Digital-Twin-Enabled Multi-Spacecraft On-Orbit Operations

Sebastian Henao-Garcia*, Michael Kapteyn†, and Karen E. Willcox‡
The University of Texas at Austin, Austin, TX 78712, USA

Marco Tezzele§
Emory University, Atlanta, GA 30322, USA

Miguel Castroviejo-Fernandez¶, Taehyeun Kim‖, Michele Ambrosino**, Ilya Kolmanovsky††
University of Michigan, Ann Arbor, MI 48109-2140, USA

Himadri Basu‡‡, Piyush Jirwankar§§, and Ricardo G. Sanfelice¶¶
University of California Santa Cruz, Santa Cruz, CA 95060, USA

This paper proposes a digital twin formulation for on-orbit servicing operations involving multiple spacecraft operating in resource-constrained and uncertain scenarios. Our problem setup considers multiple spacecraft, each comprised of multiple subsystems, and develops a modular digital twin formulation that incorporates state estimation, information sharing, and compatible control strategies across the various subsystem digital twins, with a goal of driving the system-of-systems towards mission success. The proposed formulation is implemented and demonstrated for a simulated on-orbit servicing mission featuring a controllable, health-aware chaser spacecraft performing a rendezvous with an uncontrollable target spacecraft. The modular digital twin formulation handles uncertainty in the kinematic states as well as the health of the propulsion subsystem, while producing optimal control strategies that are robust to the considered modes of failure.

I. Introduction

Digital twins provide an integrated framework for translating observed data into control decisions via an asset-specific representation of the physical systems of interest and their unique characteristics [1–4]. A predictive digital twin integrates predictive models, data assimilation, model calibration, uncertainty quantification, and control algorithms to optimize the monitoring and operation of an autonomous system [5, 6]. Specifically, one begins with a predictive model of the system dynamics (either physics-based, data-driven, or a hybrid combination of both). As observations of the physical system are collected, parameters of the model and their associated uncertainty are calibrated in real-time so that the digital twin continually reflects the current understanding of the system's state. These calibrated models are then used to issue predictions about future system states, conditioned on past observations and future control decisions. Such a predictive capability can be leveraged to explore, validate, and/or optimize these control decisions via optimization or control-theoretic approaches.

In this work, we consider autonomous on-orbit spacecraft operations, such as on-orbit assembly, decommissioning and towing, and on-orbit refueling of satellites. These applications pose some unique challenges for digital twins. Consider the on-orbit servicing mission illustrated in Fig. 1. A potentially non-cooperative target spacecraft (e.g., space debris or a spacecraft that has undergone damage or system failure) needs to be towed to a disposal orbit by a controllable chaser spacecraft equipped with a robotic manipulator. As described in [7, 8], this close-proximity

*Ph.D. Student, Oden Institute for Computational Engineering and Sciences, AIAA Student Member
†Research Associate, Oden Institute for Computational Engineering and Sciences
‡Director, Oden Institute for Computational Engineering and Sciences, AIAA Fellow
§Assistant Professor, Department of Mathematics
¶Ph.D. Student, Department of Aerospace Engineering
‖Ph.D. Student, Department of Aerospace Engineering
**Postdoctoral Fellow, Department of Aerospace Engineering
††Professor, Department of Aerospace Engineering, AIAA Associate Fellow
‡‡Postdoc, Department of Electrical and Computer Engineering
§§Ph.D. Student, Department of Electrical and Computer Engineering
¶¶Professor, Department of Electrical and Computer Engineering

operation can be broken down into three main phases. In the pre-capture phase, the two spacecraft will be on potentially different orbits and the chaser must rendezvous with the target (Fig. 1a). In the capture phase, a docking maneuver must be carried out, starting with the two spacecraft separated by some distance and ending with the two spacecraft attached using the robotic arms on the chaser (Fig. 1b). Finally, in the post-capture phase the docked assembly will be towed to the disposal orbit while being actuated by the propulsion system on the chaser (Fig. 1c).

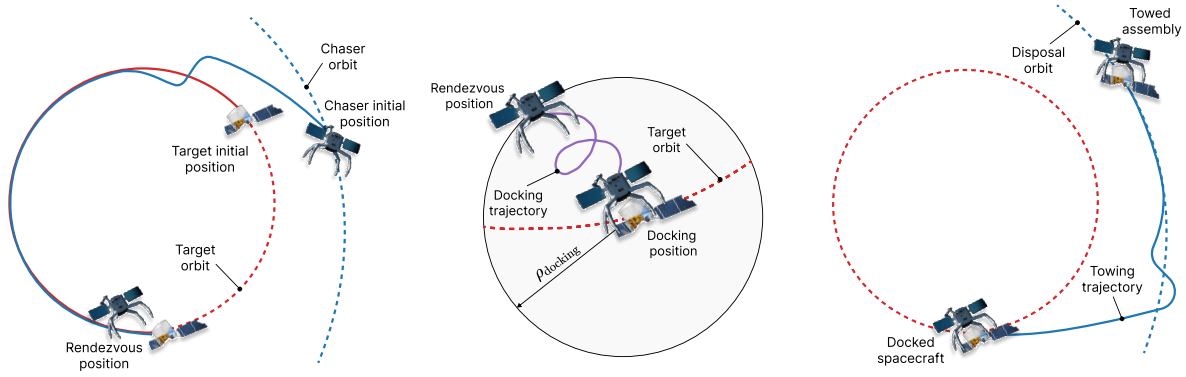

(a) Pre-capture phase. (b) Capture phase (Zoomed in). (c) Post-capture phase.

Fig. 1 An on-orbit servicing mission broken down into three phases: pre-capture, capture, and post-capture.

The success of the mission described above depends on precise control of the chaser spacecraft. This is challenged by uncertainty in the estimated position, velocity, and attitude of both the target and chaser spacecraft. These estimates are informed by noisy, temporally sparse, and potentially intermittent or corrupted measurements, thus increasing reliance on model-based predictive capabilities for the systems involved. Moreover, the control decisions of each spacecraft could depend on their instantaneous condition, for example, the health of the propulsion or power subsystems could necessitate the application of more conservative control inputs. The digital twin framework seeks to address these challenges by continually calibrating predictive models of each relevant subsystem in order to maintain real-time estimates of their state and associated uncertainty. However, it remains a challenge that the control algorithms must be capable of handling the uncertain state estimates provided by digital twins, as well as constraints resulting from subsystem health predictions.

The coupled dynamics between different spacecraft, or between the subsystems that compose a single spacecraft, presents a major challenge in both the digital twin design and the control design. On the digital twin side, there are modeling challenges. For example, the capture phase dynamics may involve complex contact physics that require a nonlinear high-fidelity contact model. When coupled with other aspects of the system models, this can lead to a monolithic coupled system of equations that is computationally expensive to solve. In addition, the models may be uncertain; for example, the post-capture dynamics of the spacecraft may depend on the exact docked configuration of the two spacecraft, which may be unknown *a priori*. Further, treating the virtual representation as a monolithic coupled set of nonlinear equations may lead to a digital twin for which continual dynamic updating becomes computationally intractable. On the control side, simultaneously determining control inputs for all coupled subsystems, while accounting for uncertainty in their current and future states, may lead to a large multi-objective optimization under uncertainty problem that becomes intractable. These challenges make a monolithic approach to digital-twin-based estimation and control not ideal for such complex and interconnected systems.

In this work we aim to overcome these challenges by developing a *modular* digital twin abstraction tailored to multi-spacecraft missions. Instead of considering a monolithic digital twin of the entire mission or scenario at-hand, we focus on constructing subsystem-level digital twins that can be coupled together to represent a single spacecraft, and coupled with subsystems of other spacecraft to represent a complete mission or scenario. The predictive dynamic models for each subsystem-level digital twin are tractable to derive, implement, update, and solve, while approximate coupling can be achieved via on-demand communication between the subsystem-level digital twins.

Section II presents the proposed digital twin formulation and discusses the associated challenges from a controls perspective. Section III applies the formulation to a simulated on-orbit rendezvous scenario between a controllable chaser and a non-cooperative target. We formulate the modular digital twins for this scenario, focusing on their coupling and interaction, and how their individual updates inform health-aware control inputs for the chaser spacecraft. Simulation

results demonstrate the health-aware rendezvous under uncertainty. Finally, Section IV concludes the paper.

II. Motivation: Modular Digital Twins

A digital twin, as defined in [1, 2], is an asset-specific set of computational models that mimic the structure, context, and behavior of a system of interest. These models are dynamically updated from observations of such a system, provide predictive capabilities beyond the current state, and ultimately inform decisions that realize value throughout the lifetime of the system. A key component of a digital twin is the bidirectional flow of information between the physical and virtual worlds, in terms of observations coming from the system and decisions being sent in return. Such a bidirectional interaction translates into improved monitoring, planning, and control for critical on-orbit spacecraft operations. Complex spacecraft operations often involve multiple interacting spacecraft, and each individual spacecraft is comprised of multiple interacting subsystems (structures, power, kinematics, propulsion, etc.). Motivated by this application, this section presents a modular digital twin formulation and outlines the need to develop methods for modeling, analyzing, and enabling interaction among multiple digital twin instances.

A. Modular Digital Twin Formulation for Multi-spacecraft Operations

Consider the on-orbit servicing scenarios depicted in Fig. 1, where a controllable chaser spacecraft, which is composed of multiple subsystems, interacts with a non-cooperative target spacecraft, also composed of multiple subsystems. Depending on the phase of the mission, these subsystems may be coupled in the sense that the decisions made by one subsystem will affect the state of another subsystem. Consider, for instance, how a sudden maneuver directed by the kinematics subsystem might require extended thrust from the propulsion subsystem, thus drawing more power and generating more heat on-board the spacecraft. If we consider building digital twins of these interacting spacecraft, then it becomes clear that the digital twin abstraction must be capable of handling the coupling and interconnection among the different subsystems.

One approach to achieving this would be to create a single monolithic digital twin that includes the state of all involved subsystems. However, this would require a model of the fully-coupled dynamics, something that is typically expensive to solve and difficult or impossible to update in real-time operational settings. On the other hand, employing a set of subsystem-specific models is typically more tractable, and provides a modular formulation to augment with additional subsystem models as needed. Thus, we instead seek a *modular* digital twin formulation that leverages each subsystem's models to achieve tractability and scalability.

To this end, we propose a formulation in terms of subsystem-level digital twins as depicted in Fig. 2. In this setting, each subsystem composing the spacecraft (such as kinematics, propulsion, thermal, etc.) is represented as a single digital twin instance whose underlying models are updated given subsystem-specific prediction, observation, and control frequencies. The chaser spacecraft's subsystem-level digital twins feature a bidirectional flow of information with their physical counterparts, which can be commanded with control inputs. The target spacecraft's subsystem-level digital twins, on the other hand, rely only on limited observations of the evolution of the target, such as tracking and localization data, and may not be able to communicate control decisions to their physical counterparts. Modeling the interaction between subsystem-level digital twins is an application-dependent choice that responds to the required level of modeling fidelity. For instance, in some scenarios, sharing information on-demand via polling among digital twin instances will suffice, while in others, formulating coupled dynamics and controls problems may be required instead.

