



POLITECNICO
MILANO 1863

**SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE**

EXECUTIVE SUMMARY OF THE THESIS

Early-Stage Design and Trade-off Analysis of Hyper-redundant Space Manipulators for On-Orbit Servicing

LAUREA MAGISTRALE IN SPACE ENGINEERING - INGEGNERIA SPAZIALE

Author: BEATRICE ALESSANDRA TUROTTI

Internal advisor: PROF. MAURO MASSARI

External advisor: PROF. ANTON DE RUITER

Co-advisor: DR. NICCOLÒ FARACO

Academic year: 2024-2025

1. Introduction

In-Orbit Servicing (IOS) is becoming a key capability for more sustainable and flexible space operations, including inspection, docking, refuelling, repair, relocation, and end-of-life disposal. Robotic manipulators enable controlled interaction with non-cooperative targets in constrained spacecraft environments. In this context, hyper-redundant manipulators (HRM) are particularly suitable: additional degrees of freedom (DoF) increase workspace and dexterity, improving obstacle avoidance and posture selection during proximity operations [1, 2].

However, these advantages also create design challenges, especially when the arm is mounted on a free-floating spacecraft. Additional joints or longer links typically increases mass and inertia, strengthening the dynamic coupling with the base. Therefore, arm motion can induce reaction forces and moments that disturb spacecraft attitude and increase attitude-control effort (e.g., thruster activity or reaction-wheel momentum management). Higher redundancy also makes motion planning and control more complex, and structural compliance and vibration can limit performance beyond rigid kinematic

assumptions [1, 3].

This thesis supports early-stage architectural selection of hyper-redundant space manipulators in free-floating conditions. The objective is not to identify a single optimal design, but to provide a systematic framework to compare candidate architectures and understand how key design choices influence mission-level performance. The investigated choices include the number of joints, link-length distribution, and geometry-driven mass and inertia properties. The evaluation focuses on mission-relevant indicators such as base disturbance (reaction wrenches), actuator effort and energy-like consumption, peak power demand, reachability, and feasibility along representative task trajectories [3].

To achieve this goal, a modular simulation and optimisation toolchain is developed. Representative end-effector task trajectories are generated at pose level and discretised with consistent time scaling. Two spatial inverse kinematics solvers for generic hyper-redundant chains are implemented and compared, including safeguards near ill-conditioned Jacobians and redundancy exploitation through null-space objectives. The resulting joint trajectories are

evaluated through a microgravity-oriented inverse dynamics model based on a Recursive Newton–Euler Algorithm (RNEA), to extract base reaction histories and trajectory metrics. A multi-objective optimisation explores trade-offs, with single-objective refinement of selected candidates and a proof-of-concept geometry–trajectory co-design extension; representative cases are reproduced in Simscape Multibody and MSC Adams with joint-level PD tracking for cross-tool consistency [3].

2. Problem formulation and methodological approach

A design scenario is defined through a compact set of high-level inputs that represent the mission context and constraints. The goal is to convert these inputs into feasible candidate architectures and to map the resulting trade-offs among conflicting performance drivers.

2.1. Scenario inputs, design variables and outputs

The input vector captures the main drivers that influence both feasibility and sizing. In this work, the key scenario descriptors include: the mission type (inspection, docking, refuelling, repair), orbit type (LEO or GEO), available power budget, target satellite mass and geometry, and an accuracy/precision requirement. From these elements, the toolchain derives baseline geometric scaling, mass budget ranges, and preliminary material and cross-section requirements.

