Optimal Design of High-Mass Systems for Entry Descent and Landing on Mars

Executive Summary of the Thesis
Laurea Magistrale in Space Engineering - Ingegneria Spaziale

Author: **Parazzi Luca**

Advisor: Prof. Michèle Roberta Lavagna
Co-advisors: Ing. Fabio Calantropio
Academic Year: 2024-2025

1. Problem Statement

Exploring the Red Planet has been one of the greatest aspirations since the beginning of the human adventure in space. For over five decades, robotic missions have explored the surface of Mars and have established a robust and reliable foundation for the Entry, Descent and Landing (EDL) architectures. However, these heritage solutions have reached their physical limits in terms of achievable landing mass. The Martian atmosphere is too thin to provide sufficient aerodynamic deceleration for heavy vehicles; therefore, a larger entry surface is needed to reduce the ballistic coefficient. However, this required increase in entry area conflicts with the maximum allowable diameter imposed by the launch vehicle fairing, resulting in a critical scalability bottleneck for the heritage architectures. Therefore, as mission requirements evolve towards more ambitious payload masses, which are necessary to sustain human presence on the Martian surface, new EDL technologies must be investigated to increase the mission cargo capability of future missions to Mars. In fact, the *Moon to Mars Architecture* definition document produced by NASA [1] identifies 75 metric tons as the minimum payload capability for a short-duration crewed mission, almost two orders of magnitude greater than the heaviest mass ever landed on Mars (Perseverance, 1 ton), underscoring the necessity for a paradigm shift.

2. Objective of the Thesis

The present thesis explores different solutions aimed at bridging the gap between the heritage architectures and human-class mission requirements. The primary objective is to investigate whether novel combinations of aerodynamic and propulsive deceleration technologies can enable a substantial increase in the delivered payload mass.

To this end, three architectures are proposed, each solving the Entry, Descent and Landing phases differently. The first, Architecture A, employs a Hypersonic Inflatable Aerodynamic Decelerator (HIAD) to increase the entry area of a seventy-degree sphere-cone capsule, while relying on the heritage Disk-Gap-Band (DGB) parachute for the descent phase. Architecture B shares the same entry technology but replaces the subsequent unguided parachute phase with supersonic retropropulsion combined with powered descent guidance. Architecture C, instead, adopts a different capsule geometry, and in particular a biconic capsule while relying on a Supersonic Inflatable Aerodynamic Decelerator (SIAD) for the descent phase. The landing phase is common to all the architectures and relies on powered descent guidance and on a sky-crane maneuver to safely deliver the payload, avoiding the severe plume-surface interaction effects. Figure 1 presents the Concept of Operations (ConOps) for the three investigated architectures.

