



Region of Attraction Estimation for Free-Floating Systems under Time-Varying LQR Control

- Lasse Shala** Researcher, Robotics Innovation Center, German Research Center for Artificial Intelligence (DFKI GmbH), 28359 Bremen, Germany.
- Shubham Vyas** Researcher, Robotics Innovation Center, German Research Center for Artificial Intelligence (DFKI GmbH), 28359 Bremen, Germany. Shubham.Vyas@dfki.de
- Mohamed Khalil Ben-Larbi** Postdoctoral Fellow and Group Lead, Chair of Space Technology, TU Berlin, 10587 Berlin, Germany.
- Shivesh Kumar** Researcher, Robotics Innovation Center, German Research Center for Artificial Intelligence (DFKI GmbH), 28359 Bremen, Germany.
- Enrico Stoll** Full Professor, Chair of Space Technology, TU Berlin, 10587 Berlin, Germany.

ABSTRACT

Future Active Debris Removal (ADR) and On Orbit Servicing (OOS) missions demand for elaborate closed loop controllers. Feasible control architectures should take into consideration the inherent coupling of the free floating dynamics and the kinematics of the system. Recently, Time-Varying Linear Quadratic Regulators (TVLQR) have been used to stabilize underactuated systems that exhibit a similar kinodynamic coupling. Furthermore, this control approach integrates synergistically with Lyapunov based region of attraction (ROA) estimation, which, in the context of ADR and OOS, allows for reasoning about composability of different sub-maneuvers. In this paper, TVLQR was used to stabilize an ADR detumbling maneuver in simulation. Moreover, the ROA of the closed loop dynamics was estimated using a probabilistic method. In order to demonstrate the real-world applicability for free floating robots, further experiments were conducted onboard a free floating testbed.

Keywords: Active Debris Removal; Space Robots; Region of Attraction; Time-Varying LQR; Free-Floating Robot Control

1 Introduction

Future spaceflight endeavors, undertaken by government agencies and private companies, envision a sustainable use of Low and Geostationary Earth Orbits (LEO/GEO) and further venturing out into the solar system. May it be the life extension vehicles by Northrop Grumman [1], manned missions to Mars by SpaceX [2], or the international Lunar Gateway project [3] involving NASA, ESA, JAXA and CSA, a common challenge for many upcoming projects is the in-orbit operations process. This will require capabilities to manipulate and maneuver accurately with and around cooperative and non-cooperative objects in orbit.

Corresponding Author: Shubham Vyas Shubham.Vyas@dfki.de

Active debris removal (ADR) and On-Orbit servicing (OOS) are required to reduce or at least stabilize the current space debris environment [4, 5]. In such a mission an active spacecraft (Chaser) would approach a debris object (Target), dock to it (e.g. using a robotic manipulator), optionally detumble, and finally remove it from the region of concern (ADR) or extend its operational life (OOS) through servicing measures, i.e. refueling or repair. This will help to prevent collisions that would otherwise render parts of the orbit inoperable for decades to come [4]. Furthermore, projects involving large structures in space require powerful and dedicated launchers, which constitute a technical and financial challenge. However, this limitation can be overcome by innovative techniques, such as additive manufacturing, which makes it possible to build components and assemble them directly in space, putting virtually no constraints on the structure's geometry.

For all the aforementioned applications, sophisticated control systems are required to take care of the simultaneous movement of a mobile base, i.e. the spacecraft, and a robotic arm (potentially grasping a target object) to allow for efficient manipulation, flawless assembly, or continuous printing. For a free-floating manipulator, any end effector position change in an inertial frame is a function of the change in the joint positions as well as the configuration-dependent inertia properties [7]. Therefore, in general, the inverse and forward kinematics for space manipulators are also considered dynamics problems [7]. This kino-dynamic coupling is illustrated in Figure 1. Moreover, the time-varying motion constraints imposed by a moving, potentially tumbling, target further complicates the grasping task and requires agile planning capabilities within an operationally useful time frame.

While real-time dynamic re-planning is an increasingly plausible alternative [8–10], the majority of the proposed techniques still face significant performance challenges because of the high computational requirements for on-board processors and the inability to convexify certain constraints [8]. To address these issues, the control problem can be split into two parts. In this case, the trajectory optimization problem is solved offline, or at a non-real-time rate, and a linear controller is used for online stabilization. Within the field of optimal control of underactuated systems, the usage of Time-Varying Linear Quadratic Regulators (TVLQR) as locally optimal feedback control policy to stabilize planned trajectories has gained popularity [11–15]. TVLQR requires little computational capacity, can be run at high frequencies, and integrates synergistically with Lyapunov-based methods for Region of Attraction (RoA) estimation that can be used to verify closed-loop stability. So-called "funnels" describe time-varying regions of certifiable stability around trajectories [15–17]. In the context of ADR and OOS, such analyses are especially interesting when considering space qualification of control algorithms. Furthermore, RoA analysis could be an interesting tool enabling controller composition in the spirit of [18], where it could be used to allow for safe transitioning from one phase of a maneuver to the next. For example, considering a robotic ADR scenario consisting of a final approach, capture, and a post-capture/detumble phase [19], determining the set of stabilizable states of the closed-loop dynamics during each phase is beneficial. Doing so allows the connection of each individual phase to the subsequent one while ensuring the stability of the entire ADR maneuver (depicted in Figure 2). Within this work, a probabilistic RoA estimation pipeline for free-floating spacecraft under TVLQR control was conceived and applied to two distinct scenarios, namely: *scenario I*, a simulated multibody scenario, and *scenario II*, a simple trajectory tracking scenario that was realized onboard the ELISSA free floating testbed [20, 21]. Within *scenario I*, the goal was to investigate the usefulness of probabilistic RoA estimation in an ADR context, whereas for *scenario II*, the focus was put on investigating the real-world applicability of RoA estimates.

This paper is organized as follows: Section 2 provides background information on the dynamical models and the control approach used throughout this work and shortly introduces the concept of RoA.




Figure 1 Coupling of kinematics and dynamics, adopted from [6]. For manipulators fixed to the ground, there is no significant motion of the base link, when moving individual joints. The contrary holds for manipulators mounted on a spacecraft.

In particular, the control approach used within this work consists of stabilizing a precomputed nominal open loop trajectory obtained using trajectory optimization (Section 2.2) with TVLQR (Section 2.3). This method has proven to be capable of stabilizing a range of systems with similar kinodynamic couplings, i.e. severely underactuated systems [14, 22], fully-actuated [23], and 3-DoF freeflyers [24]. Section 3 introduces the probabilistic RoA estimation methodology that is then applied to the two aforementioned scenarios (Section 4). The methodology and results are then discussed in Section 5.

2 Background

2.1 Equations of Motion

The equations of motion (EoM) for free-floating multibody systems have been studied in detail in the literature. Hence, we only provide the relevant equations here. For in-depth information on the derivation of the EoM, we direct the reader to the following works: [25–28].

