derivatives yield a solution in less than one second. The minimum throttle is 10%, and the thrust direction may rotate by at most 5 deg/s, as for Altair [11].

The descent flies two phases from PDI to the low gate. Braking ends at 2 km altitude with throttle above 90%; without this bound, the fuel optimum idles at minimum throttle for minutes after perilune, which is equivalent to moving PDI. The approach phase nulls the horizontal velocity and reaches the low gate,

$$r_f = R_{\mathrm{M}} + h_{\mathrm{LG}}, \qquad v_{t,f} = 0, \qquad v_{r,f} = -v_{\mathrm{LG}}, \tag{15}$$

with $h_{\mathrm{LG}} = 150$ m and $v_{\mathrm{LG}} = 5$ m/s. The sink rate may only decrease during the approach, and the thrust must be within 20 deg of the vertical at the low gate, so the lander reaches it upright. A guidance law then flies the soft landing of the Altair concept [11]: it decelerates along a constant deceleration profile to 1 m/s at 30 m, optionally hovers, descends at 1 m/s and cuts the engines at $h_c = 1$ m. The touchdown speed is

$$v_{\mathrm{TD}} = \sqrt{v_{\mathrm{td}}^2 + 2 g_{\mathrm{M}} h_c} \approx 2.06\ \mathrm{m/s}. \tag{16}$$

The guidance integrates the actual mass flow, so the budget includes gravity loss during near-hover terminal descent.

The ascent flies a vertical rise to 100 m and then one continuous burn at full thrust, as the ascent engines of the Apollo LM and Altair [11]. For an insertion ellipse with perilune r_p and apolune r_a, the powered flight boundary conditions are

$$r_f = r_p, \quad v_{r,f} = 0, \quad v_{t,f} = \sqrt{\frac{2\mu_{\mathrm{M}} r_a}{r_p(r_p + r_a)}}. \tag{17}$$

For elliptic insertion, the vehicle then coasts to apolune and circularises. The circularisation magnitude follows Eq. (9), using the ascent ellipse parameters, and its propellant demand is evaluated using the remaining vehicle mass. The earlier shooting models of both legs remain in the routine as a fallback when the OCP does not converge, and the result reports which model answered.

3.2.4 Mission duration and return opportunity

The timing of a return mission depends on the sum of its transfers, powered legs, coasts, orbital loiters and surface stay:

$$\begin{aligned} t_{\mathrm{arr,N}} = t_{\mathrm{dep,N}} &+ \sum t_{\mathrm{transfer}} + \sum t_{\mathrm{powered}} \\ &+ \sum t_{\mathrm{coast}} + \sum t_{\mathrm{loiter}} + t_{\mathrm{surface}}. \end{aligned} \tag{18}$$

A loiter of n revolutions in a circular LLO has duration

$$t_{\mathrm{loiter}} = n\,2\pi\sqrt{\frac{r_{\mathrm{L}}^3}{\mu_{\mathrm{M}}}}. \tag{19}$$

The Gateway keeps moving along the NRHO while the crew is on the surface, so the surface stay fixes the Gateway's position when the lander can leave. The return OCP may add up to 24 h of extra LLO loiter to reach a cheaper geometry, and it imposes the Gateway state at arrival, so the reported rendezvous is exact within the 2 km propagation check. The crewed duration of the mission,

$$t_c = t_{\mathrm{arr,N}} - t_{\mathrm{dep,N}}, \tag{20}$$

is the quantity that sizes the ECLSS consumables in the loop.

3.3 Vehicle sizing

3.3.1 Mass estimation

The mass estimation module closes the vehicle mass in five steps, with the MERs of Isaji et al. [9]. Step 1 iterates on the total dry mass of structure and thermal protection, descent and ascent propulsion, power, avionics, ECLSS and other items, the last taken as a fraction c_o of the dry mass. In general form, the dry mass estimate is written as

$$m_{\mathrm{dry}} = \sum_{i=1}^{n_s} \mathcal{M}_i(\boldsymbol{q}, m_{\mathrm{dry}}), \tag{21}$$

where $\boldsymbol{q}$ contains the mission and architecture inputs and $\mathcal{M}_i$ is the selected mass model for subsystem i. The dry mass on both sides reflects the internal iteration needed by mass-estimating relationships that depend on vehicle size. Step 2 iterates on the descent-stage dry mass with the stage-level MERs, Step 3 obtains the ascent dry mass by difference, and Steps 4 and 5 compute the propellant for the ascent and descent stages. With $R_j = \exp[\Delta v_j/(I_{\mathrm{sp},j} g_0)]$, the usable propellant of the ascent stage is

$$m_{p,a} = \frac{(m_{\mathrm{dry},a} + m_{\mathrm{pl},a})(R_a - 1)}{1 - k_a(R_a - 1)}, \tag{22}$$

where $k_a = 0.159$ includes a 2.7% reserve and a 13.2% reaction-control allowance of the usable propellant [9]. The ascent stage then weighs $m_a = m_{\mathrm{dry},a} + m_{\mathrm{pl},a} + (1 + k_a)m_{p,a}$, and the descent stage carries it:

$$m_{p,d} = \frac{(m_a + m_{\mathrm{dry},d} + m_{\mathrm{pl},d})(R_d - 1)}{1 - k_d(R_d - 1)}, \tag{23}$$

with a 4.3% reserve, $k_d = 0.043$. In an NRHO-staged mission, the descent stage flies the transfer to LLO, DOI and the powered descent, and the ascent stage flies the ascent, the circularisation and the transfer back to the Gateway, so the loop assigns

$$\begin{aligned} \Delta v_d &= \Delta v_{\mathrm{NL}} + \Delta v_{\mathrm{DOI+PDI}}, \\ \Delta v_a &= \Delta v_{\mathrm{asc+circ}} + \Delta v_{\mathrm{LN}}. \end{aligned} \tag{24}$$

Inside the loop, the ECLSS module replaces the ECLSS MER of Step 1, and the dry-mass fixed point runs to a relative tolerance of 10^{-6}, so it does not limit the outer-loop 1% criterion.