Design variables are grouped to reflect how design decisions are made in practice. At the architectural level, the number of DoF controls redundancy and manipulability and strongly impacts overall mass and actuator count. At link level, geometry and mass properties are tuned through link-length scaling, defined by a scale factor s applied either globally or per link, cross-sectional sizing (with material density as an associated parameter), and a link-wise mass distribution model. The latter is controlled via a center-of-mass placement and a lumped non-structural distribution parameter, γ , which accounts for added hardware masses (e.g., actuators, harness, sensors and mounted components) and biases that would be allocated in a detailed design. At hardware level, torque/power ratings are selected to satisfy worst-case loads, while

joint stiffness and damping are used as simplified proxies for compliance to support tracking-oriented assessments. Finally, a limited set of trajectory-related parameters (total time, via points, smoothness and phase weights) is introduced, since the motion profile directly affects loads, base disturbance, and actuator sizing.

The main outputs of the toolchain are mission-relevant performance indicators that can be compared across candidates: reachability and feasibility along representative task trajectories, joint torque histories, peak power demand and energy-like proxies, base reaction forces and moments (used as measures of disturbance), and accuracy-related indicators. These outputs are used both to generate trade-off maps (e.g., heatmaps over design grids) and to support multi-objective comparisons.

2.2. Toolchain structure and modelling principles

Given a scenario definition, the workflow is implemented as a forward pipeline that converts mission inputs into comparable, trajectory-level metrics. First, representative end-effector trajectories are generated at pose level as a sequence of task frames with a mission-dependent multi-phase structure (approach, task execution, retreat). The approach/retreat phases include scenario-specific offsets (used to define a safe corridor and to avoid abrupt terminal conditions), while the execution phase can include a short hold segment for contact-like tasks (e.g., docking/refuelling) when required. All phases are discretised with consistent time scaling, so that candidates are evaluated under the same temporal resolution and comparable motion profiles.

Two IK methods are then used to compute the arm pose while the End-Effector (EE) follows these trajectories. In particular, Method A uses a Damped Least Squares formulation to track the desired end-effector twist while remaining stable near ill-conditioned Jacobians; redundancy is exploited through a null-space term to enforce secondary objectives (e.g., joint-limit avoidance, minimum joint motion, and reduction of base reaction moments). Method B adopts a weighted resolved-rate scheme that computes a minimum Q -weighted joint velocity to track the same twist; robustness is improved

through a guarded line-search on the integration step and optional randomized restarts when convergence is not achieved from the nominal initial guess. Both solvers use consistent stopping criteria and feasibility checks (e.g., tolerance on pose residuals, iteration limits, and joint-limit compliance) to ensure repeatable comparisons across designs.

For each candidate, the resulting joint trajectories are evaluated through a microgravity-oriented inverse dynamics model based on the RNEA. In this step, joint torques are computed along the discretised motion, and the base reaction wrench induced by the manipulator motion is extracted from the net wrench balance at the floating base. These outputs are reduced to trajectory-level indicators such as base reaction force/moment histories, mean base disturbance measures, peak actuator demands, and energy-like effort metrics obtained from discrete integration of joint power.

All modules are integrated in a modular MATLAB toolchain designed for repeated execution inside parameter sweeps and optimisation loops. The same evaluation pipeline is used both for systematic grid exploration and for multi-objective optimisation, with optional single-objective refinement of selected candidates.

3. Performance metrics and optimisation setup

A key aspect of early-stage manipulator design is the definition of performance indicators that are both mission-relevant and comparable across different architectures. In this work, each candidate design is evaluated along the same reference task trajectory (or sequence of task poses) and reduced to a compact set of scalar metrics. This enables systematic trade-off exploration through parametric sweeps (heatmaps) and through multi-objective optimisation.

3.1. Primary performance indicators

The optimisation runs adopted three main objectives, all formulated as minimisation goals:

$$f_1(\mathbf{x}) = \text{mean } M_b, \quad f_2(\mathbf{x}) = m_{\text{tot}}, \quad f_3(\mathbf{x}) = E \quad (1)$$

The first objective, mean M_b , is the mean magnitude of the base reaction moment computed along the discretised trajectory and used as a

compact indicator of base disturbance. The second objective, m_{tot} , is the total manipulator mass, computed consistently with the selected mass and inertia model. The third objective, E , is an energy-like integral based on the absolute joint mechanical power:

$$E = \int_0^T \sum_{i=1}^N |\tau_i(t) \dot{q}_i(t)| dt,$$

implemented numerically through discrete time integration using the stored time step Δt . This definition is simple and robust, and it remains directly comparable across different candidates because the same trajectory timing is used during the evaluation.