With $\mathbf{q} = [\mathbf{q}_0^T \ \mathbf{q}_m^T]^T$ denoting the pose of the freeflyer that is composed of the attitude of the base and the joint angles of the manipulator, respectively, the EoM can be written as ([29]):

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) = \mathbf{u} \quad (1)$$

Here, \mathbf{q} and $\dot{\mathbf{q}}$ are the generalized positions and velocities of the system, respectively, $\mathbf{M} \in \mathbb{R}^{6+n \times 6+n}$ is the invertible generalized mass-inertia matrix of the system, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^{6+n}$ is the generalized Coriolis and Centrifugal effort, and $\mathbf{u} \in \mathbb{R}^{6+n}$ is the vector of generalized forces with n being the number of degrees of freedom in the robot arm attached to the base. The equations of motion can also be represented in state-space form by defining $\mathbf{x} = [\mathbf{q}, \dot{\mathbf{q}}]^T$. Doing so, the derivatives of \mathbf{x} can be written as:

$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{\mathbf{q}} \\ \ddot{\mathbf{q}} \end{bmatrix} = \begin{bmatrix} \dot{\mathbf{q}} \\ \mathbf{M}(\mathbf{q})^{-1}[\mathbf{u} - \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})] \end{bmatrix} \quad (2)$$

The mass inertia matrix of the system can be written as:

$$\mathbf{M} = \begin{bmatrix} \mathbf{M}_0 & \mathbf{M}_{bm} \\ \mathbf{M}_{bm}^T & \mathbf{M}_m \end{bmatrix} \quad (3)$$




Figure 2 Composition of a multi-phase ADR scenario using funnels in the spirit of [18]. Each funnel represents the stabilizing closed-loop dynamics during a given phase of ADR. Stability over the entire maneuver can be achieved by ensuring that states at the outlet of a preceding funnel are a subset of the states of the subsequent one.

where $\mathbf{M}_0 \in \mathbb{R}^{6 \times 6}$ is the base-spacecraft mass-inertia matrix component, $\mathbf{M}_m \in \mathbb{R}^{n \times n}$ is the manipulator mass-inertia matrix component, and $\mathbf{M}_{bm} \in \mathbb{R}^{6 \times n}$ is the dynamic-coupling inertia matrix component ([27]). Note, that in Section 4.1, the mass inertial properties of the debris are also modeled in \mathbf{M}_m . Using the components of the mass-inertia matrix, the linear and angular momentum of the free-floating system can be written as:

$$\begin{bmatrix} \mathbf{L} \\ \mathbf{P} \end{bmatrix} = \mathbf{M}_0 \dot{\mathbf{q}}_0 + \mathbf{M}_m \dot{\mathbf{q}}_m \quad (4)$$

Here, $\mathbf{L} \in \mathbb{R}^3$ and $\mathbf{P} \in \mathbb{R}^3$ represent the total angular and linear momentum of the system respectively. $\dot{\mathbf{q}}_0 \in \mathbb{R}^6$ and $\dot{\mathbf{q}}_m \in \mathbb{R}^n$ are the generalized velocities of the base and manipulator respectively. The angular momentum equations form a non-holonomic constraint on the system which makes the control of the system complex.

2.2 Trajectory Optimization

For every phase of ADR/OOS, a nominal state-space trajectory of the system has to be found which accomplishes the task for the given phase. In addition to being able to realize the phase's task, the trajectory also has to be dynamically feasible i.e. satisfy the dynamics constraints, minimize certain costs, and also allow for actuation and state constraints. Trajectory optimization based on Direct Transcription [30, 31] was employed to find such a trajectory. The general formulation of the trajectory optimization problem can be written as:

$$\mathbf{q}(t), \dot{\mathbf{q}}(t), \mathbf{u}(t) = \arg \min \int_0^{t_f} \left(w_t \Delta t + \mathbf{u}^T \mathbf{W}_u \mathbf{u} \right) dt + \mathbf{x}_f^T \mathbf{W}_{x_f} \mathbf{x}_f \quad (5a)$$

subject to:

$$\mathbf{M}(\mathbf{q})\dot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} = \mathbf{u} \quad (5b)$$

$$\mathbf{q}_{min} \leq \mathbf{q}(t) \leq \mathbf{q}_{max} \quad (5c)$$

$$\dot{\mathbf{q}}_{min} \leq \dot{\mathbf{q}}(t) \leq \dot{\mathbf{q}}_{max} \quad (5d)$$

$$\mathbf{u}_{min} \leq \mathbf{u}(t) \leq \mathbf{u}_{max} \quad (5e)$$

$$\mathbf{q}(t_0) = \mathbf{q}_0, \quad \mathbf{q}(t_f) = \mathbf{q}_f \quad (5f)$$

$$\dot{\mathbf{q}}(t_0) = \dot{\mathbf{q}}_0, \quad \dot{\mathbf{q}}(t_f) = \dot{\mathbf{q}}_f \quad (5g)$$

Equation 5a shows the general cost function used with the running cost on time (with weight w_t as tuning parameter, time step Δt as a decision variable, and final time t_f) to find a minimum time solution and a quadratic cost on actuation \mathbf{u} with weight $\mathbf{W}_u \in \mathbb{R}^{6+n \times 6+n}$. The terminal cost shown here is a quadratic cost on the final state of the system with weight $\mathbf{W}_{x_f} \in \mathbb{R}^{12+2n \times 12+2n}$. The cost on final velocity using the given terminal cost formulation can be used for the detumbling phase scenario in Section 4.1. Equation 5b - Equation 5g represent the constraints on the optimization problem. Equation 5b represents the dynamics constraints, Equation 5c, Equation 5d, and Equation 5e represent the position, velocity, and actuation limit constraints respectively, and Equation 5f and Equation 5g represent the initial and final position and velocity constraints, if any.

The above-given non-linear optimization problem is transcribed into a non-linear program using direct transcription with Euler's integration scheme. The resulting non-linear problem is solved using the *SNOPT* (*S*parse *N*onlinear *O*PTimizer) solver [32] in Drake [33]. While higher order methods for trajectory optimization can be found in literature [30, 34, 35], we reason that for the purpose of this work, the accuracy provided by the Euler integration scheme combined with $0.01 \leq dt \leq 0.2$ is sufficient. The

dynamics are satisfied at each knot point in the trajectory and the errors due to system identification of the model of the physical system are generally much larger than errors from the integration scheme in trajectory optimization.

The trajectory optimization method described here is used throughout both of the example scenarios in Section 4. For the detumble scenario in Section 4.1, the aim is to bring the post-capture system to a stationary state with zero angular velocity i.e. detumbling. Thus, the terminal cost along with the terminal constraint on the velocity assists in the detumble trajectory optimization. For the experimental results in Section 4.2, the position constraints along the trajectory were used to enforce a circular trajectory.

2.3 Trajectory Stabilization using Time-Varying Linear Quadratic Regulators

For trajectory stabilization using TVLQR, as the trajectory from trajectory optimization in the previous section leads to a stabilizable fixed point of the system, the infinite horizon stabilization of the trajectory can be guaranteed within the computed RoA. The optimal open-loop trajectories $(\mathbf{x}^*(t), \mathbf{u}^*(t))$ from Section 2.2 cannot be directly run in a dynamics simulator or on the real-system due to integration errors, modeling inaccuracies, sensor inaccuracies and disturbances. Hence, a closed-loop control feedback policy is needed to stabilize the given trajectory. Here, TVLQR is employed to stabilize a trajectory leading to a fixed point \mathbf{x}_f , which is then stabilized using an infinite horizon LQR. For the TVLQR synthesis, a time-varying linearization of the system along the trajectory is needed. For this, a moving reference frame is introduced along the trajectory:

$$\bar{\mathbf{x}}(t) = \mathbf{x}(t) - \mathbf{x}^*(t), \quad \bar{\mathbf{u}}(t) = \mathbf{u}(t) - \mathbf{u}^*(t) \quad (6)$$

The nonlinear dynamics are first linearized along the nominal trajectory $\mathbf{x}^*(t)$ and $\mathbf{u}^*(t)$ using first-order Taylor expansion yielding a linear time-varying system in the new coordinates described by [23, 36]:

$$\dot{\bar{\mathbf{x}}} = \mathbf{f}(\mathbf{x}, \mathbf{u}) - \mathbf{f}(\mathbf{x}^*, \mathbf{u}^*) \approx \underbrace{\frac{\partial \mathbf{f}(\mathbf{x}^*, \mathbf{u}^*)}{\partial \mathbf{x}}}_{:=\mathbf{A}(t)} \bar{\mathbf{x}}(t) + \underbrace{\frac{\partial \mathbf{f}(\mathbf{x}^*, \mathbf{u}^*)}{\partial \mathbf{u}}}_{:=\mathbf{B}(t)} \bar{\mathbf{u}}(t) = \mathbf{A}(t)\bar{\mathbf{x}}(t) + \mathbf{B}(t)\bar{\mathbf{u}}(t) \quad (7)$$