Since each subsystem is represented as a single digital twin instance, we can leverage previous work on the mathematical abstraction of single digital twins and extend the representation to consider the aggregate of interacting subsystems during a given mission phase. This approach results in the mathematical abstraction for modular digital twins summarized in Table 1. This abstraction considers the six interacting quantities proposed by Kapteyn et al. to define a digital twin [5], here extended to be defined for each interacting subsystem with label ℓ.

Consider the lifespan of a mission discretized as $[t_1, t_2, ..., t_k, ..., t_{\text{final}}]$, where t_k denotes the time at timestep k and t_{final} denotes the mission's final time. Let $\boldsymbol{s}_k^\ell$ represent the physical state of the subsystem ℓ at time $t = t_k$. In this setting, the physical state $\boldsymbol{s}_k^\ell$ evolves according to the dynamics of the physical system and is indirectly observed via observations $\boldsymbol{o}_k^\ell$. At time t_k, estimates of the physical state are computed from the observations, thus producing the digital state $\boldsymbol{d}_k^\ell$, which can then be used to compute quantities of interest $\boldsymbol{q}_k^\ell$ (e.g., system performance characteristics that depend upon the digital state estimate). Furthermore, the predictive capabilities of the calibrated computational models that define the digital twin can be used to make predictions about the future evolution of the digital state, i.e., $\boldsymbol{d}_{k+1}^\ell, ..., \boldsymbol{d}_{\text{final}}^\ell$. The estimated and predicted digital states and quantities of interest inform current and future control

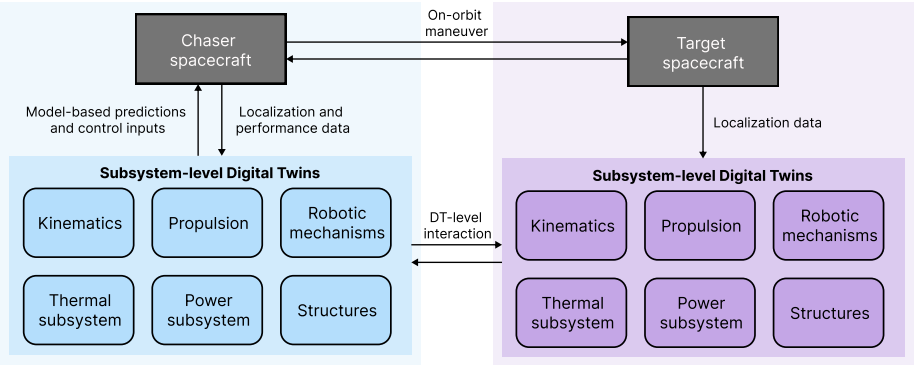

Fig. 2 Multi-spacecraft modular digital twin formulation. Each spacecraft is represented as the aggregate of subsystem-level digital twins that interact to achieve mission success.

Table 1 Abstraction of the modular digital twin formulation, adapted from [5]. All quantities defined for each interacting subsystem with label ℓ at time $t = t_k$. Uppercase script letters denote the space where each quantity in the abstraction evolves, and lowercase letters denote realizations of such quantities.

Quantity	Notation	Description
Physical State	$\boldsymbol{s}_k^\ell \in \mathcal{S}^\ell$	True parameterized state of the subsystem
Digital State	$\boldsymbol{d}_k^\ell \in \mathcal{D}^\ell$	Estimate of the parameterized state of the subsystem
Quantities of Interest (QoIs)	$\boldsymbol{q}_k^\ell \in \mathcal{Q}^\ell$	Additional quantities describing the subsystem that are estimated from the digital state and model outputs
Observations	$\boldsymbol{o}_k^\ell \in \mathcal{O}^\ell$	Measurements that are informative of the true state of the subsystem
Control Inputs	$\boldsymbol{u}_k^\ell \in \mathcal{U}^\ell$	Commanded actions or decisions that influence the evolution of the subsystem's state
Reward	$\boldsymbol{r}_k^\ell \in \mathcal{R}^\ell$	Mission goals and objective functions that quantify the performance and costs of the subsystem

decisions $\boldsymbol{u}_k^\ell, \boldsymbol{u}_{k+1}^\ell, \ldots$ that drive the evolution of the physical system subject to maximizing desired rewards $\boldsymbol{r}_k^\ell, \boldsymbol{r}_{k+1}^\ell, \ldots$. So far, the presented abstraction has assumed that each subsystem is represented by an independent digital twin that can be updated independently. For cases when coupling between subsystems is involved, communication occurs on-demand between the digital twin instances in order to share the most recent digital state estimate at time t_k. Consider, for instance, the scenario where the control input of subsystem ℓ depends on another subsystem ℓ', then the state of system ℓ' can be queried for the purpose of evaluating $\boldsymbol{u}_k^\ell$ as a function of both $\boldsymbol{d}_k^\ell$ and $\boldsymbol{d}_k^{\ell'}$. Such a flow of information hence contributes to the expressivity and value of the mission-level digital twin abstraction.

There are additional components of a digital twin that enable the evolution of the quantities underlying the abstraction outlined in Table 1. Firstly, a predictive model describing the evolution of the physical state is needed. The level of fidelity of such a model needs to be tailored to the application-specific context of each mission phase, allowing for first-principle models of the subsystem's relevant dynamics, physics-aware reduced-order models, and surrogate models derived via data-driven approaches. Secondly, a data assimilation strategy is needed to solve the inverse problem of estimating the states from observations. In the context of spacecraft operations, sequential state estimation (i.e., filtering) techniques are suitable for performing this task. Note that the concept of physical state in the context of digital twins extends the traditional concept of system state; in particular, the physical state, as defined in Table 1, allows for any parametrization of the subsystem's condition. Hence it can include, but is not limited to, the state vector arising in a state-space representation of a given dynamical system, which is commonly the state to be estimated using filtering techniques. Lastly, a control strategy is needed to drive the evolution of the physical state towards mission success. For example, optimal control techniques can be used to produce control policies that maximize given reward functions. Coupling challenges arise when formulating such control problems in a modular setting, as discussed in Section II.B.

B. Control Approaches for Modular Digital Twin Systems

A modular system, in contrast with a monolithic one, may involve multiple controllers each tasked with controlling a specific subsystem. Moreover, each of these controllers may not have access to the full system state, and may not be aware of other controllers. These challenges mean that controller design in the context of the modular digital twin formulation proposed in Section II.A requires careful consideration. Fortunately, such challenges have long been faced in the controls community when attempting to design controllers for large, complex, and potentially distributed systems, and have resulted in the formulation of a wide range of control strategies. Here we compare and contrast several of these approaches, as defined in [9]. In a centralized approach (Fig. 3a), a single controller receives system-wide information and is tasked with defining control inputs for all interacting subsystems, yet this approach has limited scalability due to the complexity of the coupled system dynamics, and the number of control variables to be determined in the optimization. Instead, in a distributed setting there are multiple controllers, where each is in charge of a specific subsystem and has access to partial information about the overall system. Ideally, these controllers would actively communicate with their neighbors to determine control inputs (Fig 3b), but there are also scenarios where neighboring controllers are not able to cooperate (Fig. 3c) and their commanded control inputs may be opposing due to being completely unaware of other subsystems. Additionally, a hierarchical scheme may be in place (Fig. 3d), where a supervisor agent would receive system-wide information and is charged with providing control objectives to each subsystem-level controller, which can cooperate to fulfill these objectives.

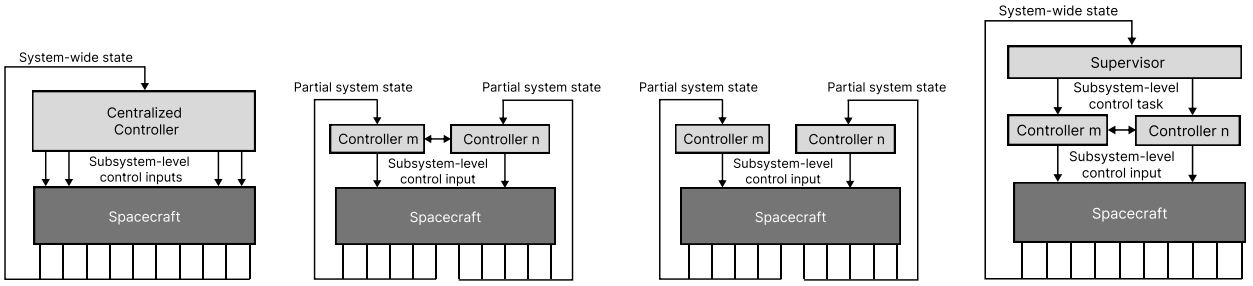

(a) Centralized. (b) Distributed with communication (c) Distributed and decentralized. (d) Distributed and hierarchical.

Fig. 3 Different control strategies for modular digital twins adapted from [9] In contrast with the monolithic approach, subsystem-level digital twins may not have access to system-wide state information. Thus, a supervisory, distributed, or fully decentralized control approach may be required.

The question that remains is how to tailor the control schemes of the interacting subsystems on each mission phase to work in the context of modular digital twins in support of the overall success of on-orbit servicing missions. Feedback-based control schemes, while fundamental, are excellent candidates for performing the intended on-orbit operations and can be naturally integrated into the digital twin formulation [10, 11]. Considering that a digital twin by definition features continual state estimation and updating from indirect observations, a feedback control law would be a complementary approach for controls, with such state estimates being used to inform the control inputs required to achieve the desired performance measured by a given reward function. There are multiple implementations of feedback-based control strategies that are suitable for integrating with modular digital twins in the context of on-orbit spacecraft operations. For instance, the Linear Quadratic Regulator (LQR) has been demonstrated to perform well when augmented with a supervisory control scheme to coordinate control algorithms for the different mission phases outlined in Fig. 1, as described in [8]. Alternatively, Model Predictive Control (MPC) could also be complementary to digital twins. MPC-based techniques could leverage the continually calibrated predictive models with quantified uncertainty that encompass the digital twin to define a feedback control law through the solution of an optimal control problem subject to state and control constraints [12]. Therefore, the predictive capabilities of digital twins combined with the optimality of MPC are an opportunity to realize guidance and control for on-orbit space operations under stringent and time-evolving state and control constraints.

III. Application: Digital-twin-enabled Rendezvous for On-orbit Servicing

This section presents an example of modular digital twins for an on-orbit servicing mission. Section III.A introduces the simulated pre-capture mission scenario. In Section III.B we formulate the problem into individual subsystem-level digital twins. A summary of the mission-level integration is presented in Section III.C, leading to the simulation results presented in Section III.D.