In addition to the three objectives in Eq. (1), the toolchain stores diagnostic outputs that are useful for engineering interpretation and for down-selection. These include peak joint torques, peak power demand, base reaction force histories, and feasibility flags (e.g., IK convergence and joint-limit compliance). These indicators are not always treated as explicit objectives, but they support screening and sanity checks during trade studies.

3.2. Constraints and feasibility safeguards

The optimisation problem is primarily bounded. Bounds are imposed directly on the decision variables (e.g., per-link scaling factors, mass-distribution parameters, and geometry-driven cross-section variables) to ensure that the solver explores physically meaningful and numerically stable configurations. Optional safeguards can also be enabled to avoid pathological candidates. In particular, a smoothness bound can limit excessive oscillations in the per-link scaling profile by constraining $|s_{i+1} - s_i|$, and simple geometric inequalities can be used to prevent unrealistic cross-sections under slender-link assumptions. Accuracy requirements are included in the evaluation through feasibility and penalty logic. Conceptually, the tracking error can be constrained by a precision threshold ϵ_r :

$$\max_{k=1,\dots,K} \|\mathbf{p}_{d,k} - \mathbf{p}_k\| \leq \epsilon_r,$$

with violations converted into a soft penalty to preserve solver robustness during exploration. Finally, before the IK stage, the task sequence

can be pre-checked through a reachability safeguard so that the downstream solver receives only reachable end-effector poses.

3.3. Trade-off exploration and optimisation workflow

The optimisation section couples the design variables with the full evaluation pipeline (IK \rightarrow dynamics \rightarrow metrics). Three complementary modes are used.

First, low-dimensional parametric sweeps are performed over the most influential continuous variables to visualise the cost landscape and identify trends and promising regions. These exploratory maps also help define meaningful bounds and provide physically consistent baseline candidates.

Second, when explicit trade-offs are required, a population-based multi-objective optimisation is executed with `gamultiobj` to approximate a non-dominated set in the objective space (m_{tot} , $\text{mean } \|\mathbf{M}_b\|$, E). To improve robustness, the initial population is generated within bounds and can be seeded with a baseline design and perturbed near-seed individuals.

Finally, after the multi-objective search, an optional local refinement stage is applied to a limited subset of Pareto solutions. The refinement uses `fmincon` on a weighted-sum scalarisation of the objectives. To avoid dominance due to unit scaling, the weighted sum is built using normalised objectives (each objective divided by an estimate of its range over the computed Pareto set). This hybrid step preserves the global exploration provided by `gamultiobj` while allowing local improvements around selected non-dominated candidates.

4. Results and discussion

This section summarises the main outcomes of the developed toolchain, focusing on the behaviour of the two IK solvers, the structure of the optimisation landscape and the resulting design trends, the multi-objective and co-design trade-offs, and the consistency checks performed with multibody simulations.

4.1. Inverse kinematics robustness and computational behaviour

The two IK approaches were first tested on fixed target poses for two reference scenarios: a

docking-like case (GEO-oriented, 10-DoF chain) and an inspection-like case (LEO-oriented, 8-DoF chain). In both scenarios, the solvers converged to feasible configurations with residual errors well below the accuracy bands adopted at mission level.

A consistent trend is observed in the accuracy-speed balance. Method A typically achieves smaller position residuals with a stable balance between translation and rotation errors, while Method B is generally faster in wall time while remaining within the adopted tolerance bands. Importantly, the two methods often return different joint configurations for the same end-effector target, which is expected for redundant chains and remains compatible with optimisation workflows as long as feasibility and numerical stability are preserved.