For this time-varying linear system, we can now define a quadratic cost function to drive the error to zero in the error coordinates thereby following the nominal trajectory. The cost function can be written as:

$$J = \int_0^{t_f} (\bar{\mathbf{x}}^T \mathbf{Q} \bar{\mathbf{x}} + \bar{\mathbf{u}}^T \mathbf{R} \bar{\mathbf{u}}) dt + \bar{\mathbf{x}}_f^T \mathbf{Q}_f \bar{\mathbf{x}}_f \quad (8)$$

Where $\mathbf{Q} = \mathbf{Q}^T \geq 0$ and $\mathbf{R} = \mathbf{R}^T > 0$ are positive semi-definite weight matrices that penalize deviations in state and actuation, respectively, along the trajectory and $\mathbf{Q}_f = \mathbf{Q}_f^T \geq 0$ penalizes the error on the final state which is a fixed point. A closed loop optimal control policy for the linearized problem can be obtained by solving the differential Riccati Equation [37]:

$$-\dot{\mathbf{S}}(t) = \mathbf{S}(t)\mathbf{A}(t) + \mathbf{A}^T(t)\mathbf{S}(t) - \mathbf{S}(t)\mathbf{B}(t)\mathbf{R}^{-1}\mathbf{B}^T(t)\mathbf{S}(t) + \mathbf{Q} \quad (9)$$

Integrating Equation 9 backwards with $\mathbf{S}(t_f) = \mathbf{Q}_f$ yields $\mathbf{S}(t)$, the optimal cost-to-go. The feedback law can be written as:

$$\mathbf{u} = \mathbf{u}^*(t) - \mathbf{K}(t)(\mathbf{x}(t) - \mathbf{x}^*(t)) \quad (10)$$

Where

$$\mathbf{K}(t) = -\mathbf{R}^{-1}\mathbf{B}(t)^T\mathbf{S}(t) \quad (11)$$

An advantage of utilizing TVLQR is, that the linear optimal cost-to-go $J^*(t) = \bar{\mathbf{x}}^T(t)\mathbf{S}(t)\bar{\mathbf{x}}(t)$ is a suitable candidate for a time varying Lyapunov function, which enables performing subsequent stability and robustness analysis of the closed loop dynamics.

2.4 Region of Attraction

The region of attraction (RoA) \mathcal{R}_a around a fixed point \mathbf{x}^* of a nonlinear system $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}(t))$ is the set of states that evolve to the fixed point as times approaches infinity. Oftentimes, there is no simple, closed-form description of the RoA for closed-loop systems. However, there exist multiple methods to estimate it, many of which build upon the notion of stability in the sense of Lyapunov. More specifically, the RoA can be estimated by considering sublevel sets of Lyapunov functions [38]. Doing so, we define an inner approximation of $\mathcal{B} \subset \mathcal{R}_a$ by writing:

$$\mathcal{B} = \{\mathbf{x} | V(\mathbf{x}) < \rho\} \quad (12)$$

Assuming a Lyapunov function $V(\mathbf{x})$ is known, estimating the RoA is a matter of finding the greatest scalar $\rho > 0$, such that the Lyapunov conditions are satisfied, i.e. $V(\mathbf{x}) > 0$ and $\dot{V}(\mathbf{x}) = \nabla V\mathbf{f}(\mathbf{x}) \leq 0$ for $x \in \mathcal{B}$. There exist multiple ways to estimate ρ . One can use Sum of Squares programming (SoS) and convex optimization to reason about the problem algebraically and obtain a formal certificate for stability [39]. However, formulating SOS programs can often be challenging, not only because of the background knowledge that is required but also because doing so requires expressing the closed-loop dynamics in a polynomial, non-rational form. Besides this optimization-based approach, there exist probabilistic methods for certifying stability in the sense of Lyapunov [40]. For time-invariant systems, one can evaluate the right-hand side of $\mathbf{f}(\mathbf{x})$ for states $\tilde{\mathbf{x}}$ randomly sampled from the current estimate \mathcal{B} and check whether the Lyapunov conditions are satisfied. If this is not the case, ρ is reduced to $\rho = V(\tilde{\mathbf{x}})$ and the next $\tilde{\mathbf{x}}$ is sampled from the shrunk estimate. This process is then repeated until convergence and yields a probabilistic certificate of stability.

For time-varying dynamics $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}(t), t)$, the notion of a funnel described by a sublevel set of a time-varying Lyapunov function can be seen as an extension of the RoA to time-varying systems [15–17]:

$$\mathcal{B}(t) = \{\mathbf{x} | V(\mathbf{x}, t) < \rho(t)\} \quad (13)$$

$$\mathbf{x}(t) \in \mathcal{B}(t) \implies \mathbf{x}(t_f) \in \mathcal{B}_f \quad (14)$$

where in contrast to Equation 12, $\mathcal{B}(t)$, $V(\mathbf{x}, t)$ and $\rho(t)$ are now time varying. In Equation 13 $\mathcal{B}(t)$ describes a time-varying set of states leading to a set of goal states \mathcal{B}_f . Similar as for the time invariant case, $\mathcal{B}(t)$ can be estimated by using SoS based [14–16], or probabilistic [12, 13] methods. Instead of evaluating the right-hand side as was done sampling based time invariant RoA estimation by [40], the approach presented by [12] relies on exhaustive simulations of the dynamics.

3 Region of Attraction Estimation

In order to estimate funnels according to the definition given in Section 2.4 for closed loop dynamics that incorporate TVLQR feedback policies, a probabilistic approach based on [12] was chosen, mainly due to the greater flexibility when it comes to adding constraints or making changes to the model. As stated in Section 2.3, considering the nonlinear closed loop dynamics \mathbf{f}_{cl} of a multibody-system under

TVLQR control, the quadratic cost-to-go J^* can be used as a Lyapunov function, i.e. $V = J^* = \mathbf{x}^T \mathbf{S} \mathbf{x}$. Doing so yields the time-varying RoA (funnel) around \mathbf{x}^* :

$$\mathcal{B}(t) = \{\bar{\mathbf{x}} | \bar{\mathbf{x}}^T(t) \mathbf{S}(t) \bar{\mathbf{x}}(t) < \rho(t)\} \quad (15)$$

which, since $\mathbf{S}(t) > 0$, describes a time-varying hyperellipsoid. In practice, a discrete time version of Equation 15 is used and ρ is estimated around the collocation points of \mathbf{x}^* , which results in a set of hyperellipsoids as depicted in Figure 3.

The estimation process is based on simulating the closed loop dynamics \mathbf{f}_{cl} forward in time for initial states randomly sampled from a continuously shrinking estimate of the inlet of the estimated funnel $\mathcal{B}_0 = \mathcal{B}(t = 0)$. More specifically, initial states that after integration of \mathbf{f}_{cl} do not lead to a final state with a linear cost-to-go $J^*(t_f) < \rho_f$ or violate constraints, cause shrinking of the hyperellipsoids that make up the funnel, including its inlet \mathcal{B}_0 . Using this methodology, funnels can only be shrunk, never expanded. Pseudocode for the algorithm used in this work can be found in Algorithm 1.