A. Problem Description

For this demonstration we consider the initial pre-capture phase of a simulated on-orbit servicing mission, as depicted in Fig 1a. Following [7], both spacecraft are modeled as point masses featuring two degrees of freedom for in-plane motion. Our goal is to control the chaser spacecraft to minimize its relative position and velocity with respect to the target spacecraft, while considering the individual orbital motion of each spacecraft around the Earth. In addition, we incorporate the task of monitoring and accounting for the health of the chaser's thrusters. We model this scenario via the modular digital twin formulation outlined in Section II via three interacting subsystems: (1) the target spacecraft kinematics, (2) the chaser spacecraft kinematics, and (3) the chaser's propulsion subsystem. The position and velocity of each spacecraft are estimated based on range and bearing angle measurements. Estimates of the thrusters' health and the presence of failure in the propulsion subsystem are inferred based on the deviation between the predicted and measured kinematics, with the commanded thrust adjusted accordingly to minimize the error in the intended rendezvous trajectory. The unifying component among these digital twin instances is the computation of a health-aware control input for the on-orbit chaser spacecraft, which is informed by both the relative position and velocity between the two spacecraft, and the current health estimates of the propulsion subsystem.

B. Subsystem-level Digital Twin Formulation

The following subsections introduce the prediction, state estimation, and control approaches that will comprise the digital twin instances of the interacting subsystems in the pre-capture phase. Despite the specific choices made regarding the spacecraft dynamics model, estimation algorithm, and controls technique, recall that this is a modular implementation and such models and algorithms can be replaced with higher fidelity versions aligned with the modeling requirements of the digital twin implementation.

1. Kinematics

Throughout this section, all quantities denoted with the superscript ℓ apply to either the target or the chaser spacecraft, i.e., $\ell \in \{\text{target}, \text{chaser}\}$. Let $\boldsymbol{s}_k^\ell$ represent the physical state of the kinematics subsystem of each spacecraft. The target kinematics subsystem digital twin is concerned with tracking the position and velocity of the target spacecraft, hence $\boldsymbol{s}_k^{\text{target}} \coloneqq \boldsymbol{x}_k^{\text{target}} \in \mathbb{R}^4$ represents the physical state of the target spacecraft kinematics, where $\boldsymbol{x}_k^\ell = \left[x_k^\ell,\ y_k^\ell, \dot{x}_k^\ell,\ \dot{y}_k^\ell\right]$ is the kinematic state vector containing the position and velocity components of the spacecraft in an Earth-centered, Earth-fixed frame. The chaser kinematics digital twin, in addition to tracking the position and velocity of the spacecraft, is also tasked with estimating an acceleration bias term that could be indicative of the presence of a failure in the propulsion subsystem. Let $\boldsymbol{s}_k^{\text{chaser}} \coloneqq \left[\boldsymbol{x}_k^{\text{chaser}},\ \boldsymbol{\tau}_k^{\text{chaser}}\right]^\top \in \mathbb{R}^6$ represent the physical state of the chaser spacecraft kinematics, which augments the kinematic state vector $\boldsymbol{x}_k^{\text{chaser}} \in \mathbb{R}^4$ with the acceleration bias vector $\boldsymbol{\tau}_k^{\text{chaser}} \in \mathbb{R}^2$ containing the bias components. The acceleration bias term is described further in Section III.B.2.

Let $\boldsymbol{u}_k^\ell$ be the control input to each of the spacecraft. Since the target spacecraft is assumed to be non-cooperative, it has no control inputs and this quantity only applies to the chaser spacecraft. The control input for the chaser spacecraft $\boldsymbol{u}_k^{\text{chaser}}$ is discussed in Section III.B.3. For observations, we consider discrete and noisy measurements of the range ρ and bearing angle θ for both spacecraft separately, i.e., $\boldsymbol{o}_k^\ell \coloneqq \left[\rho_k^\ell, \theta_k^\ell\right]^\top \in \mathbb{R}^2$, expressed in an Earth-centered, Earth-fixed frame.

Define the digital state $\boldsymbol{d}_k^\ell$ as an estimate of the physical state $\boldsymbol{s}_k^\ell$ of the kinematics subsystem of each spacecraft. A key aspect of the digital twin abstraction we adopt from [5] is the treatment of uncertainty. We represent the digital state probabilistically via the random variable

$$D_k^\ell \sim p\left(\boldsymbol{d}_k^\ell\right) \coloneqq \mathcal{N}\left(\hat{\boldsymbol{s}}_k^\ell, P_k^\ell\right), \tag{1}$$

where $\mathcal{N}$ denotes a Gaussian distribution whose mean $\hat{\boldsymbol{s}}_k^\ell$ corresponds to the maximum a posteriori (MAP) estimate of the corresponding spacecraft's physical state $\boldsymbol{s}_k^\ell$, and its covariance P_k^ℓ corresponds to the state estimate covariance [13].

The estimate of the digital state is computed via a standard Extended Kalman Filter (EKF) setup [14], where prediction and correction steps are alternated to maintain a posterior belief of the digital state as new observations are acquired. In this context, process noise is added to the predictive dynamics to model uncertainty in our prediction of the physical state evolution, thus introducing the random variable

$$S_{k+1}^{\ell} \sim p\big(\boldsymbol{s}_{k+1}^{\ell}\big), \quad S_{k+1}^{\ell} \coloneqq \boldsymbol{f}^{\ell}\big(\Delta t_k^{\ell},\, \boldsymbol{s}_k^{\ell},\, \boldsymbol{u}_k^{\ell},\, \boldsymbol{V}_k^{\ell}\big), \tag{2}$$

where $\boldsymbol{f}^{\ell}$ represents the discrete-time function capturing nonlinear orbital dynamics of the spacecraft (presented in Appendix A.1.1 Eq. 23 and Appendix A.1.2 Eq. 25 for the target and chaser spacecraft, respectively), $\Delta t_k^{\ell} = t_{k+1}^{\ell} - t_k^{\ell}$ is the integration time step, and $\boldsymbol{V}_k^{\ell} \sim \mathcal{N}\big(\boldsymbol{0}, Q_k^{\ell}\big)$ represents the discrete-time process noise with covariance matrix Q_k^{ℓ}. Additionally, the EKF incorporates uncertainty in the observed data via the measurement model

$$O_k^{\ell} \sim p\big(\boldsymbol{o}_k^{\ell}\big), \quad O_k^{\ell} \coloneqq \boldsymbol{h}\big(\boldsymbol{s}_k^{\ell}\big) + \boldsymbol{W}_k^{\ell}, \tag{3}$$

where $\boldsymbol{W}_k^{\ell} \sim \mathcal{N}\big(\boldsymbol{0}, R_k^{\ell}\big)$ represents measurement noise with covariance matrix R_k^{ℓ}, and $\boldsymbol{h}$ denotes the nonlinear state-to-measurement map

$$\boldsymbol{h}\big(\boldsymbol{s}_k^{\ell}\big) \coloneqq \begin{bmatrix} \sqrt{\big(x_k^{\ell}\big)^2 + \big(y_k^{\ell}\big)^2} \\ \arctan\left(\frac{y_k^{\ell}}{x_k^{\ell}}\right) \end{bmatrix}. \tag{4}$$

In the prediction step, we seek to update our belief for the digital state at $t = t_{k+1}$ via the discrete-time dynamics model. The EKF uses a first-order linearization of the dynamics around the MAP state estimate $\boldsymbol{s}_k^{\ell} = \hat{\boldsymbol{s}}_k^{\ell}$ and the mean of the process noise $V_k^{\ell} = \boldsymbol{0}$, resulting in the predicted digital state

$$\bar{D}_{k+1}^{\ell} \sim p\big(\boldsymbol{d}_{k+1}^{\ell} \mid D_k^{\ell},\, \boldsymbol{u}_k^{\ell}\big) \coloneqq \mathcal{N}\left(\bar{\boldsymbol{s}}_{k+1}^{\ell}, \bar{P}_{k+1}^{\ell}\right), \tag{5}$$

where the mean $\bar{\boldsymbol{s}}_{k+1}^{\ell}$ corresponds to the predicted state obtained by propagating the MAP state estimate $\hat{\boldsymbol{s}}_k^{\ell}$ as

$$\bar{\boldsymbol{s}}_{k+1}^{\ell} = \mathbb{E}\left[S_{k+1}^{\ell} \mid D_k^{\ell},\, \boldsymbol{u}_k^{\ell}\right] = \boldsymbol{f}^{\ell}\left(\Delta t_k^{\ell}, \hat{\boldsymbol{s}}_k^{\ell},\, \boldsymbol{u}_k^{\ell}, \boldsymbol{0}\right), \tag{6}$$

and the covariance matrix is obtained by propagating the state covariance matrix P_k^{ℓ} via the state-transition matrix F_k^{ℓ}, and the process noise covariance Q_k^{ℓ} via the process noise-transition matrix Γ_k^{ℓ} as

$$\bar{P}_{k+1}^{\ell} = \mathrm{Var}\left[S_{k+1}^{\ell} - \bar{\boldsymbol{s}}_{k+1}^{\ell} \mid D_k^{\ell},\, \boldsymbol{u}_k^{\ell}\right] = F_k^{\ell}\, P_k^{\ell}\, F_k^{\ell\top} + \Gamma_k^{\ell}\, Q_k^{\ell}\, \Gamma_k^{\ell\top}. \tag{7}$$

We can follow this approach to predict the future evolution of the kinematics digital twins before a new measurement becomes available. Uncertainty in these predictions is quantified through the propagated state covariance matrices.