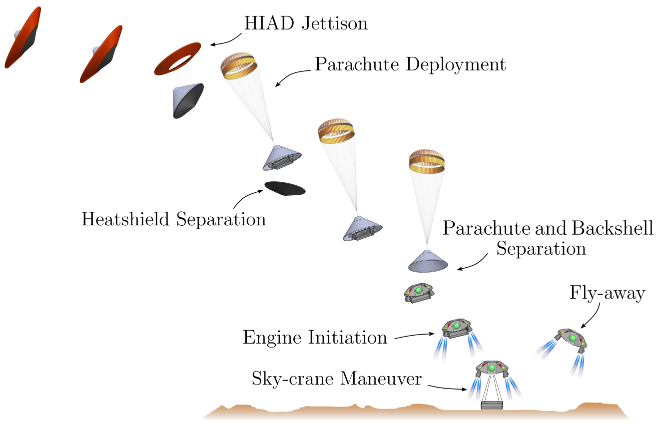

(a) Architecture A

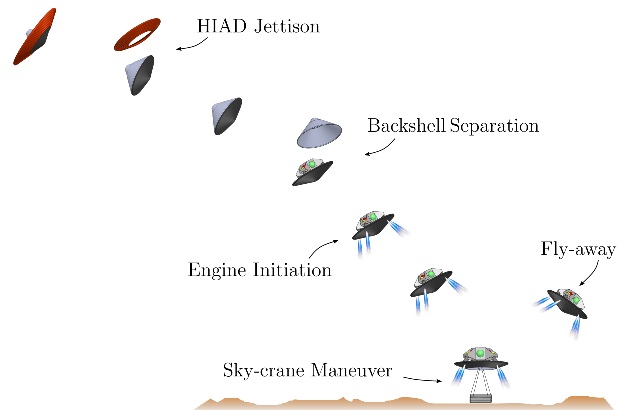

(b) Architecture B

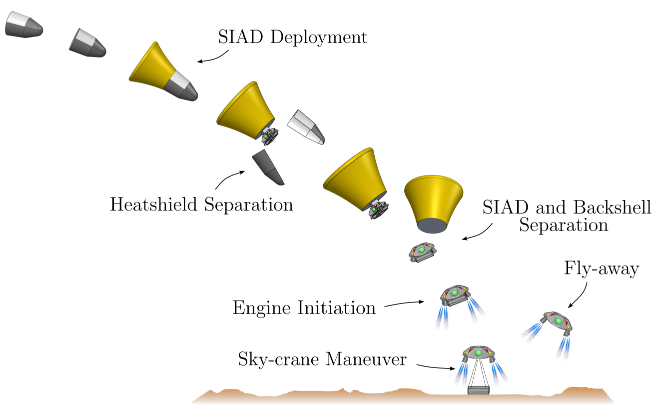

(c) Architecture C

Figure 1. Concept of Operations (ConOps) for the three architectures proposed, with indication of the main events.

To assess the potential of each architecture, a comprehensive analysis is carried out at a preliminary design level. In fact, the intent of the study is not to provide a detailed system design, but rather to explore the architectural trade space and identify configurations able to exceed the payload capability of heritage systems. Moreover, a multiobjective optimization framework is developed to properly compare the performance of the proposed designs.

3. Modelling Approach

A preliminary design study is based on two main principles that govern its modelling approach: the adopted models should be reliable and accurate enough to capture the main aspects of the problem considered, while avoiding costly high-fidelity computations.

3.1. Physical Modelling

The design of an Entry, Descent, and Landing architecture is strongly influenced by the atmospheric model adopted. For this purpose, the Mars Climate Database (MCD) [2] is used to retrieve the main properties of the Martian atmosphere. This aspect is particularly critical during the entry phase, which is highly sensitive to atmospheric prediction accuracy as the trajectory is controlled only aerodynamically through bank-angle modulation. The guidance strategy must steer the vehicle toward the target while satisfying the imposed constraints on acceleration and heat flux. In particular, due to the prohibitive cost of high-fidelity CFD simulations within a preliminary design framework, the aerodynamic heating is estimated through an approximate correlation, namely the Sutton-Graves equation.

The entry trajectory is modelled using a three-degree-of-freedom (3-DOF) dynamics expressed in corotating spherical coordinates. For the landing phase, which relies on a three-axis thruster control, a 6-DOF model is adopted to include the attitude dynamics of the vehicle. This is formulated in a Cartesian reference frame (East-North-Up) centered on the target to ease the formulation of the guidance problem.

3.2. System Design

The three architectures present different key aspects to be correctly modelled in order to provide a meaningful assessment of their potential. An overview of the model adopted for the different phases is presented in Table 1.

The aerodynamics mainly depend on the geometry of the system. Because CFD simulations are computationally prohibitive in a preliminary design, existing aerodynamic databases are leveraged whenever possible. This can be achieved, for example, by using reference entry capsules available in the literature and scaling them to match the maximum allowed diameter (8 m) while preserving the geometric ratios and, consequently, the aerodynamic coefficients. When the databases are found incomplete in the hypersonic region, the modified Newtonian theory is employed to approximate the needed coefficients. In cases involving complex geometries or specific operating conditions,