3.3.2 Life support

Crew support directly depends on crew size and mission duration. At the level of a resource balance, the demand for a consumable j can be expressed as

$$m_{j,\mathrm{demand}} = N_c q_j t_c + N_{\mathrm{EVA}} q_{j,\mathrm{EVA}}, \tag{25}$$

where q_j is the demand per crewmember and unit time, t_c is the supported crew duration, and N_{EVA} counts individual crew excursions over that duration. The second term is included only for resources associated with extravehicular activity. The NASA Life Support Baseline Values and Assumptions Document provides a reference basis for selecting and documenting such inputs [16]. The module takes 0.84 kg of oxygen, 2.0 kg of drinking water, 0.5 kg of food preparation water, and 2.39 kg of packaged food per crew member per day; 5 kg of cooling water per crew member per EVA; the nitrogen that the cabin leaks; and the stowed waste [16, 17]. A database row describes each technology through coefficients on four scaling bases,

$$\begin{aligned} \{m, V, P\}_i &= \sum_b c_{i,b}\, s_b, \\ s_b &\in \{1,\ N_c,\ N_c t_c,\ m_{j,\mathrm{demand}}\}, \end{aligned} \tag{26}$$

that is fixed, per crew member, per crew day and per kilogram of the stored consumable [17, 18]. The user selects one row for each function: carbon dioxide removal, oxygen supply, water supply, trace contaminant control, temperature and humidity control, nitrogen make-up, and EVA suits. The ECLSS mass is the sum of hardware, expendables and consumables with a 20% growth allowance. In the loop, the consumables follow the crewed duration t_c from Eq. (20), while the EVA water follows the surface stay only.

Regenerative technologies trade expendables for power. The power system MER of the mass estimation covers an open-loop ECLSS, so for a regenerative option the loop adds the battery that stores its extra energy,

$$\Delta m_{\mathrm{b}} = \frac{(P_i - P_{\mathrm{ref}})\, t_c}{e_{\mathrm{b}}}, \tag{27}$$

with an effective specific energy $e_{\mathrm{b}} = 120$ Wh/kg. The results also report the bound $\Delta m_{\mathrm{b}} = 0$, which represents a lander that draws the extra power from a source already on board.

3.3.3 Propulsion, envelope and other subsystems

The propulsion interface is defined in thrust (newtons). When a thrust-to-weight ratio is used, its gravity reference must be stated. With λ_{M} referred

to lunar surface gravity and λ_0 referred to standard terrestrial gravity,

$$T = \lambda_{\mathrm{M}} m g_{\mathrm{M}} = \lambda_0 m g_0, \qquad g_{\mathrm{M}} = \frac{\mu_{\mathrm{M}}}{R_{\mathrm{M}}^2}. \tag{28}$$

The propulsion module takes λ_0 as input, applies a 25% thrust margin and sizes the cluster of each candidate engine e of the database:

$$\begin{aligned} T_{\mathrm{des}} &= 1.25\,\lambda_0 m g_0, \\ n_{\mathrm{req}} &= \left\lceil \frac{T_{\mathrm{des}}}{T_e} \right\rceil, \qquad n_{\mathrm{tot}} = n_{\mathrm{req}} + 1. \end{aligned} \tag{29}$$

The spare engine provides single-engine-out capability. A candidate must also deliver the peak mass flow of the flown trajectory, $n_{\mathrm{req}} T_e / (I_{\mathrm{sp},e} g_0) \geq \dot{m}_{\max}$. The module sizes the stage for every candidate and ranks them by dry propulsion mass. In the loop, the lightest cluster's installed thrust, $T_{\mathrm{cl}} = n_{\mathrm{req}} T_e$, replaces the nominal thrust of the next trajectory, and only throttleable engines qualify for the descent stage, which must hover before touchdown. Given an oxidiser-to-fuel mass ratio χ, the allocated propellant mass is divided as

$$m_{\mathrm{ox}} = \frac{\chi}{1+\chi} m_{\mathrm{prop}}, \qquad m_{\mathrm{fu}} = \frac{1}{1+\chi} m_{\mathrm{prop}}. \tag{30}$$

The associated tank volumes are

$$V_{\mathrm{tank},j} = (1+u_j)\frac{m_j}{\rho_j}, \qquad j \in \{\mathrm{ox}, \mathrm{fu}\}, \tag{31}$$

where u_j is an added volume allowance relative to liquid volume. The module then sizes titanium tanks for the maximum expected operating pressure of the feed system, their insulation and boil-off, the pressurant and its bottle, the feed lines and the reaction control system, following NASA TP-3388 and the ESA margin philosophy. The envelope module packages the ascent stage and the cabin, places the descent tanks in a ring around them with bays for the cargo, and selects the smallest launcher fairing that contains the stowed lander. The communication module closes the link budget to the Gateway at worst-case range, and the avionics module sizes the guidance, navigation, and control sensors, computers, and data handling with N+1 redundancy.

3.4 Cost estimation and technology roadmapping

For the lander, HyCost applies the AMCM [3, 6],

$$C = \alpha\, Q^{\beta} M^{\xi}\, \delta^{S}\, \varepsilon^{1/(\mathrm{IOC}-1900)}\, B^{\phi} \gamma^{D}, \tag{32}$$

where M is the dry mass in pounds, $Q = 3$ the number of units, $S = 2.46$ the specification of a planetary lander, $\mathrm{IOC} = 2024$ the first year of operations, $B = 1$ the block number of a new design and $D = 0$ an average difficulty; $\alpha = 5.65 \times 10^{-4}$, $\beta = 0.5941$, $\xi = 0.6604$, $\delta = 80.599$, $\varepsilon = 3.8085 \times 10^{-55}$, $\phi = -0.3553$ and $\gamma = 1.5691$. The first version extended the result to fiscal year 2020 and checked it against an analogy estimate that interpolates the cost of the closest historical systems using a normalised Euclidean distance on crew, destination, gross mass, and dry mass [3, 5].

TRIS ranks technologies based on stakeholder criteria and plans their TRL transitions. Its database provides the current TRL and the CaC of each technology, and the first version derives the typical share of cost c_k and time τ_k for each transition to TRL k for space exploration systems (Table 1) [3]. The cost and the time still needed from the current TRL L_0 are

$$C_{\mathrm{rem}} = \mathrm{CaC} \sum_{k=L_0+1}^{9} c_k, \qquad t_{\mathrm{rem}} = t_{\mathrm{CaC}} \sum_{k=L_0+1}^{9} \tau_k, \tag{33}$$

where t_{CaC} is the time at completion from TRL 1 to TRL 9. A technology matches the design when t_{rem} fits in the programme window W. Otherwise, the programme either waits, at the delay cost κ per year of the first version, or the loop selects a more mature technology for the same function.

Table 1: Share of the Cost at Completion and of the time at completion spent in the transit to each TRL, for a space exploration system [3].