Algorithm 1 Probabilistic RoA Estimation for TVLQR

```

1:  $\mathbf{x}^*(t), \mathbf{u}^*(t) \leftarrow$  nominal trajectory
2:  $\mathbf{A}(t), \mathbf{B}(t) \leftarrow$  linearized dynamics along  $\mathbf{x}^*(t), \mathbf{u}^*(t)$ 
3:  $\mathbf{S}(t), \mathbf{K}(t) \leftarrow$  Finite Horizon LQR for  $\mathbf{A}(t), \mathbf{B}(t)$ 
4:  $n \leftarrow$  Nr. of Simulations
5:  $m \leftarrow$  Nr. of Evaluation Steps
6:  $\rho_{0,k} = \infty, k \in \{0, \dots, n-1\}$ 
7:  $\rho_{0,f} = \rho_n = \bar{\mathbf{x}}_{f,\max}^T \mathbf{S}_f \bar{\mathbf{x}}_{f,\max}$  ▷ Definition of goal region
8: for  $j = 1, 2, \dots, n$  do
9:   shrink=false
10:   $\mathbf{x}_0 \leftarrow$  random initial state from  $\{\mathbf{x} | \bar{\mathbf{x}}_0^T S_0 \bar{\mathbf{x}}_0 < \rho_{j-1,0}\}$ 
11:   $J_{j,0}^* = \bar{\mathbf{x}}_0^T S_0 \bar{\mathbf{x}}_0$ 
12:  for  $k = 1, 2, \dots, m$  do
13:    Integrate  $\mathbf{f}_{cl}$  until  $t = t_k$ 
14:    if Constraint Violation then
15:      shrink=true
16:      break
17:    else
18:       $J_{j,k}^* = \bar{\mathbf{x}}_k^T S_k \bar{\mathbf{x}}_k$  ▷ compute cost-to-go
19:      if  $J_{j,k}^* > \rho_{j-1,k}$  then
20:        shrink=true
21:        break
22:      end if
23:    end if
24:  end for
25:  if shrink == true then
26:     $\rho_{j,\kappa} = J_{j,\kappa}^*, \kappa \in \{0, 1, \dots, k-1\}$  ▷ update estimates
27:  end if
28: end for

```

We initialize our estimate of ρ with $\rho = \infty$ along the trajectory, except for ρ_f , associated to the outlet $\mathcal{B}_f = \mathcal{B}(t = t_f)$ (Algorithm 1, lines 6 and 7). The value of ρ_f remains fixed throughout the estimation process and can be defined by considering the cost-to-go J_{\max}^* of a maximum allowed deviation to the nominal final state $\bar{\mathbf{x}}_{\max}^T(t_f)$ as described in [22]:

$$\rho_f = \rho(t_f) = J_{\max}^*(t_f) = \bar{\mathbf{x}}_{\max}^T(t_f) \mathbf{S}(t_f) \bar{\mathbf{x}}_{\max}(t_f) \quad (16)$$

where $\bar{\mathbf{x}}_{\max}^T$ is the user-defined, tolerable maximum deviation from the final state. In order to ensure infinite horizon stability, ρ_f must be chosen, such that the outlet $\mathcal{B}_f \subset \mathcal{B}_\infty$, where \mathcal{B}_∞ is the RoA associated to the closed loop dynamics (infinite horizon LQR) around a fixed point at the end of \mathbf{x}^* . After initializing ρ_0 and ρ_f , a random initial condition is sampled from \mathcal{B}_0 (Algorithm 1, line 10) using a similar method as was used in [41]. We found simple rejection sampling to be infeasible for sampling from a constantly shrinking \mathcal{B}_0 in a potentially high dimensional state space. The reason is the decline in acceptance rate of samples with increasing dimensionality of the sampling space [42]. Hence, direct sampling was employed, which is more complicated, but also more efficient and even required for systems with higher dimensional state spaces, such as the one that will be discussed in Section 4.1. States are sampled from the n dimensional unit sphere using the *direct-sphere* algorithm found in [42] and then mapped to \mathcal{B}_0 , the hyperellipsoid that describes the inlet of \mathcal{B} in Equation 15:

$$\mathcal{B}_0 = \{ \mathbf{x} | \bar{\mathbf{x}}_0^T \mathbf{S}_0 \bar{\mathbf{x}}_0 \leq \rho_0 \}. \quad (17)$$

Consider the sphere $\mathcal{S}_n = \{ \mathbf{y} | \mathbf{y}^T \mathbf{y} \leq 1 \}$ and a linear mapping $\mathbf{T} \in \mathbb{R}^{n \times n}$, so that $\mathbf{y} = \mathbf{T}\mathbf{x}$. This allows to express \mathcal{S} in terms of \mathbf{x} :

$$\mathcal{S}_n = \{ \mathbf{x} | \mathbf{x}^T \mathbf{T}^T \mathbf{T} \mathbf{x} \leq 1 \} \quad (18)$$

Comparing Equation 17 with Equation 18, we now seek to find \mathbf{T} , such that $\mathbf{T}^T \mathbf{T} = \rho_0^{-1} \mathbf{S}_0$. A diagonalization of $\mathbf{T}^T \mathbf{T}$ yields:

$$\mathbf{T}^T \mathbf{T} = \mathbf{W}^T \mathbf{\Lambda} \mathbf{W} \quad (19)$$

where \mathbf{W} and $\mathbf{\Lambda}$ are matrices of eigenvectors and -values of $\rho_0^{-1} \mathbf{S}_0$, respectively. With $\mathbf{T} = \mathbf{\Lambda}^{\frac{1}{2}} \mathbf{W}$, the mapping from a random $\mathbf{y} \in \mathcal{S}_n$ to a random $\mathbf{x}_0 \in \mathcal{B}_0$ is:

$$\mathbf{x}_0 = \mathbf{x}_0^* + \left(\mathbf{\Lambda}^{\frac{1}{2}} \mathbf{W} \right)^{-1} \mathbf{y} \quad (20)$$

In the following, we describe the shrinking operation (visualized in Figure 4). For this, let m and n denote the number of simulations made during the estimation process and the number of knot points of \mathbf{x}^* , respectively. Furthermore, $j \in [0, \dots, m]$ and $k \in [0, \dots, n]$, denote the index associated with the current simulation and knot point, respectively. After sampling an initial state $\mathbf{x}_{j,0}$, the closed loop dynamics \mathbf{f}_{cl} are integrated piece-wise in between knot points that are evenly spaced wrt. time. At each knot-point, where $t_k = k\Delta t$, the linear optimal cost-to-go $J_{j,k}^*$ of the current state $\mathbf{x}_{j,k}$ is compared to that of the associated value of $\rho_{j-1,k}$ made in the previous, that is $(j-1)$ th, simulation. If $\mathbf{x}_{j,k}$ is outside the existing funnel estimate, that is, $J_{j,k}^* > \rho_{j-1,k}$, all estimates of ρ made for previous knot-points within




Figure 3 Time varying RoA estimate using a set of ellipses (blue). \mathcal{B}_0 and \mathcal{B}_f are the inlet and the outlet of the funnel, respectively. \mathcal{R}_a shows the actual, albeit unknown RoA (green). The set of ellipses \mathcal{B}_k is an approximation of the funnel $\mathcal{B}(t)$.




Figure 4 Estimation process for a hypothetical, two-dimensional system. Whenever the dynamics lead outside the previously estimated funnel (light blue ellipses), or a constraint has been violated (red "X") all preceding estimates have to be shrunk (darker blue ellipses). The green area shows the real funnel, that is iteratively being approximated by the set of ellipses.

preceding simulations ($\rho_{j-1,k-1}, \rho_{j-1,k-2}, \dots, \rho_{j-1,0}$) have to be adjusted to the optimal cost-to-go values estimated during this simulation, i.e. $\rho_{j,k} = J_{j,k}^*$ for $\kappa \in \{0, \dots, k\}$ (Algorithm 1, line 18). Similarly, if during the integration of \mathbf{f}_{cl} a constraint is violated, we resize the funnel (Algorithm 1, line 14). Thus, the shrinking policy can be written as:

$$\rho_{j,\kappa} = \begin{cases} J_{j,\kappa}^* & \text{for } \kappa \in \{0, \dots, k\}, & \text{if constraint violated (alg 1., line 15)} \\ & \text{or } J_{j,k+1}^* > \rho_{j-1,k+1} \text{ (alg 1., line 21)} \\ \rho_{j-1,\kappa} & \text{for } \kappa \in \{0, \dots, n-1\}, & \text{else} \end{cases} \quad (21)$$

After aborting or finishing a simulation successfully, we continue by sampling the next initial state and start over until convergence can be seen, or $j = m$.