Phase	Feature	Architecture A	Architecture B	Architecture C
Entry	Vehicle Aerodyn.	Blunt body	Blunt body	Bluff biconic
	Decelerators	HIAD	HIAD	-
	Guidance	Bank angle	Bank angle	Bank angle
	Dynamics	3DOF spherical	3DOF spherical	3DOF spherical
Descent	Vehicle Aerodyn.	Flat plate	-	-
	Decelerators	Parachute	SRP	SIAD
	Guidance	Unguided	Powered Descent	Unguided
	Dynamics	3DOF cartesian	6DOF cartesian	3DOF cartesian
Landing	Vehicle Aerodyn.	-	-	-
	Decelerators	Thrusters	Thrusters	Thrusters
	Guidance	Powered Descent	Powered Descent	Powered Descent
	Dynamics	6DOF cartesian	6DOF cartesian	6DOF cartesian
	Landing Tech.	Sky-crane	Sky-crane	Sky-crane

Table 1. Features of each architecture for the different EDL phases. This includes vehicle aerodynamic approximation, decelerator employed (both aerodynamic and propulsive), guidance system implemented, and the set of dynamics equations used. When the thrusters are firing, the aerodynamic forces are considered negligible.

such as during supersonic retropropulsion, appropriate assumptions are introduced to simplify the problem while maintaining the required accuracy.

Mass estimation is performed using a *size-as-you-fly* methodology, with the objective of coupling trajectory design and system sizing. Approximate mass relationships for the main components derived from heritage missions are adopted, and their dependence on key dynamic parameters enables the system to be sized consistently with the trajectory. Special attention is devoted to the modelling of the inflatable devices considered in this thesis, namely HIAD and SIAD. Since these technologies are characterized by low maturity levels and only a few models have been constructed, the sizing is intentionally kept at a preliminary level using dimensionless quantities to scale with geometry parameters and environmental conditions.

As can be noted in Figure 1, multiple events characterized the EDL trajectory of the three architectures. The complexity of the operations is taken into account by introducing free-fall periods during separation transients, namely HIAD jettison, SIAD and parachute separation. This introduces a margin to reduce the risk of recontact, allowing for seamless operations.

4. Optimization

The modelling framework is integrated within the trajectory generation process. To this end, the Entry, Descent, and Landing trajectory is solved phase by phase, while an external optimization layer couples the different parts by tuning the interface conditions and the key geometric parameters. This optimization framework, which is graphically represented in Figure 2, enables the generation of various feasible designs allowing for a thorough comparison of different configurations.

4.1. Trajectory Optimization

The EDL sequence begins at the Mars atmospheric interface, at approximately 125 km altitude, with an entry velocity that may vary depending on the interplanetary trajectory followed and on the presence (or absence) of an aerocapture phase. Here, the dominant deceleration mechanism is the aerodynamic drag. While early Mars missions relied on a fully ballistic entry, a higher precision landing capability translates to the need for lift modulation to steer the vehicle during atmospheric flight. In this work, the corresponding entry guidance is formulated as an optimal control problem. This is solved using a direct collocation approach, specifically a fifth-order Gauss-Lobatto transcription, which provides a numerically efficient and robust method to compute the optimal bank-angle profile while satisfying the imposed constraints on decelerations and peak heat flux.

Depending on the architecture, the entry phase may be followed by an aerodynamic descent phase with the deployment of a parachute or a supersonic inflatable decelerator. This phase is unguided; therefore, it is not necessary to solve any optimal control problem, and the states of the system are propagated through the integration of the equations of motion (3-DOF) by dedicated solvers in MATLAB, such as `ode45`.

For all architectures, the final phase consists of a powered descent maneuver, aimed at decelerating the lander to a condition in which the sky-crane mechanism can safely operate. To do so, four gimbaled retrorockets are employed. As in the entry phase, this translates into an optimal control problem, which is solved by the state-of-the-art algorithm for powered descent guidance: successive convexification. In fact, given the 14-state dynamics and the three-axis control on each engine, convex optimization allows the optimal control problem to be solved in polynomial time with respect to the problem size. This involves multiple steps, such as time normalization, convexification of the nonlinear dynamics and constraints, discretization of the continuous-time formulation, and handling of artificial unboundedness and infeasibility, to transcribe the original control problem into a Second-Order Cone Programming (SOCP) problem. This is then solved through successive iterations of the aforementioned passages until the convergence criteria are satisfied.