TRL k	2	3	4	5	6	7	8	9
c_k [%]	2	4	5	7	15	26	26	15
τ_k [%]	7.3	3.6	2.3	1.8	15	40	18	12

4 Results and discussion

4.1 Verification of the updated modules

Table 2 compares the mission analysis with three reference missions. The Apollo LM and the LSAM fly from LLO; the NASA HLS case flies the full NRHO mission with a 6 day stay and 12 h transfers. The model returns an ideal optimum without dispersions, landing-site redesignation, or piloting margins, so it sits below the flown and published budgets: by 8.8% on the Apollo descent and by less than 5% on every other leg. The NASA HLS reference is the patched baseline from Mahajan and Condon [10], which optimises the whole mission with a higher-fidelity model and its own constraints. The largest difference, 4.9% on the outbound transfer, reflects those constraints: for a 12 h transfer of the same study, the NASA design presentation reports 704.9 m/s [20], within 1% of the 698.7 m/s computed here. The Gateway rendezvous is exact in every case, and the touchdown speed is 2.06 m/s.

Table 3 verifies the sizing modules on the same historical vehicles, with the Isaji et al. reference masses [9]. With the dry mass fixed point converged, the mass estimation returns the lift-off mass of LSAM within 5.9% and of the Apollo LM within 8.4%, and the dry mass within 3.0%. The ECLSS module, with LiOH, stored oxygen and stored water, is 12.1% and 6.3% lighter than the two references. The propulsion

Table 2: Mission analysis against reference velocity increments. Descent includes DOI and powered descent; ascent includes circularisation. NASA HLS: 6 d stay, 12 h transfer cap.

Case	Leg	Ref. [m/s]	Model [m/s]	Diff. [%]
Apollo LM [12]	Descent	2104.0	1919.7	−8.8
	Ascent	1958.0	1863.2	−4.8
LSAM [19]	Descent	1985.0	1899.5	−4.3
	Ascent	1821.0	1818.8	−0.1
NASA HLS [10]	NRHO to LLO	734.9	698.7	−4.9
	Descent	1971.9	1896.5	−3.8
	Ascent	1837.0	1818.8	−1.0
	LLO to NRHO	713.3	698.7	−2.0

module is within 12% across three stages; the Apollo descent stage remains 33.5% heavy, and validation is still ongoing.

Table 3: Sizing modules against the reference vehicles [9]. Masses in kg.

Vehicle	Quantity	Ref.	Model	Diff. [%]
LSAM	Lift-off mass	45,862	48,569	+5.9
	Dry mass	11,264	10,927	−3.0
	Usable propellant	29,820	31,137	+4.4
	ECLSS	1,412	1,242	−12.1
	Descent propulsion	2,834	3,153	+11.3
	Ascent propulsion	1,072	945	−11.8
Apollo LM	Lift-off mass	16,372	14,998	−8.4
	Dry mass	4,234	4,120	−2.7
	Usable propellant	10,582	9,630	−9.0
	ECLSS	584	548	−6.3
	Descent propulsion	493	658	+33.5
	Ascent propulsion	324	318	−1.7

4.2 Mission analysis of the NASA HLS sortie

The mission analysis is the new element of the routine, so this section shows how it behaves before it enters the loop. All results in this section use the NASA HLS preset of iDREAM 2.0, whose masses and thrust stay fixed, isolating the trajectory from sizing. With a 6 day stay and 12 h transfers, the lander leaves the Gateway with about 45 t and burns 23.8 t of propellant over the round trip. The powered descent takes 60% of it, 14.2 t in 7 min from PDI to touchdown. The descent stage then stays on the surface, so it lifts off at only 10.0 t, and the ascent burns 4.4 t in 5.2 min. The return transfer needs the same 699 m/s as the outbound leg but only 1.1 t of propellant, because it pushes a 5.6 t vehicle instead of 45 t. This asymmetry explains why the descent stage drives the lander size, while the return velocity increment drives the ascent stage size.

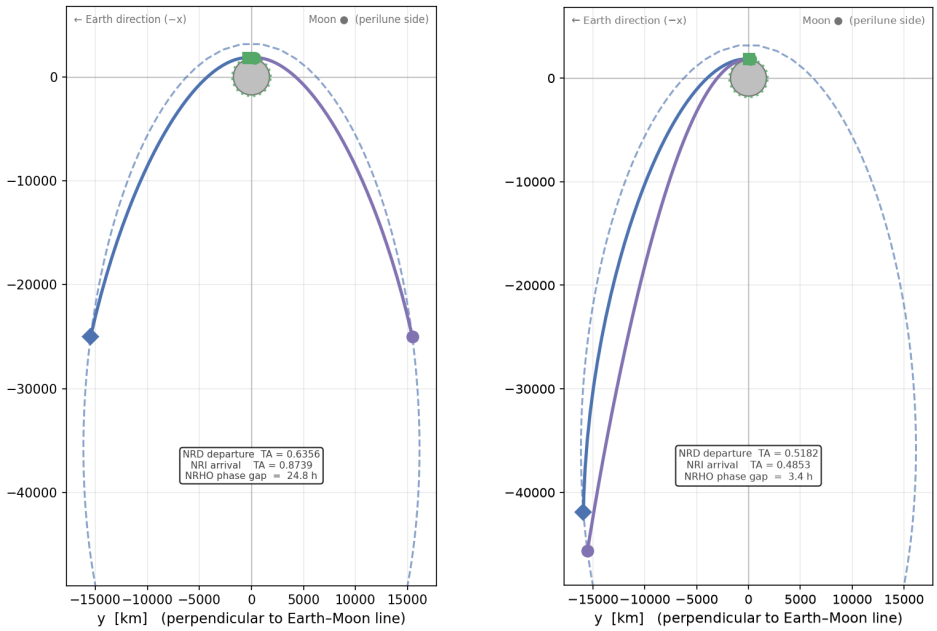
Figure 1 shows the transfers in the Moon-centred synodic frame. In both cases, the outbound leg leaves the Gateway on the apolune side of the NRHO, falls toward the Moon, and inserts into the polar LLO near the NRHO perilune, where the two orbits are closest. With a 6 day stay, panel (a), the return arc is the mirror image of the outbound arc through the $y = 0$ plane, the symmetry of the CR3BP, and it costs the same 699 m/s. With a 3 day stay and 24 h transfers, panel (b), the whole mission lasts 6.41 days, almost exactly one NRHO period of about 6.6 days. The Gateway is therefore back near where the lander left it, and the lander must rejoin it on the same side of the orbit, 3.4 h of phase from its departure point. The mirror geometry is not available, and the return costs 1309 m/s against 686 m/s for the outbound leg, even after 24 h of extra loiter.