4.2. Cost-surface structure and optimisation outcomes

Grid-based exploration showed that the objective landscape is clearly non-convex, with smooth basins separated by ridges and transitions. This behaviour is physically justified because small geometry variations can move the arm closer to singular regions and can change inertia distributions that drive base reaction moments. At the same time, the landscape remains structured (i.e., not dominated by discontinuities), indicating that the pipeline is numerically stable and suitable for early-stage trade studies. When embedded in a population-based optimisation loop, the toolchain produced consistent improvements in the objective. In the reported representative runs, the best scalar objective values (corresponding to a knee-type compromise) were $J_{\text{best}} = 17.94$ for the docking-like case and $J_{\text{best}} = 12.05$ for the inspection-like case, with total runtimes on the order of 10^4 s. The docking-oriented scenario required longer evaluations, consistent with the higher DoF count and tighter disturbance/robustness requirements. The design outputs highlight repeatable trends. In both tests, the optimised inertia distribution parameter γ is negative, corresponding to designs that are heavier proximally and lighter toward distal links. This behaviour is especially pronounced for the inspection-like case, where low disturbance and agility are dominant drivers. Structural proxy parameters follow the

expected direction: the docking-like case converges to a stiffer, larger-section solution, while the inspection-like case favours a lighter and more compliant configuration, still consistent with the assumed tracking tolerances.

4.3. Multi-objective trade-offs and co-design effects

Multi-objective runs confirmed that the selected objectives are competing. In particular, pushing the solution toward very low base disturbance typically requires either increasing mass (to redistribute inertia and increase stiffness proxies) or accepting higher actuation effort, depending on how the motion is executed. The resulting non-dominated set provides a practical basis to identify knee regions, where a moderate mass increase yields a comparatively larger reduction in disturbance or energy-like cost.

The co-design proof of concept further showed that trajectory parameters contribute directly to performance. By allowing the optimiser to adjust a limited set of trajectory-level variables (such as total duration and commanded increments), the framework can reduce the trajectory-averaged base moment without relying only on geometry changes. This supports the key engineering idea that, for IOS operations, the best design is not universal but depends on how the arm is commanded, how phases are timed, and how conservative the terminal approach must be.

4.4. Multibody validation in Simscape Multibody and MSC Adams

Selected candidate designs were reproduced in multibody environments to verify that the MATLAB-level results remain meaningful beyond a purely kinematic view. In Simscape Multibody, closed-loop simulations with joint-level PD control showed smooth tracking, bounded velocity profiles, and stable sensor-based error behaviour, indicating that the controller–plant setup is well-posed for the chosen parameters.

A further consistency check was performed through an MSC Adams model and a Simulink–Adams co-simulation: while perfect overlap is not expected due to different numerical integration and constraint handling, the joint trajec-

tories remained close in shape and scale across tools, supporting the coherence of the kinematic and inertial modelling assumptions.

5. Conclusions and future developments

The main goal of the thesis was to support repeatable comparisons across candidate architectures of HRM and mission contexts with a modelling level that remains compatible with concept-phase iterations.

A first key outcome of the unified toolchain developed for this scope is a structured evaluation pipeline that produces comparable indicators for different designs when the same task trajectory is applied. This enables both grid-based mapping (to interpret sensitivities and cost-surface structure) and optimisation-based exploration (to identify non-dominated solutions and quantify trade-offs). Across the analysed cases, clear and repeatable trends emerge. Limiting the inertia of the outer links by keeping them lightweight and avoiding unnecessary distal extensions helps reduce base reaction moments and improves the numerical conditioning of the system. In contrast, more demanding phases (e.g., docking- or contact-like final approaches) generally benefit from smoother motion profiles (longer manoeuvres and reduced accelerations) and/or higher structural rigidity proxies, although these choices typically increase mass and can raise base reactions if implemented through distal oversizing. The co-design proof of concept further highlights that motion strategy and timing are not secondary details: even limited trajectory parameter tuning can meaningfully change disturbance metrics, and therefore should be considered together with geometry already in early trade studies.