The correction step occurs whenever a new measurement $\boldsymbol{o}_{k+1}^{\ell}$ becomes available. The posterior belief for the digital state is updated as

$$D_{k+1}^{\ell} \sim p\big(\boldsymbol{d}_{k+1}^{\ell} \mid \bar{D}_{k+1}^{\ell},\, \boldsymbol{o}_{k+1}^{\ell}\big) \coloneqq \mathcal{N}\left(\hat{\boldsymbol{s}}_{k+1}^{\ell}, P_{k+1}^{\ell}\right), \tag{8}$$

where the mean is given by correcting the predicted state $\bar{\boldsymbol{s}}_{k+1}^{\ell}$ with information from the new observation $\boldsymbol{o}_{k+1}^{\ell}$ as

$$\hat{\boldsymbol{s}}_{k+1}^{\ell} = \mathbb{E}\left[S_{k+1}^{\ell} \mid \bar{D}_{k+1}^{\ell}, \boldsymbol{o}_{k+1}^{\ell}\right] = \bar{\boldsymbol{s}}_{k+1}^{\ell} + \mathcal{K}_{k+1}^{\ell}\, \boldsymbol{\upsilon}_{k+1}^{\ell}, \tag{9}$$

with $\mathcal{K}_{k+1}^{\ell}$ denoting the Kalman gain matrix as defined by [14], and $\boldsymbol{\upsilon}_k$ denoting the innovation, i.e., the residual between the received measurement and the predicted measurement,

$$\boldsymbol{\upsilon}_{k+1}^{\ell} = \boldsymbol{o}_{k+1}^{\ell} - \bar{\boldsymbol{o}}_{k+1}^{\ell}, \tag{10}$$

where $\bar{\boldsymbol{o}}_{k+1}^{\ell}$ denotes the expected measurement corresponding to the predicted state $\bar{\boldsymbol{s}}_{k+1}^{\ell}$, i.e.,

$$\bar{\boldsymbol{o}}_{k+1}^{\ell} = \mathbb{E}\left[O_{k+1}^{\ell} \mid \bar{D}_{k+1}^{\ell}\right] = \boldsymbol{h}\big(\bar{\boldsymbol{s}}_{k+1}^{\ell}\big). \tag{11}$$

The posterior covariance matrix is corrected by accounting for uncertainty in the measurement model as

$$P_{k+1}^{\ell} = \mathrm{Var}\left[S_{k+1}^{\ell} - \hat{\boldsymbol{s}}_{k+1}^{\ell} \mid \bar{D}_{k+1}^{\ell}, \boldsymbol{o}_{k+1}^{\ell}\right] = \left[I - \mathcal{K}_{k+1}^{\ell}\, H_{k+1}^{\ell}\right] \bar{P}_{k+1}^{\ell}, \tag{12}$$

where I denotes the identity matrix, and H_{k+1}^{ℓ} denotes the measurement model Jacobian evaluated at $\bar{\boldsymbol{s}}_{k+1}^{\ell}$.

We compute the target's orbital rate $\hat{n}_k$ as a quantity of interest for the target's kinematics based on the current MAP estimate of its physical state,

$$q_k^{\text{target}} \coloneqq \hat{n}_k = \sqrt{\frac{\mu}{\hat{a}_k^3}}, \quad \hat{a}_k = \left(\frac{2}{r_k} - \frac{v_k^2}{\mu}\right)^{-1}, \tag{13}$$

where μ is the Earth's gravitational constant, $\hat{a}_k$ denotes the current estimate of the target's orbit semi-major axis, and $\hat{r}_k, \hat{v}_k$ are the target's position and velocity magnitudes computed from its current MAP digital state estimate $\hat{\boldsymbol{s}}_k^{\text{target}}$.

Additionally, define the innovation presented in Eq. 10 as a quantity of interest for the chaser's kinematics that is informative of the deviation between the predicted and observed dynamics, i.e.,

$$\boldsymbol{q}_k^{\text{chaser}} \coloneqq \boldsymbol{\upsilon}_k^{\text{chaser}} = \boldsymbol{o}_k^{\text{chaser}} - \bar{\boldsymbol{o}}_k^{\text{chaser}}. \tag{14}$$

where $\bar{\boldsymbol{o}}_k^{\text{chaser}} = \boldsymbol{h}\left(\bar{\boldsymbol{s}}_k^{\text{chaser}}\right)$ is the measurement estimate obtained with the current prior belief of the digital state, as defined in Eq. 11.

2. Propulsion Subsystem Health Monitoring

We define the propulsion subsystem of the chaser spacecraft to be equipped with a set of four thrusters positioned such that they are capable of applying an arbitrary in-plane thrust vector, which in turn accelerates the spacecraft along the desired direction. Assume that each thruster is continually aligned with one of the axes in the Earth-centered, Earth-fixed frame {+X, +Y, -X, -Y}. We consider a mode of failure in the propulsion subsystem including a constant leak of propellant from the tank to the mounted thrusters and a misalignment in the thrusters' nominal pointing directions. This failure results in a constant additive acceleration bias that affects the motion of the chaser spacecraft independently of the commanded acceleration magnitude. We design a propulsion subsystem health monitoring digital twin concerned with estimating whether the spacecraft is operating in such a failure condition and which thruster is the most likely to be experiencing the failure. Define the the physical state for this subsystem as $\boldsymbol{s}_k^{\text{propulsion}} \coloneqq [f_k, i_k]^\top$, where

$$f_k \coloneqq \begin{cases} 0 \Leftrightarrow \text{No current failure,} \\ 1 \Leftrightarrow \text{The system is experiencing a failure,} \end{cases} \tag{15}$$

denotes the failure status at time $t = t_k$ and

$$i_k \in \begin{cases} \{0\} & \text{if } f_k = 0 \\ \{\text{+X, +Y, -X, -Y}\} & \text{otherwise ,} \end{cases} \tag{16}$$

identifies the thruster that is most likely to be faulty among all four available thrusters once the failure has been detected. Indirect observations of this state are available via both the MAP estimate of the acceleration bias and the innovation coming from the chaser kinematics digital twin, hence we define the observations $\boldsymbol{o}_k^{\text{propulsion}} \coloneqq \left[\hat{\boldsymbol{\tau}}_k^{\text{chaser}}, \boldsymbol{\upsilon}_k^{\text{chaser}}\right]^\top$.

Let the digital state be defined as the estimate of the physical state, i.e., $\boldsymbol{d}_k^{\text{propulsion}} \coloneqq \left[\hat{f}_k, \hat{i}_k\right]^\top$, which we seek to compute via an online Fault Detection, Isolation, and Recovery (FDIR) algorithm adapted from [15]. Consider a sequence length w of the most recent innovations $\boldsymbol{\upsilon}_{k-w}^{\text{chaser}}, ..., \boldsymbol{\upsilon}_k^{\text{chaser}}$ in the chaser's kinematic state estimation algorithm. This sequence of observed innovations corresponds to a realization of the sequence of independent random variables $Z_1^{\text{chaser}}, ..., Z_k^{\text{chaser}}$. In the absence of a failure, it can be shown that the random sequence is identically distributed with $Z_j^{\text{chaser}} \sim \mathcal{N}(\mathbf{0}, R^{\text{chaser}}),\ \forall j \in \{k-w, ..., k\}$, which follows from the definition of innovation given in Eq. 10 and the fact that the additive measurement noise is Gaussian with mean $\mathbf{0}$ and variance R^{chaser}. In the presence of a failure, however, the random sequence will no longer be identically distributed, and the mean of some Z_j^{chaser} will shift to an unknown nonzero value, i.e., $\exists\, j : \mathbb{E}\left[Z_j^{\text{chaser}}\right] \neq \mathbf{0},\ j \in \{k-w, ..., k\}$. Hence, an estimate of the the health of the propulsion subsystem in terms of the presence of a failure can be inferred from a sequence of innovations via the hypothesis test for the mean of the variables in the random sequence

$$\begin{aligned} \boldsymbol{H}_0 &: \forall\, j \in \{k-w, ..., k\},\ \ \mathbb{E}\left[Z_j^{\text{chaser}}\right] = \mathbf{0}, \\ \boldsymbol{H}_1 &: \exists\, j \in \{k-w, ..., k\},\ \ \mathbb{E}\left[Z_j^{\text{chaser}}\right] \neq \mathbf{0}. \end{aligned} \tag{17}$$

Formulating this problem in the context of the Generalized Likelihood Ratio algorithm [16], our goal is to identify the combination of timestep index j and mean vector $\bar{\boldsymbol{\upsilon}}$ that will yield the greatest cumulative sum of the log-likelihood ratios over a subsequence starting at index j. We introduce the test statistic

$$\lambda_k = \max_{k-w\le j\le k} \sup_{\bar{\boldsymbol{\upsilon}}\neq\mathbf{0}} \sum_{i=j}^{k} \ln \frac{L\left(\bar{\boldsymbol{\upsilon}} \mid \boldsymbol{\upsilon}_i^{\text{chaser}}\right)}{L\left(\mathbf{0} \mid \boldsymbol{\upsilon}_i^{\text{chaser}}\right)} = \max_{k-w\le j\le k} \frac{k-j+1}{2}\, \hat{\boldsymbol{\upsilon}}(j)^{\top} \text{inv}\left(R^{\text{chaser}}\right)\hat{\boldsymbol{\upsilon}}(j), \tag{18}$$

where $L\left(\bar{\boldsymbol{\upsilon}} \mid \boldsymbol{\upsilon}_i^{\text{chaser}}\right)$ denotes the likelihood of observing $\boldsymbol{\upsilon}_i^{\text{chaser}}$ given $\mathbb{E}\left[Z_i^{\text{chaser}}\right] = \bar{\boldsymbol{\upsilon}}$. Since the measurement noise is Gaussian, there exists an exact solution to the inner maximization problem over $\bar{\boldsymbol{\upsilon}}$ that yields the right-most equality with $\hat{\boldsymbol{\upsilon}}(j) = 1/(k-j+1)\sum_{i=j}^{k} \boldsymbol{\upsilon}_i^{\text{chaser}}$. The outer maximization problem over j can be solved by brute force since the window size w is small. Comparing the test statistic λ_k against a set threshold value ε, produces the failure status estimate

$$\hat{f}_k = \begin{cases} 0 \;\text{ if } \lambda_k < \varepsilon & \Rightarrow \text{accept } \boldsymbol{H}_0 \Leftrightarrow \text{No current failure,} \\ 1 \;\text{ otherwise} & \Rightarrow \text{reject } \boldsymbol{H}_0 \;\Leftrightarrow \text{The system is experiencing a failure.} \end{cases} \tag{19}$$

Finally, the failure is isolated by identifying the thruster that is most likely to be experiencing the failure, thus producing an estimate $\hat{i}_k$. We use the scalar projection to compute the angle that the current acceleration bias estimate makes with each of the thrusters' pointing directions. This angle serves as a measure of each thruster's likelihood of being faulty. In this context, the smaller the angle, the higher the likelihood of being faulty.

3. Controls

Let $\boldsymbol{u}_k^{\text{chaser}} \in \mathbb{R}^2$ be the control input for the chaser corresponding to an commanded acceleration described in the Earth-centered, Earth-fixed frame. We adopt a feedback control law (described in Appendix A.2, Eq. 29) designed to achieve rendezvous with the target. In particular, the implemented feedback controller applies a pre-computed controller gain to minimize the relative kinematic state between the chaser and the target spacecraft, i.e., maximizing the reward $r_k^{\text{chaser}} = -\|\hat{\boldsymbol{x}}_k^{\text{chaser}} - \hat{\boldsymbol{x}}_k^{\text{target}}\|$, with the minimum cost over the entire chaser's trajectory. In this setting, cost is defined as the cumulative magnitude of the applied accelerations, which is a measure of the fuel required to complete the maneuver.