Figure 2 gives the price of a shorter transfer. The outbound leg costs 682 m/s when the time of flight is free, and the optimiser finds this minimum at 35 h. A 24 h cap raises the cost to 686 m/s, a 12 h cap to 699 m/s, and a 6 h cap to 720 m/s: halving the time of flight from 24 to 12 h costs only 13 m/s per leg, so 12 h is a good compromise, close to the NASA design choice. The 12 h value is within 1% of the 704.9 m/s of the NASA design presentation [20], and 4% below the 731 m/s in 13 h of Whitley and Martinez [21]. The return front depends on the surface stay. With a 6 day stay, it repeats the outbound front, from 720 to 681 m/s. With a 3 day stay it rises steeply as the cap shortens: 754 m/s in 48 h, 1308 m/s in 24 h, 2379 m/s in 12 h and 4946 m/s in 6 h. A short stay therefore needs either a long return or a large ascent stage.

Table 4: Standard mission profiles of the routine NASA HLS preset. Descent 1897 m/s and ascent 1819 m/s in every profile. Velocity increments in m/s.

Stay [d]	Cap [h]	Out	Return	Total	Prop. [t]	Loiter [h]	Mission [d]
6.0	12	699	699	5113	23.8	4.2	7.59
6.0	24	686	686	5087	23.7	4.7	8.61
3.0	12	699	2379	6793	25.6	24.0	5.41
3.0	24	686	1308	5710	24.5	24.0	6.41
6.5	12	699	1161	5575	24.4	0.0	7.91
6.5	24	686	841	5243	23.9	0.0	8.91

Table 4 runs the six standard profiles of the routine, which combine three surface stays with the two caps. The descent and ascent do not change because the vehicle and site are the same, so the difference lies entirely in the return. The 6 day stay gives the cheapest missions, 5087 and 5113 m/s, because its return mirrors the outbound leg. The 3 day stay with 12 h transfers is the most expensive, 6793 m/s, although it is also the shortest mission. The 6.5 day stay, which matches one NRHO period, leaves LLO just after the mirror geometry has passed: waiting only makes the return more expensive, so the optimiser adds no extra loiter, and the return costs 841 to 1161 m/s. For an NRHO-staged lander, surface stay is therefore as much a trajectory parameter as an operational one: the cheapest stays make the Gateway reach the mirror of the departure point when the lander is ready to leave. Section 4.4 shows what this means for the converged vehicle.



(a) 6 d stay, 12 h transfers

(b) 3 d stay, 24 h transfers

Fig. 1: Optimal transfers in the Moon-centred synodic frame, y and z in km, Earth towards $-x$. Dashed: 9:2 NRHO; black: NRHO to LLO, from the Gateway (diamond); purple: LLO to NRHO, to the Gateway (circle); green: 100 km polar LLO and the insertion and departure points. The boxes give the departure and arrival epochs along the NRHO in CR3BP time units and the phase between them. The view stops at $z = -49{,}000$ km; the NRHO apolune is at about $-70{,}000$ km. (a) NASA HLS preset, 699 m/s for each leg; (b) converged reference sortie of Section 4.3, 686 and 1309 m/s.

The terminal descent is the second new element. Figure 3 (top) follows the NASA HLS preset through the last minute of the descent. The approach phase flies at full throttle until a few seconds before the low gate. Guidance then takes over at 150 m and 5 m/s, tracks the constant-deceleration profile down to 1 m/s at 30 m, descends at constant speed, and cuts the engines at 1 m, so the lander touches down at 2.06 m/s after 70 s. During this phase the throttle stays at about 30% and the engines only balance the lunar weight. The soft landing therefore costs 116 m/s, and 113 m/s of them are simply the lunar gravity, 1.62 m/s^2, integrated over 70 s. The optimised mode of the routine, which solves the same profile as a further collocation phase, returns the same 116 m/s in 70 s.

Figure 3 (bottom) varies the terminal descent assumptions. Raising the low gate from 75 to 300 m adds 122 m/s when the lander crosses it at 5 m/s and 243 m/s at 2 m/s, because a slower gate means a longer near-hover descent. Each 10 s of hover adds about 16 m/s, again the lunar gravity times the hover time. Over the trade, the descent budget spans 1855 to 2128 m/s, a band that contains the 1972 m/s of the NASA reference: a 150 m gate crossed at 2 m/s already gives 1966 m/s. The terminal descent assumptions therefore explain most of the 3.8% gap of Table 2, and they are the first parameters to align with a specific lander concept.

4.3 Reference short sortie

The reference mission follows the ESA requirements for the human lander test case of iDREAM [8]: four crew members, a 72 h stay at the lunar south pole, departure from and return to the Gateway in NRHO. The lander has two stages: a LOX/LH$_2$ descent stage and a storable NTO/MMH ascent stage that carries the cabin, with thrust-to-weight ratios of 0.365 and 0.45 referred to Earth gravity. It delivers 3500 kg to the surface and returns 834 kg, the LSAM allocations [9]. The ECLSS uses LiOH for carbon dioxide removal and stores oxygen and water, with one EVA per day. The transfers may last up to 24 h, and the return may wait up to 24 h in LLO. The loop starts from the velocity increments of the NASA HLS preset of iDREAM 2.0.

Table 5: Convergence of the closed loop for the reference sortie. Δv_d and Δv_a enter the iteration; the engines and the mass leave it.

k	Δv_d [m/s]	Δv_a [m/s]	Engines (desc. / asc.)	m_0 [kg]	ϵ [%]
0	2696.9	2714.9	BE-7 ×5 / Aestus ×3	40,590	-
1	2569.1	3119.5	RL10A-3-3 ×4 / AJ10 ×4	45,749	33.9
2	2568.7	3119.5	RL10A-3-3 ×4 / AJ10 ×4	45,745	0.02

Table 5 shows that the loop converges in three iterations, well below the 1% criterion. The first mission analysis corrects the literature's velocity increments: the descent stage needs 128 m/s less. In contrast, the ascent

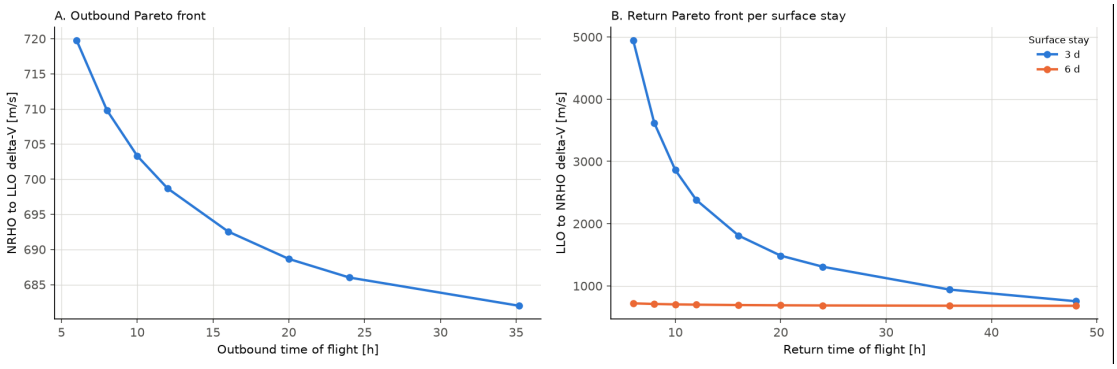


Fig. 2: Optimal transfer velocity increment against the cap on the time of flight, NASA HLS preset. (A) outbound leg; (B) return leg for a 3 day and a 6 day stay, with up to 24 h of extra LLO loiter.

stage needs 405 m/s more, because the 72 h stay forces a costly return to the Gateway. The heavier ascent stage then shifts the engine allocation to four RL10A-3-3 and four AJ10 engines, with one spare per stage. The second trajectory changes the velocity increments by less than 0.5 m/s, so the mass settles. The first iteration takes about 70 s, most of which builds the family of transfer seeds; each following iteration takes about 3 s.

