Federated Digital Twins for Space Systems

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

A federated digital twin architecture provides a path for managing complexity and computational costs, while also addressing privacy and security constraints. In contrast to a monolithic architecture, which represents a complex system via a single tightly integrated digital twin, a federated architecture comprises multiple subsystem-level digital twins that are coupled through well-defined communication interfaces. These digital twins exchange information with each other while maintaining their own models and data, potentially with built-in privacy and security constraints. This paper proposes a method based on generalized Sobol' sensitivity analysis to characterize the coupling strength between subsystem-level digital twins, which is a first step towards designing a federated digital twin architecture. Leveraging uncertainty quantification, the proposed approach identifies approximate communication strategies that improve computational efficiency while controlling the error in the federated digital twins. We illustrate the potential of this approach for representing a satellite through a federated digital twin architecture suitable for delivering improved operational health monitoring and control capabilities, which in turn lead to improved spacecraft performance, robustness, and reliability.

I. Introduction

Digital twins provide an integrated data-to-decisions framework suitable for delivering operational health monitoring and risk-aware decision-making for complex engineering systems. Underlying a digital twin is a virtual representation of the system and its unique characteristics that is dynamically updated by assimilating data from the physical system. The personalized virtual representation is used to inform optimized control decisions that improve system performance, resulting in a bidirectional flow of information between the digital twin and its physical counterpart [1–3]. Digital twins have demonstrated success across a range of applications, including structural health monitoring [4], autonomous systems [5], and personalized medicine [6–8]. In the context of aerospace vehicles, such as spacecraft, the virtual representation underlying a digital twin are typically comprised of physics-based computational models representing the dynamics of the vehicle and its components. These models are dynamically updated from data via system identification, model calibration, state estimation, or other data assimilation techniques to ensure the digital twin continually reflects the as-is state of the vehicle. Personalization of these computational models enables monitoring the evolution of the vehicle's state, and inform control decisions related to optimal mission schedules or failure recovery strategies [9–11]

To realize the promise of digital twins in the operation of a spacecraft deployed in orbit, its virtual representation needs to include the multidisciplinary subsystems that comprise the spacecraft. Representing the spacecraft via a single and tightly integrated digital twin (Fig. 1a) involves assimilating system-wide observations to dynamically update such coupled models, issuing predictions of the overall system's evolution, and commanding system-wide control inputs. However, as the complexity of the spacecraft increases, maintaining such a monolithic digital twin encompassing all interacting subsystems becomes computationally intractable and impractical. Scaling digital twin technology to represent such complex engineering systems remains an open challenge [3].

Distributed and decentralized formulations have been previously explored when dealing with the design and control of complex engineering systems, aiming to break down the complexity of the otherwise monolithic and centralized formulations. Multidisciplinary design and optimization (MDO) strategies have been developed to make the design of such complex systems computationally feasible [12–16]. From a control perspective, modular approaches combining multiple decentralized controllers have been explored to address the computational challenges associated with having a single controller in charge of optimizing for the system-wide dynamics [17, 18]

A particular type of decentralized architecture is a so-called federated architecture. Generally speaking, a federated architecture consists of a set of components that are integrated within a federation where information exchange and