We augment the feedback control system with thruster health-awareness, informed by the propulsion subsystem health monitoring digital twin described in Section III.B.2. After detecting and isolating a failure in the propulsion subsystem, a recovery strategy can be implemented in the control strategy to produce a health-aware acceleration given by

$$\boldsymbol{u}_k^{\text{propulsion}} = \begin{cases} \boldsymbol{u}_k^{\text{chaser}} & \text{if } f_k = 0 & \Leftrightarrow \text{No need to compensate for the bias,} \\ \boldsymbol{u}_k^{\text{chaser}} + \hat{\boldsymbol{\tau}}_k & \text{if } f_k = 1 \text{ and } \|\hat{\boldsymbol{\tau}}_k\| \le \tau_{\max} & \Leftrightarrow \text{Compensate for the bias if operating in failure,} \\ \boldsymbol{u}_k^{\text{abort}} & \text{if } f_k = 1 \text{ and } \|\hat{\boldsymbol{\tau}}_k\| > \tau_{\max} & \Leftrightarrow \text{Abort mission,} \end{cases} \tag{20}$$

where $\tau_{\max}$ denotes a threshold for the acceleration bias magnitude, and $\boldsymbol{u}_k^{\text{abort}}$ denotes a commanded acceleration intended to abort the rendezvous mission when the failure is deemed too critical and instead bring the chaser spacecraft to a safe location.

C. Mission-level Integration of the Interacting Digital Twin Instances

Table 2 summarizes the previously defined components, giving rise to our mission-level modular digital twin formulation to enable a controllable chaser spacecraft to perform a rendezvous with a non-cooperative target via control inputs that are aware of the propulsion subsystem's health.

In our setup, each subsystem digital twin is able to evolve independently based on specified prediction, observation, and control frequencies. However, all digital twins can be updated to a given time (via their respective prediction models) and the current state information can be shared across the independent digital twin instances to enable the coupling required when computing a control input for the chaser spacecraft. To evaluate the controller that produces the chaser acceleration $\boldsymbol{u}_k^{\text{chaser}}$, current estimates of both the chaser's and target's kinematic states, i.e, $\boldsymbol{x}_k^{\text{chaser}}$ and $\boldsymbol{x}_k^{\text{target}}$, respectively, are required to compute the relative kinematic state vector. Also, an estimate of the target's orbital rate n_k^{target} is required to define the linearized relative motion dynamics used to compute the controller gain matrix. To compute a health-aware control input $\boldsymbol{u}_k^{\text{propulsion}}$, a sequence of length w of the most recent innovations $\boldsymbol{\upsilon}_{k-w}^{\text{chaser}}, \ldots, \boldsymbol{\upsilon}_k^{\text{chaser}}$

Table 2 Summary of the interacting subsystem-level digital twins in the simulated on-orbit rendezvous.

Quantity	Notation	Subsystem: Target's kinematics (ℓ = target)	Subsystem: Chaser's kinematics (ℓ = chaser)	Subsystem: Propulsion's health (ℓ = propulsion)
Physical State	$\boldsymbol{s}_k^\ell \in \mathcal{S}^\ell$	$\boldsymbol{x}_k^{\text{target}} \in \mathbb{R}^4$	$\left[\boldsymbol{x}_k^{\text{chaser}},\ \boldsymbol{\tau}_k^{\text{chaser}}\right]^\top \in \mathbb{R}^6$	$\left[f_k,\ i_k\right]^\top \in \mathbb{R}^2$
Digital State	$\boldsymbol{d}_k^\ell \in \mathcal{D}^\ell$	$\mathcal{N}\left(\hat{\boldsymbol{x}}_k^{\text{target}}, P_k^{\text{target}}\right)$	$\mathcal{N}\left(\left[\hat{\boldsymbol{x}}_k^{\text{chaser}},\ \hat{\boldsymbol{\tau}}_k^{\text{chaser}}\right]^\top, P_k^{\text{chaser}}\right)$	$\left[\hat{f}_k,\ \hat{i}_k\right]^\top \in \mathbb{R}^2$
QoIs	$\boldsymbol{q}_k^\ell \in \mathcal{Q}^\ell$	$\hat{n}_k^{\text{target}} \in \mathbb{R}$	$\boldsymbol{v}_k^{\text{chaser}} \in \mathbb{R}^2$	-
Observations	$\boldsymbol{o}_k^\ell \in \mathcal{O}^\ell$	$\left[\rho_k^{\text{target}}, \theta_k^{\text{target}}\right]^\top \in \mathbb{R}^2$	$\left[\rho_k^{\text{chaser}}, \theta_k^{\text{chaser}}\right]^\top \in \mathbb{R}^2$	$\left[\hat{\boldsymbol{\tau}}_k^{\text{chaser}},\ \boldsymbol{v}_k^{\text{chaser}}\right]^\top \in \mathbb{R}^4$
Control Inputs	$\boldsymbol{u}_k^\ell \in \mathcal{U}^\ell$	-	$\boldsymbol{u}_k^{\text{chaser}} \in \mathbb{R}^2$	$\boldsymbol{u}_k^{\text{propulsion}} \in \mathbb{R}^2$
Reward	$\boldsymbol{r}_k^\ell \in \mathcal{R}^\ell$	-	$-\lVert\hat{\boldsymbol{x}}_k^{\text{chaser}} - \hat{\boldsymbol{x}}_k^{\text{target}}\rVert \in \mathbb{R}$	-

is needed to estimate the current state of the failure, and the current estimate of the acceleration bias $\hat{\boldsymbol{\tau}}_k^{\text{chaser}}$ is needed to identify the thruster that is most likely to be faulty. These health estimates are then used to inform how the controller acceleration $\boldsymbol{u}_k^{\text{chaser}}$ is accommodated into a health-aware control input. The computations required for this chain of events are summarized in Algorithm 1.

Algorithm 1 Mission-level integration of the subsystem-level digital twins.

Input: Chaser spacecraft's request for a control input at time $t = t_k$

1: Update target's kinematics digital twin to the requested time: $\boldsymbol{d}_k^{\text{target}}$ as described in Section III.B.1
2: Update chaser's kinematics digital twin to the requested time: $\boldsymbol{d}_k^{\text{chaser}}$ as described in Section III.B.1
3: Update thrusters' health digital twin to the requested time: $\boldsymbol{d}_k^{\text{propulsion}}$ as described in Section III.B.2
4: Evaluate chaser's controller: $\boldsymbol{u}_k^{\text{chaser}}$ as described in Appendix A.2, Eq. 29
5: Compute the health-aware commanded control input: $\boldsymbol{u}_k^{\text{propulsion}}$ as described in Section III.B.3, Eq. 20

D. Simulation Results

The simulated on-orbit rendezvous scenario considers a non-controllable target spacecraft placed on a circular orbit with an altitude of 10^3 km above the Earth's surface, corresponding to an orbital rate of 10^{-3} rad/s and a period of 6.3×10^3 s. The controllable chaser spacecraft is initially located on a nearby elliptical orbit whose semi-major axis is aligned with the horizontal axis of the Earth-centered, Earth-fixed frame. We assume that immediately before the chaser initiates the rendezvous operation, the chaser's and target's positions are connected by a line segment starting at the Earth's center and ending at the chaser's orbit perigee, i.e., the nearest point in the orbit to Earth. For simplicity, we set the prediction, observation, and control frequencies to 0.1 Hz for all the interacting digital twins, but note that the implementation can support any combination of frequencies. In this simulated scenario, the chaser spacecraft's propulsion subsystem experiences a failure at $t_{\text{failure}} = 3 \times 10^3$ s, leading to a constant additive acceleration bias of $\boldsymbol{\tau}_{\text{failure}} = [0.2, 0.05]$ km/s^2 associated with a major failure in the thruster pointing along the +X direction. In our simulated pre-capture phase, mission-success is defined as the chaser achieving a relative distance to the target of 10 km or less. This chaser's position could serve as the starting point for the capture phase depicted in Fig. 1b. The following subsections present the simulation results corresponding to the update of each subsystem-level digital twin and their integration to enable the health-aware rendezvous.

1. Localization and Tracking of the Target

The evolution of the components of the target's kinematics digital state corresponding to the position of the spacecraft is shown in Fig. 4a. Observe that the repeated updates of the kinematics digital twin successfully localize and track the target spacecraft along its ground truth orbital trajectory. Uncertainty is quantified in the posterior belief of the digital state. This posterior belief balances the uncertainty in the localization measurements with the uncertainty in the

modeled dynamics based on a standard EKF setup. This balancing strategy explains why the MAP position estimates may deviate from the measured position, as seen in Fig. 4a. The evolution of the target's orbital rate, a quantity of interest derived from the target's kinematics digital state, is shown in Fig. 4b. The time window shown corresponds to the the portion of the target's trajectory shown in Fig. 4a. Since the computation of the chaser's controller gain matrix requires a point-estimate of the orbital rate, we calculate the orbital rate using the MAP estimate of the target's digital state. Despite the variations of this quantity of interest around its ground truth value as a result of the measurement and process noise, the chosen controls strategy is robust enough to these perturbations and the stability of the chaser spacecraft dynamics is not affected, as described below.

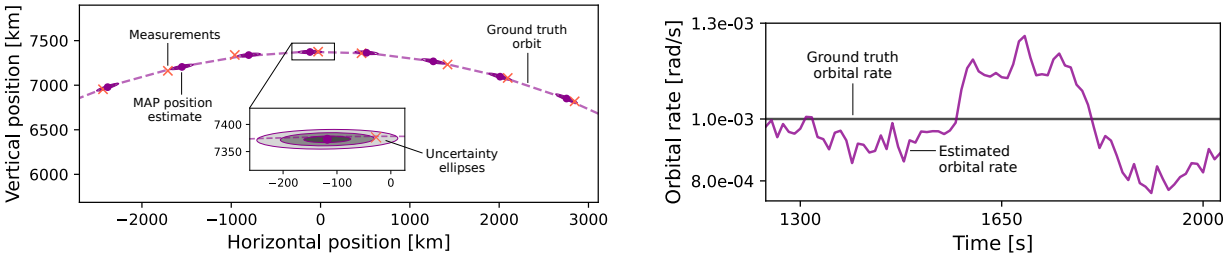

(a) Localization and tracking of the target spacecraft.

(b) Orbital rate of the target spacecraft.

Fig. 4 Updating of the target's kinematics digital twin. Uncertainty in the spacecraft's position estimates is represented using the 3σ uncertainty ellipses shown in the inset.