The allocation matters for the trajectory. The installed thrust exceeds the nominal thrust by 63% on the descent stage and by 34% on the ascent stage, because engines come in discrete sizes and the design thrust includes a 25% margin. With higher thrust, the powered descent lasts 5.0 min instead of 7.2 min and loses less to gravity: the lander is 302 kg (0.7%) lighter than in the same loop run with nominal thrust.

Table 6: Converged reference sortie: velocity increment and propellant of each leg.

Leg	Δv [m/s]	Prop. [kg]	Duration
NRHO to LLO	686.0	6,586	24.0 h
DOI	19.5	172	impulsive
Powered descent	1863.2	13,428	5.0 min
braking	1365.0	10,373	187 s
approach	382.3	2,374	45 s
soft landing	116.1	681	70 s
Powered ascent	1791.3	7,873	3.9 min
Circularisation	19.5	63	impulsive
LLO to NRHO	1308.7	3,466	24.0 h
Mission	5688.2	31,589	6.41 d

Figure 4 and Table 6 describe the converged mission. The lander leaves the Gateway with 45.7 t and touches down with 25.6 t. The braking phase takes almost three-quarters of the powered descent budget; the soft landing adds 116 m/s in 70 s, the price of the near-hover terminal descent. The ascent stage leaves the surface with 18.1 t. After a 72 h stay, the Gateway is far from the cheap return geometry, so the return waits the full 24 h of extra loiter and still needs 1309 m/s, against 686 m/s for the outbound leg. The crew therefore spends 6.41 days on board for 3 days on the surface.

Table 7 gives the converged mass budget. The dry mass is 7.14 t, of which the ECLSS is 1.12 t. Replacing the ECLSS module with its MER, which scales only with surface stay, yields 739 kg and a 2.5 t lighter lander: the MER ignores the 3.4 days the crew spends in transfers and loiters, which the NRHO architecture adds to every sortie. The detailed modules then describe the rest of the vehicle. The bottom-up propulsion hardware is 62% heavier than the MER on the descent stage and 43% on the ascent stage; the gap includes the growth allowance, the pressurant bottles and the feed system that the module books explicitly, and it is the next item to reconcile. The envelope module places the ascent stage below an 18.6 m^3 pressurised cabin and four fuel and two oxidiser tanks in a ring with two cargo bays; the stowed lander, 7.1 m in diameter and 10.6 m tall, fits the 8.4 m fairing of SLS Block 1B. The ECLSS draws 372 W, the avionics 285 W and the communication link to the Gateway 0.6 W with a 1.7 m dish.

Table 7: Converged mass budget of the reference sortie. The loop uses the MER column; the detailed column is the bottom-up sizing of the converged vehicle.

Item	Loop [kg]	Detailed [kg]
Structure and TPS	1,773	-
Descent propulsion	1,429	2,309
Ascent propulsion	1,150	1,639
Power	710	-
Avionics and communication	247	65
ECLSS (module, $t_c = 6.41$ d)	1,118	1,118
Other	714	-
Dry mass	7,141	
descent / ascent stage	3,089 / 4,052	
Payload down / up	3,500 / 834	
Usable propellant desc. / asc.	20,186 / 11,403	
RCS and reserves	2,681	
Wet mass at NRHO departure	45,745	

4.4 Sensitivity to the surface stay

Figure 5 repeats the loop for surface stays from 1.5 to 7 days; each point is a converged vehicle. The return transfer, in panel (a), dominates the trend. The Gateway completes one NRHO revolution in about 6.6 days, so stays between 5.5 and 6 days let the lander leave

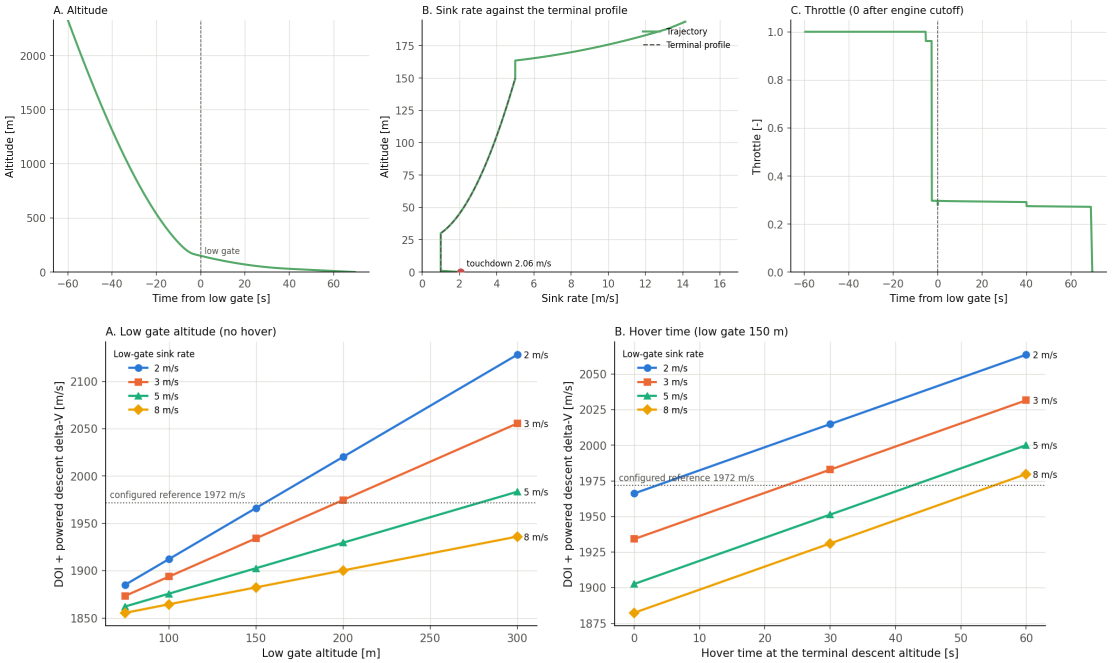


Fig. 3: Terminal descent of the NASA HLS preset. Top: soft landing in guidance mode, 116 m/s in 70 s from the low gate, with (A) altitude, (B) sink rate against the terminal profile and (C) throttle. Bottom: DOI and powered descent velocity increment against (A) the low-gate altitude without hover and (B) the hover time with a 150 m gate, for four low-gate sink rates.