4.2. Global Optimization

The aforementioned optimization processes are solved independently of each other, and no state information is exchanged. This parallelization, which may be exploited for computational efficiency, decouples the different phases but requires an external optimization framework. The latter merges the two internal optimizations to construct a dynamically feasible trajectory leveraging both geometric dimensions and interface conditions as design variables. Since the ultimate goal of an EDL architecture is to maximize the delivered payload mass, a global trajectory optimization is necessary. However, a single-objective formulation is of limited use in a preliminary design, where trade-offs are often necessary to account for technological constraints.

Given that this study includes novel deceleration technologies whose future availability and maturity remain uncertain, to not restrain the study to a feasible but not reachable design (at least in the near-future term), the dimensions of such decelerators are included in the objectives space. Accordingly, three multiobjective optimization problems are defined, one for each architecture. The optimization workflow is designed to be modular, thereby minimizing the changes required to tailor the optimization problem to each architecture.

The resulting optimization problems can be classified as Expensive Constrained Multi-objective Optimization Problems (ECMOPs), since they include two nested optimizations, which are expensive to solve, and different sets of constraints. A novel approach for this type of problem is employed in this work, which combines the use of surrogate assisted evolutionary algorithm with deep reinforcement learning (i.e., DRL-SAEA [3]). This is based on a DRL agent that selects between three different Kriging surrogate models for approximating the original expensive objective function and constraint violation function, with the aim of constructing a more accurate approximation of the feasible search space and of the Pareto front, within a limited budget for function evaluations.

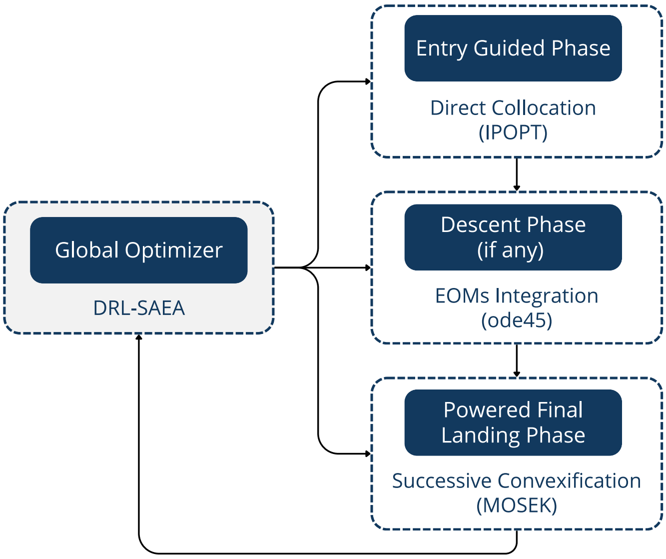

Figure 2. Optimization framework adopted in this work. For each component, the corresponding method and solver employed are indicated.

5. Results

A single run of the whole optimization procedure has been found insufficient to obtain a well-populated Pareto front, given the complexity of the optimization problem considered. Accordingly, a numerical procedure is developed to combine unseeded and seeded runs to construct a unique and complete set of non-dominated solutions. In particular, five unseeded runs are first employed to provide a sufficient spread of the solutions in the objective space. Then, the non-dominated set computed by considering all five runs together is used as the initial guess for two additional seeded runs, with the aim of confirming the Pareto front found. Although no detailed analysis was carried out on the necessary number of runs to formally achieve convergence, the resulting approximation of the Pareto front is considered sufficient for the scope of this work, enabling, with limited computational resources, a consistent comparison between the different architectures.

For the sake of completeness and to provide an example of the optimization process outcome, Figure 3 shows the Pareto front obtained for Architecture A, after the seven aforementioned runs. Similar results have been obtained for the other two architectures, except that the optimization problems they address are bi-objective due to the absence of a parachute.