2. Controlling the Chaser's Kinematics

Figure 5a shows the trajectories of both the target and chaser spacecraft starting from their initial position estimates, but before the failure occurs. The solid lines connect the MAP estimates of the digital states for each spacecraft, while uncertainty ellipses are shown every 5 timesteps. The velocity components of the digital state of each spacecraft are shown in Fig. 5b. The time window corresponds to the the portion of the trajectories shown in Fig. 5a. Note that these velocity estimates are enabled by the predictive dynamic models underlying the kinematics digital twins since the range and bearing angle measurements only inform the position of each spacecraft. We observe that the chaser's position and velocity converge towards the position and velocity of the target. This is achieved by the chaser's kinematics digital twin requesting the target's digital state on-demand, and using the relative state to compute a control input that maximizes the chaser's kinematics reward $\boldsymbol{r}_k^{\text{chaser}}$. An important condition for the stability of the chosen state estimation strategy is the choice of initial guess for the digital state estimates, since the EKF relies on a linearized version of the orbital dynamics of the spacecraft. To initialize the filter, we use the first $\boldsymbol{o}_1^\ell, \boldsymbol{o}_2^\ell, \boldsymbol{o}_3^\ell$ observations of each spacecraft $\ell \in \{\text{target}, \text{chaser}\}$ to obtain an initial guess for the position and velocity at $t = t_2$, leveraging a 3-point central difference stencil for computing an approximation of the velocity and assigning an initial diagonal covariance matrix. Additionally, for the chaser's digital state, the acceleration bias components were initialized to zero. Note that in this implementation both kinematics digital twins feature modular and independent underlying predictive dynamic models, measurements, state estimation, and controls. We emphasize that the specific choices made for those components in this specific simulation setup can be replaced by higher- or lower-fidelity counterparts depending on the mission-level modeling fidelity requirements, thus highlighting the modularity of the subsystem-level digital twins formulation.

3. Health Monitoring of the Chaser's Propulsion Subsystem

When a failure occurs in the propulsion subsystem, the chaser spacecraft will diverge from the desired rendezvous trajectory until the acceleration bias is estimated and accounted for by the health-aware control system. The main diagnostic tool that informs about the presence of the failure is the innovation in the chaser's kinematics (Eq. 14) which is visualized in Fig. 6a. Note how before the failure starts, both components of the innovation are dominated by measurement noise. When the failure occurs, the innovation component corresponding to the range measurement exhibits a deviation that is ultimately indicative of the failure, while the variation of the component corresponding to the bearing angle does not change significantly. Processing the sequence of innovations through the propulsion subsystem health monitoring strategy described in Section III.B.2 leads to detection of the failure within a few timesteps of it

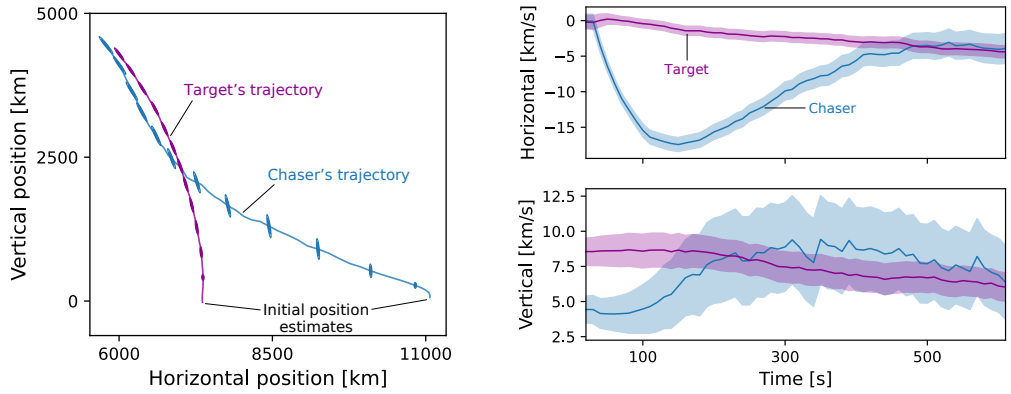

(a) Estimated trajectories.

(b) Estimated component-wise velocities.

Fig. 5 Chaser performing a rendezvous with the target before the failure occurs. Uncertainty in the spacecraft's position estimates is represented using the 3σ uncertainty ellipses, whereas uncertainty in the velocity estimates is represented using the 3σ uncertainty bands.

occurring. Simultaneously, the chaser's kinematics digital twin estimates the acceleration bias, as shown in Fig. 6b. Once the failure is detected, a recovery strategy is triggered as described in Section III.B.3, correcting the chaser's control input with the current estimate of the acceleration bias. With the recovery strategy in place, the innovation component corresponding to the range measurement decreases, indicating how the error in the predicted digital state of the chaser's kinematics is addressed by using the health-aware control inputs. A few timesteps later, the estimated acceleration bias converges to the ground truth value and both components of the innovation are back to the baseline levels observed before the failure occurs.

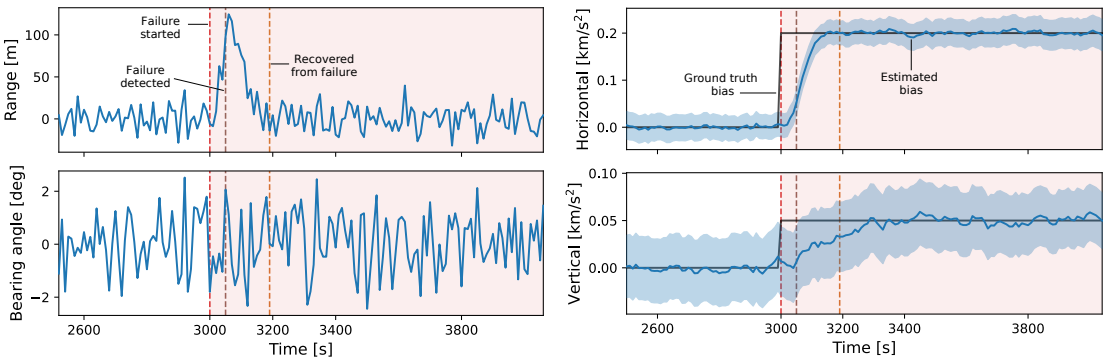

(a) Component-wise innovation in the chaser's kinematics.

(b) Estimated component-wise acceleration bias.

Fig. 6 Evidence of the failure in the chaser's kinematics. Uncertainty in the acceleration bias components estimate is represented with the 3σ uncertainty bands. The shaded region in red denotes the time window when the spacecraft is operating in the presence of a failure. The recovery from the failure during operation is enabled by health-aware control inputs.

Analyzing the failure and recovery from the perspective of the propulsion subsystem health monitoring digital twin provides additional insight about how the failure is detected and isolated. Consider the continual update of this digital twin, which involves estimating the failure status (Fig. 7a), and identifying the thruster that is most likely of being faulty (Fig. 7b) via the failure detection and isolation strategy described in Section III.B.2. Note the evolution of the failure status among the three possible states. Up to the actual failure start time, no failure has been identified. For this particular failure condition, the health monitoring algorithm took 50 s to detect the failure, after which the recovery strategy took 140 s to bring the chaser's kinematics innovation back to baseline levels, hence entering a recovered

from failure mode of operation. Since no intervention was done on the thrusters themselves, in this recovered mode of operation the failure is still present in the propulsion subsystem, but the spacecraft is able to compensate for this failure with health-aware control inputs to successfully continue performing the intended rendezvous operation. Figure 7b shows that by the time the failure is detected the thruster pointing along the +X direction has already been identified as the most likely of being faulty. Additionally, note that the ground truth acceleration bias has a small positive component along the vertical axis, which is identified by the failure isolation strategy by assigning some likelihood of being faulty to the thruster pointing along the +Y direction.

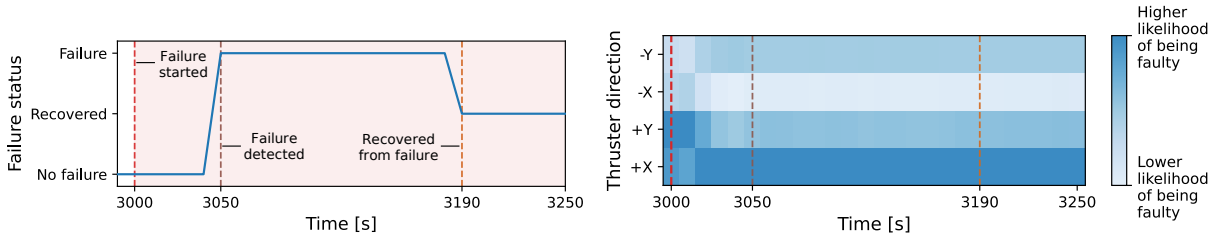

(a) Estimated failure status.

(b) Estimated thruster's likelihood of being faulty.

Fig. 7 Updating the digital twin of the propulsion subsystem health monitoring. The shaded region in red denotes the time window when the spacecraft is operating in the presence of a failure.

4. Digital-twin-enabled Health-aware Rendezvous

The overall rendezvous simulation is depicted in Fig. 8a, showing the evolution of both spacecraft from their initial positions to successful rendezvous within a 10 km radius after experiencing and recovering from a failure. Initially, the chaser maneuvers to the target's orbit successfully. After experiencing the propulsion subsystem failure, the chaser spacecraft starts drifting away from the target. If the failure is not addressed, the chaser would continue on this diverging trajectory. Instead, the failure recovery strategy provides the chaser with a sequence of health-aware control inputs that compensate for the failure and once again achieve a successful rendezvous with the target. Figure 8b shows the time history of both the desired accelerations $\boldsymbol{u}_k^{\text{chaser}}$ and the health-aware accelerations $\boldsymbol{u}_k^{\text{propulsion}}$, along with the actual acceleration the chaser spacecraft is subject to. Before the failure, the chaser's actual acceleration matches the desired acceleration commanded by the controller. At the time of failure, the chaser's actual acceleration shifts suddenly due to the acceleration bias resulting from the failure. At this point, the controller is not yet aware of the failure and continues commanding the desired accelerations as usual. Once the failure is detected, health-aware control inputs are commanded which compensate for the acceleration bias using its most recent estimate to reduce the error between the chaser's actual acceleration and the desired acceleration. After both components of the acceleration bias have converged to their ground truth value, the actual acceleration the chaser spacecraft is subject to is close enough to the desired acceleration commanded by the controller, thus enabling the chaser to successfully complete the pre-capture phase of the on-orbit servicing mission.

IV. Conclusions and Future Work

This work has proposed a modular digital twin formulation of a multi-spacecraft on-orbit servicing scenario where each spacecraft is represented as the aggregate of several subsystem-level digital twins that interact to achieve mission success. This modular formulation defines a digital twin for each subsystem using predictive models, state estimation algorithms, and control strategies that respond to the modeling fidelity requirements of the mission. In contrast with a monolithic digital twin defined by fully-coupled system-wide dynamics, which may be expensive to control and continually update, the modular approach proposes a computationally tractable and scalable alternative for improved monitoring, planning, and control.

In this work we considered achieving coupling between subsystems via on-demand information sharing, but this may not always be feasible, either due to communication limitations or two-way dependencies between subsystem dynamics. Further work is needed to explore how to rigorously manage the coupling between subsystem-level digital twins, especially in regards to controlling the aggregate system-of-systems towards mission success.