4 Results

Two scenarios have been conceived to investigate the applicability of TVLQR and RoA estimation for space robotics. First, using the RoA estimation methodology introduced in Section 2.4, the funnel of a multi-body detumbling scenario has been computed in simulation. Secondly, in order to demonstrate TVLQR on a physical free-floating system, experiments have been conducted onboard the Experimental Lab for Proximity Operations and Space Situational Awareness (ELISSA), a free-floating laboratory.

4.1 Scenario I: Simulated Multibody ADR Detumbling Maneuver

A post-capture detumbling scenario has been conceived, in which a chaser satellite with a 3-DoF robot arm has already docked to a target debris object (as shown in Figure 5), that is still rotating with an angular velocity of $\omega_0 = 5^\circ \text{s}^{-1}$. Assuming perfect synchronization of the grasping mechanism and the capture point, it can be assumed that there were no contact forces involved during the capture i.e. ideal capture. This allows the post-capture capture velocities to be determined using the Generalized Jacobian Matrix [43] when assuming the center of mass of the target is the end-effector frame for the combined post-capture system. The chaser and target were modeled as cubes with side lengths of $l_c = 2 \text{ m}$ and $l_t = 0.6 \text{ m}$, and masses of 100 kg and 50 kg , respectively and with a homogenous mass distribution. The manipulator's link masses are $m_1 = 10 \text{ kg}$, $m_2 = 8 \text{ kg}$ and $m_3 = 4 \text{ kg}$ and their respective lengths are $l_1 = 0.9 \text{ m}$, $l_2 = 0.7 \text{ m}$ and $l_3 = 0.3 \text{ m}$. The dynamics of this system, described by Equation 1 were linearized according to [36], so that the state vector \mathbf{x} can be written as:




Figure 5 Chaser (partially visible) and target satellite are connected by a 3R manipulator. In the beginning (orange) the chaser is rotating. After executing the detumbling controller, there is no relative velocity of both spacecraft (blue).

$$\mathbf{x} = \left[\hat{\mathbf{q}}_B \quad \mathbf{p}_B \quad q_1 \quad q_2 \quad q_3 \quad \boldsymbol{\omega}_B \quad \mathbf{v}_B \quad \dot{q}_1 \quad \dot{q}_2 \quad \dot{q}_3 \right]^T \quad (22)$$

where $\hat{\mathbf{q}}_B$ denotes the vector part of a quaternion, $\mathbf{p}_B \in \mathbb{R}^3$ is the position of the floating base and q_1, q_2 and q_3 specify the manipulator's configuration. Angular and translational velocities are given by $\boldsymbol{\omega}_B \in \mathbb{R}^3$ and $\mathbf{v}_B \in \mathbb{R}^3$, respectively. The manipulators joint velocities are written as \dot{q}_1, \dot{q}_2 and \dot{q}_3 . The complete input vector \mathbf{u} is:

$$\mathbf{u} = \left[\tau_x \quad \tau_y \quad \tau_z \quad f_x \quad f_y \quad f_z \quad \tau_1 \quad \tau_2 \quad \tau_3 \right]^T \quad (23)$$

where $\tau_x, \tau_y, \tau_z, \tau_1, \tau_2$ and τ_3 are torques applied to the base and manipulator, respectively. In addition, the base is assumed to be equipped with thrusters, such that forces f_x, f_y and f_z can be applied. Actuation limits were set to $\pm 50\text{Nm}$ for the base and manipulator torques and $\pm 10\text{N}$ for the base forces. A detumble trajectory was found using the above-given model and actuation using the transcription method in Section 2.2. This trajectory was stabilized using the method given in Section 2.3 and executed on the Drake dynamics simulator [33].

For RoA estimation, a generalized fuel constraint has been implemented, that does not make any specific assumptions of the actuators used. Instead, the idea of generalized forces applied at the center of mass of the base for the satellite and at the joints of the robot arm was employed. The fuel spent at time t is calculated as the integral of the sum of all elements of $\mathbf{u}(t)$:

$$F(t) = \int_{\tau=0}^t \sum_{k=0}^9 |u_k(t)| d\tau \quad (24)$$

We calculate F_0 , the fuel usage of the nominal trajectory by integrating Equation 24 over the duration of the entire nominal trajectory (with $\mathbf{u} = \mathbf{u}^*$). Furthermore, we define the maximal fuel usage for the entire maneuver as:

$$F_{\max} = (1 + \alpha) F_0 \quad (25)$$

where α is a parameter that relates the nominal amount of fuel to the share that TVLQR is allowed to use for stabilization. While carrying out the RoA estimation, Equation 24 is evaluated and compared to F_{\max} . If $F(t) > F_{\max}$, the simulation loop is aborted (Algorithm 1, line 14).

First, the RoA was calculated for $\alpha = \infty$, essentially bypassing the fuel limit. Furthermore, values of $\alpha = 1$ and $\alpha = \frac{1}{2}$ were used. For each α , 5 RoA estimations were conducted with $n = 1000$ simulations each, and $m = 100$ knot points. Figure 6 a) shows the final values of ρ for three different fuel limit settings. One can see, that while estimates for equal values of α lie close together, there is a discrepancy when seeding the random number generator differently. This is




Figure 6 a) Evolution of ρ for different fuel limit settings and over multiple estimation processes. b) to d) projections of the estimated RoA into different subspaces.

caused by the probabilistic nature of the estimation process in combination with the relatively large state space and the limited number of simulations that could be made within a reasonable time of ≈ 8 hours on a recent laptop (Intel i9, 32GB RAM). Furthermore, one can see that some of the lines stop before $k = 100$

$$\min_{\mathbf{x}_k, \mathbf{u}_k} \sum_{k=0}^{100} \Delta t_k \left(W_t + \mathbf{u}_k^T \mathbf{W}_u \mathbf{u}_k \right) \quad (27a)$$

$$\text{s.t. } \mathbf{q}_{k+1} = \mathbf{q}_k + \Delta t_k \dot{\mathbf{q}}_k \quad (27b)$$

$$\dot{\mathbf{q}}_{k+1} = \dot{\mathbf{q}}_k + \Delta t_k \ddot{\mathbf{q}}_k \quad (27c)$$

$$\ddot{\mathbf{q}}_k - \mathbf{M}^{-1} \mathbf{u}_k = 0 \quad (27d)$$

$$\Delta t_{k+1} = \Delta t_k \quad (27e)$$

$$- \mathbf{u}_{\max} \leq \mathbf{u}_k \leq \mathbf{u}_{\max} \quad (27f)$$

$$\mathbf{q}_{30} = \left[3 \quad 1 \quad \frac{\pi}{2} \right]^T \quad (27g)$$

$$\mathbf{q}_{50} = \left[4 \quad 2 \quad \pi \right]^T \quad (27h)$$

$$\mathbf{q}_{70} = \left[3 \quad 3 \quad \frac{3\pi}{2} \right]^T \quad (27i)$$

$$(27j)$$

Where Δt_k is the integration step size, and the cost consists of W_t penalizing the overall duration of trajectories and a quadratic cost on \mathbf{u} with a weighing defined by $\mathbf{W}_u \in \mathbb{R}^{3 \times 3}$. A simple box quadrature was used as an integration scheme to model the double integrator dynamics (Equation 27b and Equation 27c):

$$\mathbf{q}_{k+1} - \mathbf{q}_k = \int_{t_k}^{t_{k+1}} \dot{\mathbf{q}}(t) dt \approx (t_{k+1} - t_k) \dot{\mathbf{q}}(t_k) = \Delta t_k \dot{\mathbf{q}}_k \quad (28)$$

Furthermore, Equation 27d describes, how \mathbf{u} affects the freeflyer i.e system dynamics, 27e forces equal time steps, and 27f describes the actuation limits (not taking into account the deadband). Intermediate waypoints (Equation 27g to 27i) were set to find a nearly circular trajectory.

A TVLQR policy was obtained according to the methodology described in Section 2.3. The weighing matrices have been set to $\mathbf{Q} = \text{diag}(50, 50, 0.01, 50, 50, 0.001)$ and $\mathbf{R} = \text{diag}(1, 1, 10)$. These weights are different from those in simulation and were tuned on the hardware directly. Equation 9 was initialized with $\mathbf{S}(t_f) = \mathbf{Q}_f = \mathbf{S}_\infty$ the solution of the infinite Horizon LQR controller around the first and final state of the trajectory.