LLO when the Gateway is back in the outbound-transfer geometry; the return mirrors it at 686 m/s, while the 2 to 3 day stays need 1300 to 1440 m/s even with 24 h of extra loiter. The ECLSS mass in panel (b) grows only from 970 to 1349 kg over the range; with open-loop consumables, four crew need about 44 kg per crewed day and 24 kg of EVA water per surface day.

The converged lander, in panel (c), is lightest at 4 to 4.5 days, 41.1 t, and 10% lighter than the 72 h reference. Beyond 6 days, the return geometry degrades again. The dry mass also grows by 2.1 t between 3 and 6 days: the ECLSS accounts for 171 kg of it, and the MERs of structure, power, avionics and descent propulsion, which grow with the surface stay and with the vehicle, account for the rest. The ascent stage follows the return transfer because it flies it. For an NRHO-staged sortie, mission duration therefore sizes the lander through the return geometry and the stay-dependent subsystems more than through life-support consumables. A 72 h stay, which looks cheap in consumables, is one of the most expensive in propellant. The 2 day case needed eight iterations because the return optimiser alternated between two local solutions with and without extra loiter, a reminder that the loop inherits the discrete choices of the mission model.

4.5 Regeneration technology

Figure 6(a) compares four carbon dioxide removal technologies of the ECLSS database for four crew: LiOH canisters, which consume 1.75 kg per crew day, and three regenerative options that trade expendables for hardware and power: electrochemical depolarised concentration (EDC), solid amine water desorption (SAWD) and the four bed molecular sieve (4BMS). With free power, EDC breaks even with LiOH after 6.3 days of crewed duration, SAWD after 9.8 days and 4BMS after 22.3 days. The reference sortie, at 6.41 days, sits at the break-even of the lightest regenerative option. When a battery stores the extra energy at 120 Wh/kg, the regenerative options never break even because their power penalty grows with duration faster than LiOH consumption.

Table 8: Carbon dioxide removal technology in the loop, reference sortie ($t_c = 6.41$ d). Δm_0 is the change in the converged wet mass with respect to LiOH.

Technology	P [W]	m_{ECLSS} [kg]	Δm_0 free [kg]	Δm_0 battery [kg]
LiOH	372	1,118	0	0
EDC	614	1,117	−4	+2,441
SAWD	965	1,146	+208	+6,893
4BMS	1,560	1,251	+1,008	+18,015

Table 8 carries the four options through the loop. Every kilogram of ECLSS on the ascent stage costs about 8 kg at NRHO departure, because the ascent stage flies 3.1 km/s with storable propellants and the descent stage carries it: the 310 kg battery of the EDC option adds 2.4 t to the lander, and the 4BMS with its battery also changes the engine allocation of both stages. For short sorties, regeneration saves nothing even with free power, and it costs several tonnes when

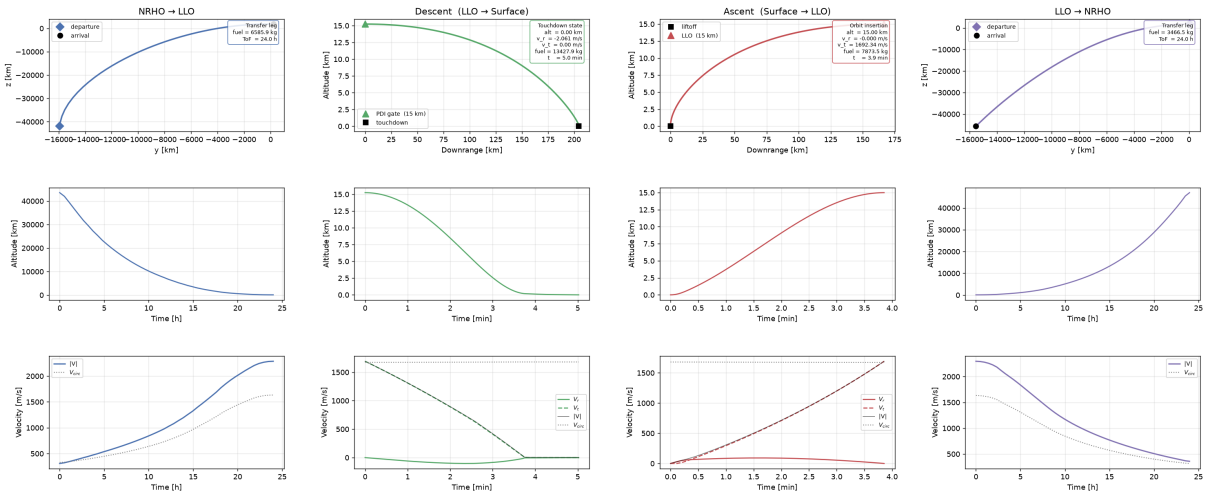


Fig. 4: Converged reference sortie. From left to right: outbound transfer in the Moon-centred synodic frame, powered descent, powered ascent and return transfer. Rows: path, altitude and velocity against time.

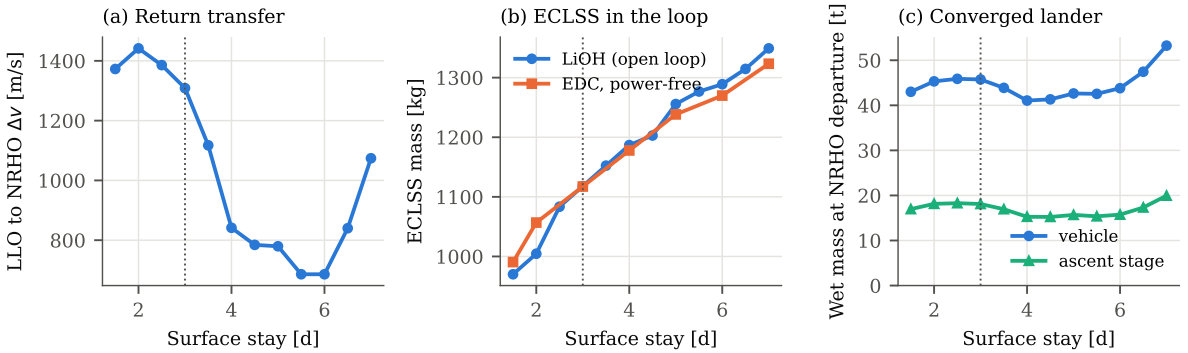


Fig. 5: Converged designs against the surface stay: (a) velocity increment of the return transfer, 24 h cap and up to 24 h of extra loiter; (b) ECLSS mass in the loop for open-loop LiOH and for EDC with free power; (c) wet mass of the lander and of the ascent stage at NRHO departure. The dotted line marks the 72 h reference.

the lander stores energy. With free power, EDC starts to save mass after 6.3 days, and SAWD after 9.8 days, so regeneration becomes worthwhile only for longer sorties with a power source that does not scale with time, such as surface solar arrays.