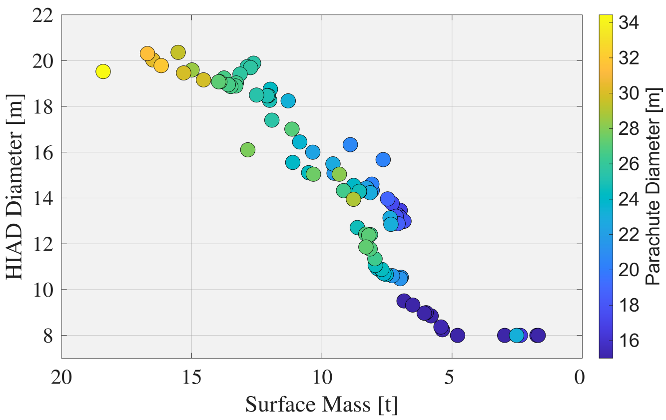

Figure 3. Pareto front obtained for Architecture A in the two-dimensional objective space. The third dimension is included through the colorbar.

To ensure meaningful insights into the results obtained, a limited number of representative solutions is selected from each Pareto front. In particular, the best point for each objective and the knee solution are selected as significant trade-off configurations. For each selected point, the corresponding EDL trajectory is reconstructed and analyzed in terms of state evolution, control history, and main trajectory parameters to both verify the correctness of each solution and identify the limits and advantages of each architecture.

In any case, all three architectures exhibit substantially higher payload capability than the past missions to the Red Planet, as highlighted in Figure 4.

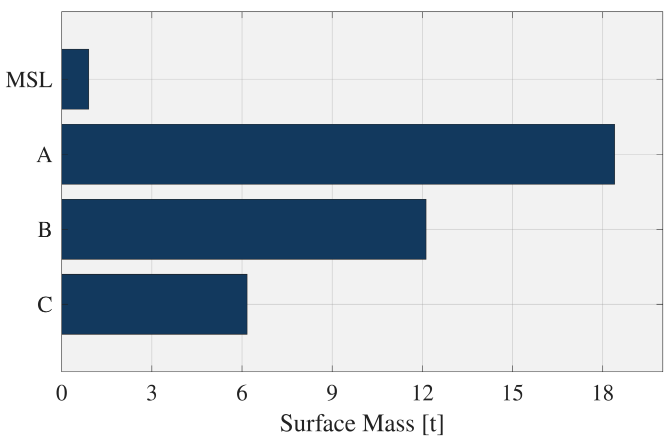

Figure 4. Surface mass of each architecture compared to the rover mass of MSL (899 kg).

The most promising solution, namely Architecture A, notably outperformed the baseline considered (Mars Science Laboratory), with a payload of about 18.4 metric tons, corresponding to an entry mass of 75 tons. The other architectures show more limited capability, with Architecture B able to land about 12 tons, due to the higher propellant consumption, and Architecture C capable of 6.2 tons, constrained by the absence of any inflatable decelerator during the entry phase.

Moreover, for each solution, the key technological drivers are identified, and their maturity and near-term feasibility are discussed. This brief assessment complements the achievable payload capability and highlights the most feasible designs to be considered for near-term high-mass missions to Mars. It is found that the best trade-off between maturity and payload capacity is a specific version of Architecture A, with a limited parachute diameter.

6. Conclusions

This thesis demonstrates that innovative Entry, Descent, and Landing architectures can significantly increase payload capability for future missions to the Red Planet (up to approximately 18 tons) compared with heritage solutions (maximum of 1 ton). Through an optimization-based comparison of three candidate architectures, the analysis highlights both opportunities and key trade-offs associated with advanced deceleration technologies for high-mass landing applications.

This work, although at a preliminary design level, offers a structured reference for trade-off studies and a baseline for future refinements and higher-fidelity assessments.

References

[1] NASA. *Moon to Mars Architecture - Revision B*. Architecture Definition Document. 2024. URL: https://www.nasa.gov/moontomarsarchitecture-architecturedefinitiondocuments/.

[2] Ehouarn Millour et al. *The Mars Climate Database, MCD version 6.1*. Tech. rep. EPSC2024-516. Copernicus Meetings, 2024. DOI: 10.5194/epsc2024-516.

[3] Shuai Shao, Ye Tian, and Yajie Zhang. "Deep reinforcement learning assisted surrogate model management for expensive constrained multi-objective optimization". en. In: *Swarm and Evolutionary Computation* (2025). DOI: 10.1016/j.swevo.2024.101817.