First, to verify and demonstrate the capabilities of TVLQR for free-floating systems, the control policy was executed for a set of off-nominal initial positions around \mathbf{x}_0 . More specifically, initial positions were sampled from a 5×5 , equally spaced, quadratic grid with a side length of $1m$ thereby giving 23 data points for the experiments. The resulting trajectories of the freeflyer as measured by the tracking system are shown in Figure 9. For every initial condition (blue lines), convergence to the nominal trajectory (black line) can be observed, after the initial positional offset was corrected.

In addition to demonstrating, that TVLQR could indeed stabilize the trajectories, a RoA estimation was carried out according to the method described in Section 2.4. The results are shown in Figure 10. The light- and dark-blue ellipses are projections of the funnel in- and outlet for the given nominal trajectory (black dashed line) and controller. Each circle represents one of the 23¹ initial states \mathbf{x}_0 , all sampled from within the estimated inlet of the funnel \mathcal{B}_0 . The solid lines are the measured trajectories of the freeflyer. Green and red color gives information on whether a given initial state/trajectory leads towards the predefined goal region \mathcal{B}_f or not. At first glance, it appears as if every measured trajectory leads to a state within \mathcal{B}_f . However, recall that the ellipses are merely a projection of \mathcal{B}_f under the assumption that $v_x = v_y = 0$. Since freeflyers had nonzero terminal velocity (as indicated by the quivers in Figure 10 b)),

¹Shortly before starting the 24th run, the freeflyers battery was drained.




Figure 8 Photographs of *Hamilcar* (left) moving along the nominal trajectory (right).




Figure 9 Measured trajectories (blue) converging to the nominal trajectory (black).




Figure 10 Trajectories in the x vs y plane. Green and red indicate that the final state was in or outside the estimated RoA, respectively. Projections of the In- and Outlet of \mathcal{B} are shown as dark and light-blue ellipses, respectively. Quivers at the end of the trajectory show the magnitude and direction of the terminal velocity.

this assumption breaks. One can see, that trials for which $\mathbf{x}_f \notin \mathcal{B}_f$ usually started at the outer domain of the sampled space of initial conditions and ended close to the bounds of the projected \mathcal{B}_f , which makes sense, since the magnitude of the terminal velocity is relatively similar across trials and has a larger influence on the cost if the positional deviation is larger.

5 Discussion

As stated in Section 2.4, probabilistic methods such as the one conceived in Section 3 have the big advantage of being applicable to a large range of systems. As long as the system can be simulated, it does not matter what kind of dynamics and constraints it exhibits. This is not the case for SoS based methods, where every dynamics constraint has to be carefully transcribed into a constraint within an optimization problem. Furthermore, the fact that only non-rational polynomial functions can be used, limits the range of applicable systems or introduces inaccuracies, if one decides to use Taylor-approximations of the dynamics.

The inherent disadvantage of the approach described in Section 3 lies in the probabilistic character of the estimate. Especially for high dimensional systems, one would have to simulate the dynamics for a plethora of initial states in order for Algorithm 1 to yield dependable results. Depending on the system, performing such exhaustive simulations of multi-body dynamics could be a time-consuming undertaking. The possibilities for parallelization exist, but are limited due to the sequential nature of Algorithm 1. One could carry out the sampling and simulation in parallel, join processes and only consider the failed simulation with the smallest initial cost-to-go, when shrinking funnels. While this approach would certainly speed up computations, it will most likely not be able to fully alleviate the *curse of dimensionality*. In order to increase the likelihood of finding initial states that after simulation of \mathbf{f}_{cl} lead outside \mathcal{B}_f , one could add heuristics to the sampling process. For example, one could modify the *direct-sphere* algorithm, such that samples are preferably drawn from the outer regions of \mathcal{B}_0 . In doing so, the likelihood of the simulation failing and therefore enhancing the estimate are increased. A similar approach was found to increase the performance of probabilistic RoA estimation for linear time invariant dynamics [40]. Furthermore, in order to obtain funnels that reach towards the end of \mathbf{x}^* , \mathbf{u}^* , one could perform the piecewise simulations in reverse direction, i.e. first sample from the penultimate ellipsoid and simulate until the last, then move from back to front. A combination of forward and backward piecewise simulations could be viable. One could perform a couple of "forward" runs to quickly shrink the funnel and then continue with the reverse piecewise simulation in order to make sure, it extends until \mathcal{B}_f . However, such an approach would make it harder to include certain constraints, e.g. integral ones, such as the fuel constraint in *scenario I*.

In *scenario I*, while the RoA estimations converged and were plausible with respect to the fuel constraint (Equation 25), there is some remaining uncertainty in the final estimates. Different seeding of the random number generator leads to results that are close to one-another, but not equal (Figure 6 a). An explanation as to why different estimation processes do not yield more similar results is, that as \mathcal{B}_0 shrinks, it contains mostly initial states that are part of the actual RoA. Hence, it becomes increasingly more unlikely for a simulation to fail, and no information is added to the estimate. Even worse, in case of high dimensional state spaces, the *curse of dimensionality* affects the estimation process twice: not only does randomly picking an initial state that after simulation leads outside \mathcal{B}_f take longer than for low dimensional systems, but also the simulation itself is more demanding. Furthermore, the fuel constraint seems to negatively affect the convergence of Algorithm 1. This can be seen in Figure 6 a), in which the final estimates much more agree with one another, when the fuel constraint is bypassed, i.e. $\alpha = \infty$.

In *scenario II*, a precomputed nominal trajectory \mathbf{x}^* , \mathbf{u}^* could successfully be stabilized using TVLQR. The estimated RoA could partially be validated in lab experiments. However, certain initial states estimated to be part of \mathcal{B}_0 lead outside of \mathcal{B}_f due to exceeding the terminal velocity limit. Most likely, this was

caused by the actuation deadband of the thrusters, that was not considered in trajectory optimization and can not properly be addressed by a linear controller, such as TVLQR. Recall that towards the end of \mathbf{x}^* , the freeflyer was supposed to break and hold the position. When close to the holding pose, such that $|v_x| = |v_y| \approx 0$, TVLQR demands forces that fall below the lower threshold of minimal actuation of the thrusters. This results in early abortion of breaking and subsequent overshooting leading to a steady-state limit cycle at the end of the trajectory (or during station keeping) which was not considered in this work. For an analysis on the limit cycle during station keeping, see [46]. For trajectory optimization in *scenario II*, one could opt for a higher dimensional $\mathbf{u} \in \mathbb{R}^8$ that consists of actuator forces rather than generalized forces. Doing so would allow to include the actuation deadband within trajectory optimization, which could lead to trajectories in which opposing thrusters are active during station holding in order to prevent a nonzero terminal velocity. Besides, in order to enable freeflyers to operate in low-thrust regimes and to reduce steady state error, sigma-delta modulation could be used [47].

6 Conclusion

In this work, a probabilistic method for computing the Regions of Attractions (RoA) around trajectories for free-floating robots under TVLQR control is presented. The probabilistic method is applied to two distinct application scenarios: multibody post-capture detumbling of space debris in simulation and single rigid body trajectory tracking in experiments onboard ELISSA. The trajectories are computed using non-linear trajectory optimization and stabilized with a time-varying LQR. The stability of this controller is evaluated using Lyapunov-based methods and an algorithm for stability guarantees within the Region of Attraction is presented. From the results, it can be seen that the detumble RoA estimate allows to find the initial conditions which guarantee a successful detumble while satisfying fuel constraints. In the real-world experiments using ELISSA, it can be seen that proper modeling of the system and actuator dynamics affect the accuracy of the RoA estimate and these should be considered for future space applications. The RoA estimation methods allow for the control of future autonomous spacecraft during proximity operations as the computed funnels provide a probabilistic guarantee of stability. Furthermore, they allow the system designers to ensure the feasibility and composability of (sub-) maneuvers. In particular, probabilistic estimation methods from this work allow for modeling a wide range of constraints and effects, that need to be taken into consideration, when the goal is to reason about real-world scenarios.