4.6 Cost and technology roadmap

HyCost evaluates the converged vehicle with the AMCM settings of the first iDREAM version (Eq. (32)). The dry mass of 7.14 t gives a development and production cost of 18.9 billion US$ in fiscal year 2020 for three units. The first version estimated 18.8 billion US$ with the AMCM and 18.6 billion US$ with the analogy method for a 7.1 t lander, and the NASA Office of Inspector General reported 21.3 billion US$ for the HLS [3]. The difficulty factor dominates the uncertainty: one step of difficulty up or down, from average to high or low, moves the estimate to 29.6 or 12.0 billion US$, far more than any mass change in the trades above.

TRIS evaluates the technologies behind the ECLSS choices with the database and the stakeholder criteria of the first version, which roadmapped the life support technologies of the LSAM case [4]. Table 9 lists the technologies that the design choices bring in, with their TRL and CaC, and applies Eq. (33) with a time at completion of 23.6 years, the pace of the first version roadmap, and a programme window of 8 years, from a roadmap start in 2027 to a first landing in 2035. Figure 6(b) generalises the check. Because the transit from TRL 6 to TRL 7 takes 40% of the time at completion, a technology must already be at TRL 7 to fit an 8 year window for any time at completion up to 26.7 years; TRL 6 fits only if the whole development takes less than 11.4 years.

The regenerative group adds 265 M€ of remaining CaC, less than 2% of the lander cost, but none of its technologies reaches TRL 7, so none fits the window: each would delay the first landing by 12 to 13 years, or 385 to 416 M€ at the delay cost of 32 M€ per year of the first version. The roadmap therefore rejects the regenerative closure, and the loop falls back to LiOH and stored consumables, which are at TRL 9. For the short sortie this fallback costs no mass (Table 8);

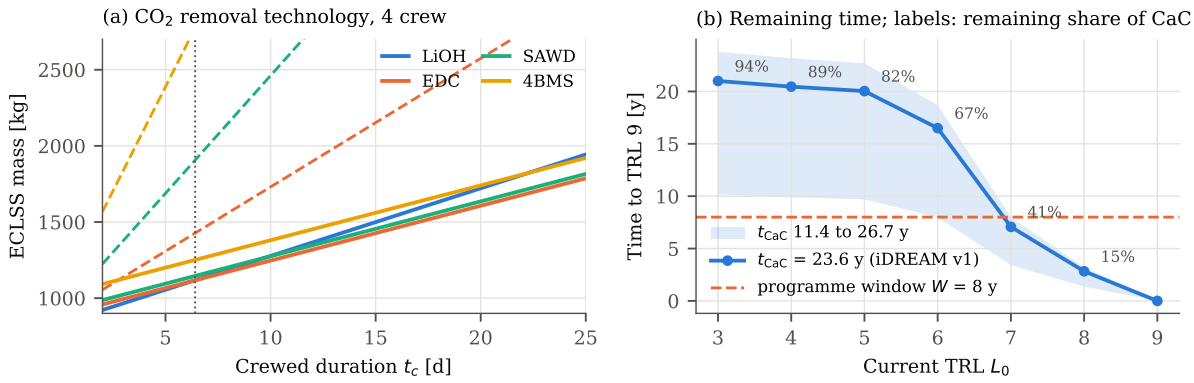


Fig. 6: (a) ECLSS mass against the crewed duration for four carbon dioxide removal technologies, four crew and 3 days on the surface; solid lines with free power, dashed lines with the extra energy stored in batteries at 120 Wh/kg. The dotted line marks the reference sortie. (b) Time to reach TRL 9 from the current TRL with the transit shares of Table 1; the band spans the times at completion for which the TRL 7 and TRL 6 thresholds hold, and the labels give the share of the CaC still to spend.

Table 9: Technologies behind the ECLSS choices: TRL and CaC of the first iDREAM database [4], remaining cost and time from Eq. (33) with $t_{\mathrm{CaC}} = 23.6$ y.

Technology	TRL	CaC [M€]	C_{rem} [M€]	t_{rem} [y]
Needed only by a regenerative ECLSS				
Closed loop air revitalisation	3	150	141.0	21.0
Oxygen recovery	3	35	32.9	21.0
Water recycling system	5	60	49.2	20.0
Waste processing	3	45	42.3	21.0
Needed by any crewed surface sortie				
EVA suit for lunar exploration	4	85	75.7	20.5
Medical and physiological monitoring	3	22	20.7	21.0
Radiation monitoring	5	10	8.2	20.0

a longer sortie that needs regeneration to close its mass would find the roadmap, not the ECLSS mass, as the binding constraint. The crew technologies stay on the critical path whatever the ECLSS choice. Among them, the lunar EVA suit, at TRL 4, has by far the highest remaining cost, 75.7 M€, so it is the first development priority for the programme schedule, followed by crew health and radiation monitoring. In cost terms, technology development is small relative to the vehicle; in schedule terms, it determines whether the vehicle can fly within the window.

5 Discussion and model applicability

The selected model fidelity also sets the scope of its conclusions. The CR3BP neglects solar perturbations, the eccentric motion of the primaries and a full ephemeris environment. The transfers use two impulsive burns and a polar LLO, and the optimiser finds local solutions. The planar powered model omits lateral motion, detailed landing site geometry, terrain clearance, navigation uncertainty and attitude control.

On the sizing side, the MERs come from a small population of landers, and the stay-dependent power and avionics terms drive part of the trend in Fig. 5(c); the EPS module, still in development, will replace them. The battery penalty of Eq. (27) stands in for that module in the regeneration trade. The cost model depends on subjective difficulty, and the roadmap relies on TRL values from the first iDREAM database, which predate recent flight experience with some life support hardware. These limits affect the absolute numbers more than the trends: the return geometry, the crewed duration and the maturity threshold at TRL 7 follow from the structure of the problem.