Overall, the framework offers a practical way to select suitable HRM concepts for IOS missions. It supports decision-making by connecting architecture and sizing choices to mission-relevant performance indicators, while remaining extensible toward higher-fidelity models and broader co-design formulations.

Future developments can strengthen the approach along four main directions. First, the optimisation scope can be expanded to include discrete architectural choices (e.g., systematic selection of the number of degrees of freedom) and

Executive synthesis of recommendation matrix	Practical implication for early-stage design
Across most scenarios, smooth trajectories (especially min-jerk) and longer final phases improve robustness and reduce peak loads; explicit hold segments are particularly useful for docking/refuelling and, when needed, repair tasks.	Trajectory law and timing should be treated as design variables already in the concept phase (not only as control details), especially for contact-like or precision-critical operations.
A negative mass-distribution trend ($\gamma < 0$) is generally preferred, i.e., keeping distal links lighter and concentrating more mass proximally. This trend is strongest for inspection-oriented cases and remains beneficial in most GEO cases.	Reducing distal inertia is one of the most effective levers to mitigate base disturbance while preserving reachability and conditioning. Avoid distal thickening unless strictly required.
LEO inspection cases favour lighter geometries and low-disturbance solutions, while LEO refuelling/repair cases require a more balanced compromise between accuracy, stiffness proxies, and disturbance.	The same orbit does not imply the same architecture: mission type (inspection vs refuelling/repair) changes the optimal balance between agility and structural robustness.
GEO cases are more sensitive to geometry scaling because long reach tends to increase mass and base reaction moments; overly long distal links and “stretched” configurations are consistently unfavourable.	For GEO servicing, distribute reach across the chain and prefer conservative distal scaling. Geometry choices have a stronger penalty on disturbance than in comparable LEO cases.
Docking and refuelling (especially in GEO) are the most disturbance- and robustness-critical scenarios; they often benefit from higher stiffness proxies (mainly on the proximal half) combined with gentler terminal motion.	When accuracy and contact robustness dominate, increase stiffness carefully (prefer proximal sections) and rely on smoother final approach timing instead of adding distal mass.
Increasing the number of degrees of freedom (e.g., 10 vs 8) generally improves feasibility, conditioning, and robustness to target variability, and can reduce local joint-motion peaks through better distribution of motion.	Higher redundancy is advantageous when precision/feasibility margins are tight, but it increases computational cost and design complexity; it is most justified for demanding GEO or contact-like tasks.
Repair-oriented scenarios often favour moderate, balanced designs: too light may reduce accuracy/stability, while too heavy increases base disturbance.	For mixed-objective tasks, avoid extreme sizing choices and prioritise balanced solutions that preserve both tracking quality and disturbance control.

Table 1: Executive synthesis (engineering guidelines) derived from the thesis recommendation matrix.

richer link-wise structural variables, allowing a more realistic representation of non-uniform designs. Second, computational efficiency can be improved through profiling and vectorisation, parallel evaluation of candidate populations, and surrogate models for expensive metrics; automatic differentiation could further support more efficient gradient-based refinement. Third, the current geometry-driven mass and inertia modelling can be refined by introducing link-wise cross-section parameters and more explicit compliance representations to capture structural effects more accurately, especially for contact-like phases. Finally, the multibody verification step can be progressively automated and integrated more tightly with the optimisation

loop, while the co-design formulation can be extended to optimise richer trajectory parametrisations (timing laws and phase-specific weights) together with geometry in a constrained joint optimisation or hierarchical design loop.

References

- [1] O. Khatib B. Siciliano. *Springer Handbook of Robotics*. Springer-Verlag, 2008.
- [2] G. S. Chirikjian. *Theory and Applications of Hyper-Redundant Robotic Manipulators*, 1992.
- [3] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo. *Robotics: Modelling, Planning and Control*, 2008.