Acknowledgments

This work has been performed in the M-RoCK project funded by the German Aerospace Center (DLR) with federal funds (Grant Number: FKZ 01IW21002) from the Federal Ministry of Education and Research (BMBF) and is additionally supported with project funds from the federal state of Bremen for setting up the Underactuated Robotics Lab (Grant Number: 201-342-04-1/2023-4-1). The second author acknowledges support from the Stardust Reloaded project which has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 813644.

The authors would also like to acknowledge the ELISSA Lab at the Institute of Space Systems, TU Braunschweig, for providing the facilities and support that contributed significantly to the results presented in this paper. The research was conducted during the third author's time at TU Braunschweig, who is currently affiliated with the SmallSat Rendezvous & Robotics Group at the Chair of Space Technology, TU Berlin.

References

- [1] Nola Taylor Redd. Bringing satellites back from the dead: Mission extension vehicles give de-funct spacecraft a new lease on life - [News]. *IEEE Spectrum*, 57(8):6–7, 2020. ISSN: 0018-9235. DOI: [10.1109/MSPEC.2020.9150540](https://doi.org/10.1109/MSPEC.2020.9150540).
- [2] Jennifer L. Heldmann, Margarita M. Marinova, Darlene S.S. Lim, David Wilson, Peter Carrato, Keith Kennedy, Ann Esbeck, Tony Anthony Colaprete, Richard C. Elphic, Janine Captain, Kris Zacny, Leo Stolov, Boleslaw Mellerowicz, Joseph Palmowski, Ali M. Bramson, Nathaniel Putzig, Gareth Morgan, Hanna Sizemore, and Josh Coyan. Mission Architecture Using the SpaceX Starship Vehicle to Enable a Sustained Human Presence on Mars. *New Space*, 10(3):259–273, 2022. ISSN: 2168-0256. DOI: [10.1089/space.2020.0058](https://doi.org/10.1089/space.2020.0058).
- [3] Jack O. Burns, Benjamin Mellinkoff, Matthew Spydell, Terrence Fong, David A. Kring, William D. Pratt, Timothy Cichan, and Christine M. Edwards. Science on the lunar surface facilitated by low latency telerobotics from a Lunar Orbital Platform - Gateway. *Acta Astronautica*, 154:195–203, 2019. ISSN: 00945765. DOI: [10.1016/j.actaastro.2018.04.031](https://doi.org/10.1016/j.actaastro.2018.04.031).
- [4] Mohamed Khalil Ben-Larbi, René Hensel, Gianfranco Atzeni, Eduard Arzt, and Enrico Stoll. Orbital debris removal using micropatterned dry adhesives: Review and recent advances. *Progress in Aerospace Sciences*, 134:100850, 2022. ISSN: 03760421. DOI: [10.1016/j.paerosci.2022.100850](https://doi.org/10.1016/j.paerosci.2022.100850).
- [5] Mohamed Khalil Ben Larbi, Benjamin Grzesik, Jonas Radtke, and Enrico Stoll. Active Debris Removal for Mega Constellations: CubeSat Possible? In *9th international workshop on satellite constellations and formation flying*, Boulder, Colorado, USA, 2017.
- [6] Yangsheng Xu and Takeo Kanade, editors. *Space Robotics: Dynamics and Control*, volume 188 of *The Kluwer International Series in Engineering and Computer Science*. Springer US, Boston, MA, 1993. ISBN: 978-1-4613-6595-2 978-1-4615-3588-1. DOI: [10.1007/978-1-4615-3588-1](https://doi.org/10.1007/978-1-4615-3588-1).
- [7] Angel Flores-Abad, Ou Ma, Khanh Pham, and Steve Ulrich. A review of space robotics technologies for on-orbit servicing. *Progress in Aerospace Sciences*, 68:1–26, 2014. ISSN: 03760421. DOI: [10.1016/j.paerosci.2014.03.002](https://doi.org/10.1016/j.paerosci.2014.03.002).
- [8] Josep Virgili-Llop, Costantinos Zagaris, Richard Zappulla, Andrew Bradstreet, and Marcello Romano. A convex-programming-based guidance algorithm to capture a tumbling object on orbit using a spacecraft equipped with a robotic manipulator. *The International Journal of Robotics Research*, 38(1):40–72, Jan. 2019. ISSN: 0278-3649. DOI: [10.1177/0278364918804660](https://doi.org/10.1177/0278364918804660).
- [9] Tomasz Rybus, Karol Seweryn, and Jurek Z. Sasiadek. Control System for Free-Floating Space Manipulator Based on Nonlinear Model Predictive Control (NMPC). *Journal of Intelligent & Robotic Systems*, 85(3-4):491–509, Mar. 2017. ISSN: 0921-0296, 1573-0409. DOI: [10.1007/s10846-016-0396-2](https://doi.org/10.1007/s10846-016-0396-2).
- [10] Keenan Albee, Monica Ekal, Rodrigo Ventura, and Richard Linares. Combining Parameter Identification and Trajectory Optimization: Real-time Planning for Information Gain. In *arXiv*. arXiv, June 2019.
- [11] Russ Tedrake. *Underactuated Robotics*. 2022.
- [12] P. Reist and R. Tedrake. Simulation-based LQR-trees with input and state constraints. In *2010 IEEE International Conference on Robotics and Automation*, pages 5504–5510, May 2010. DOI: [10.1109/ROBOT.2010.5509893](https://doi.org/10.1109/ROBOT.2010.5509893).
- [13] Philipp Reist, Pascal Preiswerk, and Russ Tedrake. Feedback-motion-planning with simulation-based LQR-trees. *The International Journal of Robotics Research*, 35(11):1393–1416, Sept. 2016. ISSN: 0278-3649, 1741-3176. DOI: [10.1177/0278364916647192](https://doi.org/10.1177/0278364916647192).
- [14] Joseph Moore, Rick Cory, and Russ Tedrake. Robust post-stall perching with a simple fixed-wing glider using LQR-Trees. *Bioinspiration & Biomimetics*, 9(2):025013, May 2014. ISSN: 1748-3182, 1748-3190. DOI: [10.1088/1748-3182/9/2/025013](https://doi.org/10.1088/1748-3182/9/2/025013).