6 Conclusions

This paper presents two enhancements to the iDREAM 2.0 lander routine. First, the vehicle sizing of ASTRID-H and reduced-order mission models run in one closed loop. Optimal two-burn transfers in the CR3BP meet the Gateway exactly for any surface stay; optimal powered descent and ascent follow the operational constraints of Apollo and Altair; and a soft-landing routine flies the terminal descent. The loop exchanges velocity increments, propellant, installed thrust and crewed duration, and it converged in three iterations to 0.02% for the reference sortie. Second, the sizing works at the subsystem level: an ECLSS module inside the loop sizes each life support function to the crewed duration, and propulsion, envelope, communication, and avionics modules detail the converged vehicle. The mission models agree with Apollo, LSAM, and NASA HLS budgets within 9%, and the sizing modules agree with the reference vehicles within 12%, except for Apollo descent propulsion. On the NASA HLS sortie, the mission analysis shows that a 12 h transfer costs only 13 m/s more than a 24 h one, that the cheapest returns mirror the outbound leg when the Gateway comes back to the mirror geometry, and that the low gate and hover assumptions of the terminal descent explain most of the gap to the NASA descent

budget, because the near-hover flight pays the lunar gravity for every second it lasts.

For a four-crew, 72 h sortie from the Gateway to the lunar south pole, the loop converges to a 45.7 t lander with a 7.1 t dry mass and 1.1 t of ECLSS. Three findings follow. The surface stay sizes the lander through the return geometry: a 4 day stay makes the lander 10% lighter than the 72 h reference, and an ECLSS MER based on the surface stay alone underestimates life support by a third. For crewed durations below about a week, regenerative carbon dioxide removal saves no mass, and with battery power it adds tonnes. Finally, the cost and roadmap routines of the first iDREAM version show that technology maturation is small in cost but decisive in schedule: technologies below TRL 7 cannot mature within an 8 year programme, the loop falls back to open-loop life support at no mass penalty, and the lunar EVA suit becomes the first development priority.

Next steps are to bring the EPS and the propulsion module into the loop, update the TRL and CaC of the life support technologies, extend transfers to non-polar sites, and connect the preliminary trajectories to ASTOS for detailed descent and ascent.

References

[1] N. Viola, R. Fusaro, G. Narducci, G. Governale, J. Rimani, D. Ferretto, S. Chiusano, A. Fiori, O. Cohen and E. Dupuy, *iDREAM: a multidisciplinary methodology and integrated toolset for flight vehicle engineering*, 73rd International Astronautical Congress, Paris, 2022, IAC-22,D1,4A,12,x73441.

[2] NASA, *NASA Systems Engineering Handbook*, NASA/SP-2016-6105 Rev. 2, 2016. NASA Technical Reports Server, record 20170001761.

[3] G. Narducci, D. Chirulli, V. Vercella, G. Governale, J. Rimani, D. Ferretto and R. Fusaro, *iDREAM Final Report*, ESA Contract 4000132573/20/NL/BJ/rk, Politecnico di Torino, July 2022.

[4] G. Narducci, D. Chirulli, V. Vercella, G. Governale, J. Rimani, D. Ferretto and R. Fusaro, *iDREAM Technical Note 3*, ESA Contract 4000132573/20/NL/BJ/rk, Politecnico di Torino, June 2022.

[5] R. Shishko, *Developing analogy cost estimates for space missions*, AIAA Space 2004 Conference and Exposition, 2004.

[6] A. Owens and O. de Weck, *Sensitivity analysis of the Advanced Missions Cost Model*, 46th International Conference on Environmental Systems, 2016.

[7] N. Viola, R. Fusaro, V. Vercella and G. Saccoccia, *Technology RoadmappIng Strategy, TRIS: methodology and tool for technology roadmaps for hypersonic and re-entry space transportation systems*, Acta Astronautica 170 (2020) 609–622.

[8] G. Narducci, J. Rimani, D. Ferretto, R. Fusaro and N. Viola, *iDREAM Maintenance, Technical Note*, ESA Contract SC 1300068123, Politecnico di Torino, November 2022.

[9] M. Isaji, I. Maynard and B. Chudoba, *A new sizing methodology for lunar surface access systems*, AIAA SciTech 2020 Forum, 2020, doi: 10.2514/6.2020-1775.

[10] B. Mahajan and G. L. Condon, *End-to-End Performance Optimization of a Crewed Lunar Landing Mission Staged from a Near Rectilinear Halo Orbit*, AAS 22-063, 2022. NASA Technical Reports Server, record 20220001869.

[11] L. D. Kos, T. P. Polsgrove et al., *Altair descent and ascent reference trajectory design and initial dispersion analyses*, AIAA Guidance, Navigation, and Control Conference, 2010. NASA Technical Reports Server, record 20100035768.

[12] F. V. Bennett, *Apollo Experience Report: Mission Planning for Lunar Module Descent and Ascent*, NASA TN D-6846, June 1972. NASA Technical Reports Server, record 19720018205.

[13] C. P. Newman, D. C. Davis, R. J. Whitley, J. R. Guinn and M. S. Ryne, *Stationkeeping, Orbit Determination, and Attitude Control for Spacecraft in Near Rectilinear Halo Orbits*, AAS 18-388, NASA report JSC-E-DAA-TN59836, 2018. NASA Technical Reports Server, record 20180006800.

[14] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings and M. Diehl, *CasADi: a software framework for nonlinear optimization and optimal control*, Mathematical Programming Computation 11 (2019) 1–36.

[15] A. Wächter and L. T. Biegler, *On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming*, Mathematical Programming 106 (2006) 25–57.

[16] M. K. Ewert, T. T. Chen and C. D. Powell (eds.), *Life Support Baseline Values and Assumptions Document*, NASA/TP-2015-218570/REV2, February 2022. NASA Technical Reports Server, record 20210024855.

[17] P. Eckart, *Spaceflight Life Support and Biospherics*, Space Technology Library, vol. 5, Microcosm Press and Kluwer Academic Publishers, 1996.

[18] W. J. Larson and L. K. Pranke (eds.), *Human Spaceflight: Mission Analysis and Design*, McGraw-Hill, 1999.

[19] NASA, *Exploration Systems Architecture Study, Final Report*, NASA-TM-2005-214062, 2005.

[20] B. Mahajan and G. L. Condon, *HLS NRHO to lunar surface and back mission design*, presentation, NASA Johnson Space Center, August 2022. NASA Technical Reports Server, record 20230002566.

[21] R. Whitley and R. Martinez, Options for staging orbits in cis-lunar space, IEEE Aerospace Conference, Big Sky, MT, 2016.