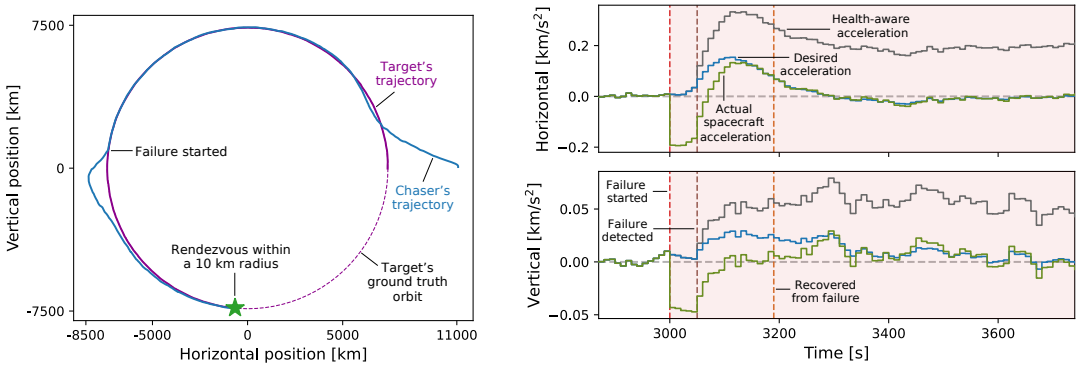

(a) Estimated trajectories.

(b) Chaser's acceleration.

Fig. 8 Overall digital-twin enabled rendezvous simulation results. The shaded region in red denotes the time window when the spacecraft is operating in failure.

The modular digital twin formulation is demonstrated in the context of the pre-capture phase of an on-orbit servicing mission where a controllable and health-aware chaser performs a rendezvous with a noncooperative target. Independent digital twins of the target's and chaser's kinematics are continually calibrated to inform control decisions that enable the rendezvous operation in a failure-free scenario. A digital twin representing the health of the chaser's propulsion subsystem triggers a recovery strategy once a failure is detected and isolated, resulting in a sequence of health-aware control inputs that enable the chaser to continue with the rendezvous operation despite the presence of the failure. Future work on this application may consider extending this example to feature rigid body dynamic models of the interacting spacecraft. Estimates with quantified uncertainty of the spacecraft attitude, inertia, and docking ports may be required to inform further control decisions in later stages of an on-orbit servicing scenario considering the capture of the target and its posterior towing to disposal orbit. Additionally, the presented modular digital twin formulation can also be leveraged to inform estimates of the performance and condition of additional subsystems and their intricate modes of failure. Further work is required to integrate observations, predictive models, and control strategies of such subsystems into continually calibrated digital twins that are coupled with other mission-critical subsystems.

Appendix A: Computational Models

A.1. Orbital Motion

In the setting of the two-body problem, the continuous-time equations of motion that describe the in-plane orbit of a spacecraft around the Earth, as defined in [17] and expressed in the Earth-centered, Earth-fixed frame, are

$$\begin{aligned} \ddot{x} &= -\frac{\mu}{r^3}x + u_x - \tau_x + v_x, \\ \ddot{y} &= -\frac{\mu}{r^3}y + u_y - \tau_y + v_y, \end{aligned} \tag{21}$$

where the position of the spacecraft is denoted by $\boldsymbol{r} = [x,\ y]^\top \in \mathbb{R}^2$, whose magnitude is $r = \sqrt{x^2 + y^2}$, and μ denotes the Earth's gravitational constant. Control inputs are denoted by the vector $\boldsymbol{u} = [u_x,\ u_y]^\top \in \mathbb{R}^2$, bias terms in the acceleration resulting from potential failures in the propulsion subsystem are represented by $\boldsymbol{\tau} = [\tau_x,\ \tau_y]^\top \in \mathbb{R}^2$. Uncertainty in the acceleration is represented by a realization of the continuous-time process noise vector $\boldsymbol{v} = [v_x,\ v_y]^\top \in \mathbb{R}^2$. We consider two modeling scenarios spanning an uncontrollable and a controllable spacecraft, referred to as the target and chaser spacecraft, respectively.

A.1.1 Target Spacecraft Kinematics

For the target spacecraft, we set both the control input and the acceleration bias vectors to zero in the equations of motion (Eq. 21) and rewrite the system of equations in a first-order continuous-time state space formulation with

kinematic state vector $\boldsymbol{x} = [x,\ y,\ \dot{x},\ \dot{y}] \in \mathbb{R}^4$ as

$$\dot{\boldsymbol{x}}(t) = \tilde{\boldsymbol{f}}(\boldsymbol{x}(t)) + D\boldsymbol{v}(t), \tag{22}$$

where $\tilde{\boldsymbol{f}}(\boldsymbol{x}(t)) =: \left[\dot{x},\ \dot{y},\ -\frac{\mu}{r^3}x,\ -\frac{\mu}{r^3}y\right]^\top$ represents the nonlinear continuous-time dynamics function, with μ and r as defined in Eq. 21. The process noise gain matrix is denoted by $D = [0_{2\times2},\ I_{2\times2}]^\top \in \mathbb{R}^{4\times2}$, with 0 and I representing the zero and identity matrices, respectively. Consider the following time discretization $[t_1, t_2, ..., t_k, ..., t_{\text{final}}]$, with $\Delta t_k = t_{k+1} - t_k$, and let $\boldsymbol{x}_k = \boldsymbol{x}(t_k)$. Using a zero-order hold across each Δt_k for the process noise, i.e., $\boldsymbol{v}(t) = \boldsymbol{v}_k\ \forall\ t \in [t_k, t_{k+1}]$, we arrive to the discretized dynamics

$$\boldsymbol{x}_{k+1} = \boldsymbol{x}_k + \int_{t_k}^{t_{k+1}} \tilde{\boldsymbol{f}}(\boldsymbol{x}(t))\, dt + \Delta t_k D\boldsymbol{v}_k =: \boldsymbol{f}(\Delta t_k, \boldsymbol{x}_k, \boldsymbol{v}_k), \tag{23}$$

where $\boldsymbol{f}(\Delta t_k, \boldsymbol{x}_k, \boldsymbol{v}_k)$ represents the discrete-time dynamics function, and the term involving the integral of the continuous-time dynamics function $\tilde{\boldsymbol{f}}(\boldsymbol{x}(t))$ is computed via a suitable numerical integration scheme. Uncertainty in the discrete-time dynamics is represented by the additive process noise vector $\boldsymbol{v}_k \in \mathbb{R}^2$, which we model as a realization of the random variable $\boldsymbol{V}_k \sim \mathcal{N}(\boldsymbol{0}, Q_k)$, with discrete-time covariance matrix $Q_k \in \mathbb{R}^{2\times2}$.

A.1.2 Chaser Spacecraft Kinematics

The chaser spacecraft admits a control input in the form of an acceleration $\boldsymbol{u}(t)$ and may be subject to a additive acceleration bias $\boldsymbol{\tau}(t)$ that takes nonzero values only in the presence of a failure in the propulsion subsystem. This acceleration bias term is a consequence of one or more thrusters being faulty, which can be identified based on the pointing direction of the resulting bias given that the pointing directions of all thrusters are fixed. In a simulation setting, we prescribe the failure by setting the starting time of the failure t_{failure} and the resulting acceleration bias $\boldsymbol{\tau}_{\text{failure}}$. However, the ground truth values of these quantities are not shared with any of the involved digital twins and their values are estimated instead to compute health-aware control inputs.

To estimate the acceleration bias, we consider an augmented version of the dynamics that includes the acceleration bias as part of the state to propagate. Consider the augmented state vector $[\boldsymbol{x},\ \boldsymbol{\tau}]^\top = \left[x,\ y,\ \dot{x},\ \dot{y},\ \tau_x,\ \tau_y\right]^\top \in \mathbb{R}^6$, and augment the equations of motion presented in Eq. 21 with a model for the evolution of the acceleration bias as $\dot{\boldsymbol{\tau}} = \boldsymbol{0}$. In this setting, the estimate of the acceleration bias would remain constant during the prediction step, but would be updated in the correction state of the state estimation algorithm by exploiting the correlation between the acceleration bias with the acceleration of the spacecraft as outlined in the Eq. 21. This augmented dynamics system can be represented in the first-order continuous-time state space formulation as

$$\begin{bmatrix}\dot{\boldsymbol{x}}(t)\\ \dot{\boldsymbol{\tau}}(t)\end{bmatrix} = \begin{bmatrix}\tilde{\boldsymbol{f}}(\boldsymbol{x}(t)) + B\left[\boldsymbol{u}(t) - \boldsymbol{\tau}(t)\right]\\ \boldsymbol{0}\end{bmatrix} + D\boldsymbol{v}(t), \tag{24}$$

where $\tilde{\boldsymbol{f}}(\boldsymbol{x}(t)) \coloneqq \left[\dot{x},\ \dot{y},\ -\frac{\mu}{r^3}x,\ -\frac{\mu}{r^3}y\right]^\top \in \mathbb{R}^4$ represents the nonlinear continuous-time dynamics function, with μ and r as defined in Eq. 21. The control gain matrix is denoted by $B = [0_{2\times2},\ I_{2\times2}]^\top \in \mathbb{R}^{4\times2}$, with 0 and I representing the zero and identity matrices, respectively. The process noise gain matrix is represented by $D = [0_{2\times4},\ I_{4\times4}]^\top \in \mathbb{R}^{6\times4}$ with process noise vector $\boldsymbol{v}(t) \in \mathbb{R}^4$, which was augmented to include model uncertainty corresponding to the dynamics of the acceleration bias. Discretizing time as $[t_1, t_2, ..., t_k, ..., t_{\text{final}}]$ and using a zero-order hold across each Δt_k for the control, acceleration bias, and process noise, i.e., $\boldsymbol{u}(t) = \boldsymbol{u}_k$, $\boldsymbol{\tau}(t) = \boldsymbol{\tau}_k$, and $\boldsymbol{v}(t) = \boldsymbol{v}_k\ \forall\ t \in [t_k, t_{k+1}]$, we arrive to the discretized dynamics

$$\begin{bmatrix}\boldsymbol{x}_{k+1}\\ \boldsymbol{\tau}_{k+1}\end{bmatrix} = \begin{bmatrix}\boldsymbol{x}_k + \int_{t_k}^{t_{k+1}} \tilde{\boldsymbol{f}}(\boldsymbol{x}(t))\, dt + \Delta t_k B\left[\boldsymbol{u}_k - \boldsymbol{\tau}_k\right]\\ \boldsymbol{\tau}_k\end{bmatrix} + \Delta t_k D\boldsymbol{v}_k =: \boldsymbol{f}(\Delta t_k, [\boldsymbol{x}_k, \boldsymbol{\tau}_k]^\top, \boldsymbol{u}_k, \boldsymbol{v}_k), \tag{25}$$

where $\boldsymbol{f}(\Delta t_k, [\boldsymbol{x}_k, \boldsymbol{\tau}_k]^\top, \boldsymbol{u}_k, \boldsymbol{v}_k)$ represents the discrete-time dynamics function. Uncertainty in the discrete-time dynamics is represented via the additive process noise vector $\boldsymbol{v}_k \in \mathbb{R}^4$, which we model as a realization of the random variable $\boldsymbol{V}_k \sim \mathcal{N}(\boldsymbol{0}, Q_k)$ with discrete-time covariance matrix $Q_k \in \mathbb{R}^{4\times4}$.