- [15] Anirudha Majumdar and Russ Tedrake. Funnel libraries for real-time robust feedback motion planning. *The International Journal of Robotics Research*, 36(8):947–982, July 2017. ISSN: 0278-3649. DOI: [10.1177/0278364917712421](https://doi.org/10.1177/0278364917712421).
- [16] Russ Tedrake, Ian R. Manchester, Mark Tobenkin, and John W. Roberts. LQR-trees: Feedback Motion Planning via Sums-of-Squares Verification. *The International Journal of Robotics Research*, 29(8):1038–1052, July 2010. ISSN: 0278-3649, 1741-3176. DOI: [10.1177/0278364910369189](https://doi.org/10.1177/0278364910369189).
- [17] Mark M. Tobenkin, Ian R. Manchester, and Russ Tedrake. Invariant Funnels around Trajectories using Sum-of-Squares Programming. *IFAC Proceedings Volumes*, 44(1):9218–9223, Jan. 2011. ISSN: 1474-6670. DOI: [10.3182/20110828-6-IT-1002.03098](https://doi.org/10.3182/20110828-6-IT-1002.03098).
- [18] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek. Sequential Composition of Dynamically Dexterous Robot Behaviors. *The International Journal of Robotics Research*, 18(6):534–555, June 1999. ISSN: 0278-3649, 1741-3176. DOI: [10.1177/02783649922066385](https://doi.org/10.1177/02783649922066385).
- [19] Angel Flores-Abad, Ou Ma, Khanh Pham, and Steve Ulrich. A review of space robotics technologies for on-orbit servicing. *Progress in Aerospace Sciences*, 68:1–26, July 2014. ISSN: 03760421. DOI: [10.1016/j.paerosci.2014.03.002](https://doi.org/10.1016/j.paerosci.2014.03.002).
- [20] Christopher Trentlage, Juntang Yang, Mohamed Khalil Ben Larbi, Carlos Antonio De Alba Padilla, and Enrico Stoll. The ELISSA Laboratory: Free-Floating Satellites for Space-Related Research. In *67. Deutscher Luft- und Raumfahrtkongress*, Bonn, Germany, 2018. Deutsche Gesellschaft für Luft- und Raumfahrt - Lilienthal-Oberth e.V.
- [21] Juntang Yang, Yogesh Khedar, Mohamed Khalil Ben-Larbi, Jan Backhaus, Astrid Lampert, Ulf Bestmann, and Enrico Stoll. Concept and Feasibility Evaluation of Distributed Sensor-Based Measurement Systems Using Formation Flying Multicopters. *Atmosphere*, 12(7):874, 2021. DOI: [10.3390/atmos12070874](https://doi.org/10.3390/atmos12070874).
- [22] Rick E Cory. *Supermaneuverable Perching*. PhD thesis, Massachusetts Institute of Technology, Boston, June 2010.
- [23] Shubham Vyas, Lasse Maywald, Shivesh Kumar, Marko Jankovic, Andreas Mueller, and Frank Kirchner. Post-capture detumble trajectory stabilization for robotic active debris removal. *Advances in Space Research*, 2022. ISSN: 0273-1177. DOI: <https://doi.org/10.1016/j.asr.2022.09.033>.
- [24] Anton Bredenbeck, Shubham Vyas, Willem Suter, Martin Zwick, Dorit Borrmann, Miguel Olivares-Mendez, and Andreas Nüchter. Trajectory Optimization and Following for a Three Degrees of Freedom Overactuated Floating Platform. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, Oct 2022.
- [25] Dubowsky Steven and Papadopoulos Evangelos. The kinematics, dynamics, and control of free-flying and free-floating space robotic systems. *IEEE Transactions on robotics and automation*, 9(5):531–543, 1993. ISSN: 1042-296X.
- [26] Evangelos George Papadopoulos. *On the dynamics and control of space manipulators*. PhD thesis, Massachusetts Institute of Technology, 1990.
- [27] Markus Wilde, Stephen Kwok Choon, Alessio Grompone, and Marcello Romano. Equations of motion of free-floating spacecraft-manipulator systems: An Engineer’s tutorial. *Frontiers Robotics AI*, 5(APR):1–24, 2018. ISSN: 22969144. DOI: [10.3389/frobt.2018.00041](https://doi.org/10.3389/frobt.2018.00041).
- [28] Thai-Chau Nguyen-Huynh. *Adaptive Reactionless Control of a Space Manipulator for Post-Capture of an Uncooperative Tumbling Target*. PhD thesis, McGill University, 2013.
- [29] Kevin M. Lynch and Frank Park. *Modern Robotics -Mechanics, Planning, and Control*. Cambridge University Press, UK, 2017. ISBN: 0816057451. DOI: [10.1017/9781316661239](https://doi.org/10.1017/9781316661239).

- [30] John T. Betts. *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*. Society for Industrial and Applied Mathematics, Philadelphia, 2010. DOI: [10.1137/1.9780898718577](https://doi.org/10.1137/1.9780898718577).
- [31] Matthew Kelly. An introduction to trajectory optimization: How to do your own direct collocation. *SIAM Review*, 59(4):849–904, 2017. ISSN: 00361445. DOI: [10.1137/16M1062569](https://doi.org/10.1137/16M1062569).
- [32] Philip E. Gill, Walter Murray, and Michael A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Review*, 47(1):99–131, 2005.
- [33] Russ Tedrake and the Drake Development Team. *Drake: Model-based design and verification for robotics*. 2019.
- [34] Fariba Fahroo and I Michael Ross. Direct trajectory optimization by a chebyshev pseudospectral method. *Journal of Guidance, Control, and Dynamics*, 25(1):160–166, 2002.
- [35] Abolfazl Shirazi, Josu Ceberio, and Jose A Lozano. Spacecraft trajectory optimization: A review of models, objectives, approaches and solutions. *Progress in Aerospace Sciences*, 102:76–98, 2018.
- [36] Shubham Vyas, Bilal Wehbe, and Shivesh Kumar. Quaternion based LQR for Free-Floating Robots without Gravity. In *6th CEAS Conference on Guidance, Navigation, and Control, (EuroGNC)*, Berlin, May 2022. CEAS.
- [37] Dimitri P Bertsekas. *Dynamic programming and optimal control: Volume. 1*. Athena scientific, Nashua, US, 2012.
- [38] Hassan K. Khalil. *Nonlinear Systems*. Prentice Hall, Upper Saddle River, N.J, 3rd ed edition, 2002. ISBN: 978-0-13-067389-3.
- [39] Pablo Parrilo. *thesis.pdf*. PhD thesis, California Institute of Technology, Pasadena, California, 2000. <http://www.mit.edu/~parrilo/pubs/files/thesis.pdf>.
- [40] Esmaeil Najafi, Robert Babuška, and Gabriel A. D. Lopes. A fast sampling method for estimating the domain of attraction. *Nonlinear Dynamics*, 86(2):823–834, Oct. 2016. ISSN: 0924-090X, 1573-269X. DOI: [10.1007/s11071-016-2926-7](https://doi.org/10.1007/s11071-016-2926-7).
- [41] Lasse Maywald, Felix Wiebe, Shivesh Kumar, Mahdi Javadi, and Frank Kirchner. Co-optimization of acrobat design and controller for increased certifiable stability, 07 2022. DOI: [10.13140/RG.2.2.36436.07043](https://doi.org/10.13140/RG.2.2.36436.07043).
- [42] Werner Krauth. *Statistical mechanics: algorithms and computations*. Number 13 in Oxford master series in physics. Oxford University Press, Oxford, 2006. ISBN: 978-0-19-851535-7.
- [43] Yoji Umetani and Kazuya Yoshida. Resolved Motion Rate Control of Space Manipulators with Generalized Jacobian Matrix. *IEEE Transactions on Robotics and Automation*, 5(3):303–314, 1989. DOI: [10.1109/70.34766](https://doi.org/10.1109/70.34766).
- [44] Richard Zappulla, Josep Virgili-Llop, Costantinos Zagaris, Hyeongjun Park, and Marcello Romano. Dynamic Air-Bearing Hardware-in-the-Loop Testbed to Experimentally Evaluate Autonomous Spacecraft Proximity Maneuvers. *Journal of Spacecraft and Rockets*, 54(4):825–839, July 2017. ISSN: 0022-4650, 1533-6794. DOI: [10.2514/1.A33769](https://doi.org/10.2514/1.A33769).
- [45] Declan Jonckers, Oliver Tauscher, R. Aditya Thakur, and Lasse Maywald. Additive manufacturing of large structures using free-flying satellites. *Frontiers in Space Technologies (2022) 879542*. - <https://doi.org/10.3389/frspt.2022.879542>, Apr 2022. ISSN: 2673-5075. DOI: [10.3389/frspt.2022.879542](https://doi.org/10.3389/frspt.2022.879542).
- [46] Franek Stark, Shubham Vyas, Georg Schildbach, and Frank Kirchner. Mixed Integer Model Predictive Control for a free-floating platform with binary and continuous actuation. In *7th CEAS Conference on Guidance, Navigation, and Control, (EuroGNC)*, Bristol, June 2024. CEAS.
- [47] Richard Zappulla, Josep Virgili-Llop, and Marcello Romano. Spacecraft Thruster Control via Sigma–Delta Modulation. *Journal of Guidance, Control, and Dynamics*, 40(11):2928–2933, Nov. 2017. ISSN: 0731-5090, 1533-3884. DOI: [10.2514/1.G002986](https://doi.org/10.2514/1.G002986).