A.2. Controlling the Relative Motion Kinematics

In the pre-capture phase, the relative motion between the controllable chaser and the target spacecraft needs to be controlled such that their relative position and velocity is minimized under the smallest possible control input magnitude. Consider a rotating frame attached to the target spacecraft's center of mass, whose x-direction pointing radially outward away from the Earth and y-direction points parallel to the target's orbital velocity direction. Assume the target moves in a circular orbit to maintain the orthogonality of this frame, while the chaser is initially at an elliptic orbit. The continuous-time relative motion of the chaser with respect to the target in this rotating frame can be approximated by the Clohessy–Wiltshire–Hill equations [18] resulting from linearizing the coupled relative rotational and translational kinematics between the target and chaser spacecraft as

$$\begin{aligned} \ddot{x} - 2n\dot{y} - 3n^2 x &= u_x, \\ \ddot{y} + 2n\dot{x} &= u_y, \end{aligned} \tag{26}$$

where $[x, y]^\top \in \mathbb{R}^2$ and $[\dot{x}, \dot{y}]^\top \in \mathbb{R}^2$ represent the position and velocity of the chaser spacecraft under control input $[u_x, u_y]^\top \in \mathbb{R}^2$, all written in the rotating target-centered frame. The parameter $n \in \mathbb{R}$ denotes the ground truth target's orbital rate. In our simulation setting, n is not prescribed a priori, and instead it is estimated from the target's position and velocity estimates. These estimates change in time, thus resulting in a discrete sequence of estimates $\hat{n}_1, \hat{n}_2, ...\hat{n}_k, ...$, which are used to recompute the equations of motion in Eq. 26 for each $\hat{n}_k$. Introducing the state vector $\boldsymbol{x} = [x,\ y,\ \dot{x},\ \dot{y}]^\top$ and assuming the following time discretization $[t_1, t_2, ..., t_k, ...]$, the first-order discrete-time relative motion dynamics function corresponding to the equations of motion in Eq. 26 is

$$\boldsymbol{x}_{k+1} = A\boldsymbol{x}_k + B\boldsymbol{u}_k\,, \tag{27}$$

where A denotes the linear discrete-time state-transition matrix and B denotes the control gain matrix. Introducing the proportional control law $\boldsymbol{u}_k = -K\boldsymbol{x}_k$ with controller gain matrix $K \in \mathbb{R}^{2\times 4}$, the control problem can be formulated as an infinite-horizon Linear Quadratic Regulator (LQR) [19]. We seek to minimize the Q-weighted norm of the relative state vector and the R-weighted norm of the control input simultaneously. The optimal control gain matrix K in the rotating target-centered frame is given by

$$K \coloneqq \arg\min_K \sum_{k=1}^{\infty} \boldsymbol{x}_k^\top \left(Q + K^\top R K\right) \boldsymbol{x}_k\,, \tag{28}$$

where Q and R are the positive-definite state and control weight matrices. An analytical solution for the control gain matrix K exists in terms of the discrete-time linear operators A and B and the associated Riccati equation, as described in [19]. Note that both the discrete-time state transition matrix A and the controller gain matrix K depend on the target's orbital rate $\hat{n}_k$, hence must be recomputed whenever a new estimate for the orbital rate becomes available.

Note that the control gain matrix in Eq. 28 is defined with respect to the relative position and velocity state vector $\boldsymbol{x}$ and produces an commanded acceleration vector $\boldsymbol{u}$, with both vectors expressed in the target-centered frame. Instead, we are interested in controlling the chaser nonlinear dynamics presented in Appendix A.1.2, which are expressed in the Earth-centered, Earth-fixed frame. Given estimates of the chaser's $\boldsymbol{x}_k^{\text{chaser}}$ and the target's $\boldsymbol{x}_k^{\text{target}}$ kinematic states expressed in the Earth-centered, Earth-fixed reference frame, we apply appropriate reference frame vector rotations to produce a control input to the nonlinear orbit dynamics of the chaser spacecraft as

$$\boldsymbol{u}_k^{\text{chaser}} = -\Theta_k^\top K \left(I \otimes \Theta_k\right)\left(\boldsymbol{x}_k^{\text{chaser}} - \boldsymbol{x}_k^{\text{target}}\right), \tag{29}$$

where $\Theta_k \in \mathbb{R}^{2\times 2}$ is a rotation matrix between the Earth-centered, Earth-fixed and the rotating target-centered reference frames, and $I \otimes \Theta_k \in \mathbb{R}^{4\times 4}$ denotes the Kronecker product between the identity matrix I of order 2 and the rotation matrix Θ_k.

Acknowledgments

This work was supported by the Air Force Office of Scientific Research (AFOSR) grant FA9550-23-1-0678 under the Space University Research Initiative (SURI). The UT Austin and Emory authors also acknowledge AFOSR Grant FA9550-22-1-0419, Department of Energy grant DE-SC0021239, and NASA ULI Cooperative Agreement 80NSSC21M0071. The UCSC authors also acknowledge National Science Foundation grants CNS-2039054 and CNS-2111688; AFOSR Grants FA9550-23-1-0145, and FA9550-23-1-0313; Air Force Research Laboratory grants FA8651-22-1-0017 and FA8651-23-1-0004; Army Research Office grant W911NF-20-1-0253; and Department of Defense grant W911NF-23-1-0158.

References

[1] AIAA Digital Engineering Integration Committee, "Digital Twin: Definition & Value," An AIAA and AIA institute position paper, American Institute of Aeronautics and Astronautics (AIAA) and Aerospace Industries Association (AIA), 2020.

[2] National Academy of Engineering and National Academies of Sciences, Engineering, and Medicine, *Foundational Research Gaps and Future Directions for Digital Twins*, The National Academies Press, 2024. https://doi.org/10.17226/26894.

[3] Rasheed, A., San, O., and Kvamsdal, T., "Digital twin: Values, challenges and enablers from a modeling perspective," *IEEE Access*, Vol. 8, 2020, pp. 21980–22012. https://doi.org/10.1109/ACCESS.2020.2970143.

[4] Grieves, M., and Vickers, J., "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," *Transdisciplinary Perspectives on Complex Systems: New Findings and Approaches*, 2017, pp. 85–113. https://doi.org/10.1007/978-3-319-38756-7_4.

[5] Kapteyn, M., Pretorius, J., and Willcox, K., "A probabilistic graphical model foundation for enabling predictive digital twins at scale," *Nature Computational Science*, Vol. 1, No. 5, 2021, pp. 337–347. https://doi.org/10.1038/s43588-021-00069-0.

[6] Tezzele, M., Carr, S., Topcu, U., and Willcox, K., "Adaptive planning for risk-aware predictive digital twins," in Physics-based and Data-driven Modeling for Digital Twins, Eds. K. Cherifi and I.V. Gosea, SEMA SIMAI Springer Series, 2024.

[7] Jewison, C., and Erwin, R. S., "A spacecraft benchmark problem for hybrid control and estimation," *2016 IEEE 55th Conference on Decision and Control (CDC)*, 2016, pp. 3300–3305. https://doi.org/10.1109/CDC.2016.7798765.

[8] Malladi, B. P., Sanfelice, R. G., and Butcher, E. A., "Robust hybrid supervisory control for spacecraft close proximity missions," *Annual Reviews in Control*, Vol. 52, 2021, pp. 316–329. https://doi.org/10.1016/j.arcontrol.2021.11.001.

[9] Negenborn, R., and Maestre, J., "Distributed Model Predictive Control: An Overview and Roadmap of Future Research Opportunities," *IEEE Control Systems Magazine*, Vol. 34, No. 4, 2014, pp. 87–97. https://doi.org/10.1109/MCS.2014.2320397.

[10] Liu, G.-P., "Control Strategies for Digital Twin Systems," *IEEE/CAA Journal of Automatica Sinica*, Vol. 11, 2024, p. 170. https://doi.org/10.1109/JAS.2023.123834.

[11] Vered, Y., and Elliott, S. J., "The use of digital twins to remotely update feedback controllers for the motion control of nonlinear dynamic systems," *Mechanical Systems and Signal Processing*, Vol. 185, 2023, p. 109770. https://doi.org/10.1016/j.ymssp.2022.109770.

[12] Rybus, T., Seweryn, K., and Sasiadek, J. Z., "Control System for Free-Floating Space Manipulator Based on Nonlinear Model Predictive Control (NMPC)," *Journal of Intelligent & Robotic Systems*, Vol. 85, 2017, pp. 491–509. https://doi.org/10.1007/s10846-016-0396-2.

[13] Casella, G., and Berger, R., *Statistical Inference*, Chapman & Hall/CRC Texts in Statistical Science, CRC Press, 2024. https://doi.org/10.1201/9781003456285.

[14] Bar-Shalom, Y., Rong Li, X., and Kirubarajan, T., *Estimation with Applications to Tracking and Navigation*, A Wiley-Interscience publication, John Wiley & Sons, Nashville, TN, 2001. https://doi.org/10.1002/0471221279.

[15] Posch, A., Schwientek, A. O., Sommer, J., and Fichter, W., "Model-based on-board realtime thruster fault monitoring," *International Federation of Automatic Control Proceedings Volumes*, Vol. 46, Elsevier BV, 2013, pp. 553–558. https://doi.org/10.3182/20130902-5-de-2040.00080.

[16] Basseville, M., and Nikiforov, I., *Detection of Abrupt Changes: Theory and Application*, Prentice-Hall information and system sciences series, Prentice Hall, Harlow, England, 1993.

[17] Montenbruck, O., and Gill, E., *Satellite Orbits: Models, Methods, and Applications*, Physics and astronomy online library, Springer Berlin, Heidelberg, 2000. https://doi.org/10.1007/978-3-642-58351-3.

[18] Clohessy, W. H., and Wiltshire, R. S., "Terminal Guidance System for Satellite Rendezvous," *Journal of the Aerospace Sciences*, Vol. 27, No. 9, 1960, pp. 653–658. https://doi.org/10.2514/8.8704.

[19] Kirk, D. E., *Optimal Control Theory*, Dover Books on Electrical Engineering, Dover Publications, Mineola, NY, 2004.