*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

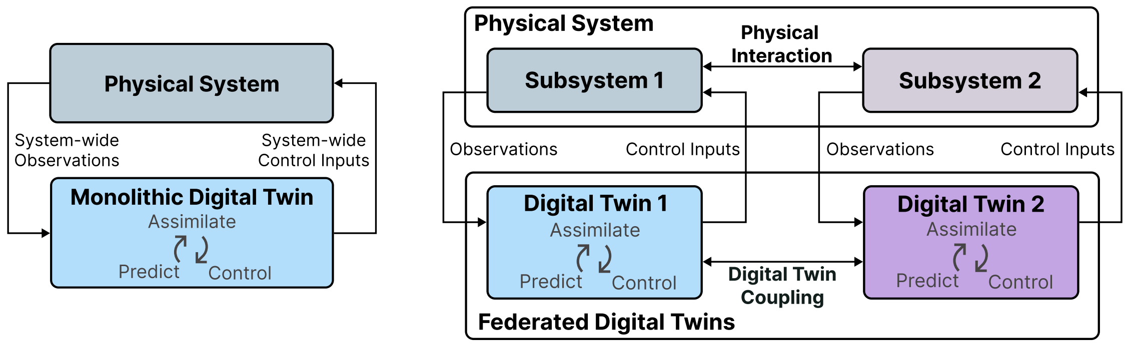

(a) Monolithic digital twin

(b) Federated digital twins

Fig. 1 Digital twin architectures for complex systems. The monolithic architecture features system-wide state, observations, and control inputs, while the federated architecture considers the coupling between digital twins featuring subsystem-level versions of these quantities.

cooperation occur under privacy and security constraints. Federated architectures have been used effectively in other fields, such as federated databases that are logically decentralized and exchange information while maintaining control over their own data [19], and federated machine learning where decentralized components learn a shared model by aggregating local updates and keeping their training data private [20, 21].

In this paper, we propose a federated architecture for digital twins that reflects a divide-and-conquer approach to break down the complexity of the monolithic formulation. In particular, we introduce a partition of the monolithic digital twin into subsystem-level digital twins (Fig. 1b). Each subsystem-level digital twin is responsible for the monitoring and control of the corresponding physical subsystem. Of course, completely decoupled subsystem-level digital twins would fail to account for interactions between subsystems occurring in the physical world. Instead, the computational models underlying each subsystem-level digital twin need to account for not only the subsystem-specific dynamics, but also the interactions with other subsystems. In the federated architecture, this is achieved by each subsystem-level digital twin featuring computational models that encode a representation of the system-wide model parameterized by both internal and external quantities, where the latter are maintained by others in the federation. The subsystem-level digital twins update their knowledge of the external quantities by exchanging information with others in the federation via well-defined communication interfaces while maintaining their internal models and data private and secure.

While the proposed federated architecture provides a path for managing complexity and computational costs, approximations are introduced due to each subsystem-level digital twin having imperfect knowledge of the global system due to infrequent communication with others in the federation. Quantifying and optimally managing this approximation error while subject to computational and communication constraints is critical to a successful federated digital twin architecture. Strategies for collaboration between digital twins have been previously explored [22], but have focused on the system architecture at the implementation level to enable such a collaboration. We instead explore a mathematical characterization of the coupling between digital twins. In particular, quantifying the strength of such coupling becomes a fundamental metric for defining the communication requirements, which is a first step towards designing a federated digital twin architecture. We address this problem in this paper.

Section II describes our proposed approach for designing federated digital twins and the required communication interfaces. We propose a principled approach based on a mathematical characterization of the strength of the coupling between digital twins obtained via generalized Sobol' sensitivity analysis. Section III illustrates the proposed formulation through the design of a federated digital twin architecture building upon a multidisciplinary model of a satellite in operation. Finally, Section IV concludes the paper.

II. Architecting Federated Digital Twins of Complex Engineering Systems

This section proposes a principled approach for designing a communication strategy within the federated digital twin architecture. Section II.A describes the information exchange between digital twins mathematically. The design strategy for the communication within the federated architecture is outlined in Sec. II.B, which relies on a quantification of the coupling strength between digital twins. We propose using the generalized Sobol' sensitivity indices to assess the coupling strength, as described in Sec. II.C.

A. Mathematically Defining the Coupling Between Digital Twins

Coupling between digital twins in the federation is defined as the coordinated information exchange needed to enable cooperation. Mathematically, we represent the exchanged information by defining $\boldsymbol{\theta}_j^i \in \mathbb{R}^{P_j}$ to be the vector containing the P_j quantities that are maintained by the j-th digital twin and that are shared with the i-th digital twin. These quantities may include a combination of subsystem-specific calibrated parameters, estimated states, commanded control inputs, and coupling variables between the i-th and j-th digital twins. In this setting, the computational models of the i-th digital twin depend on information maintained both internally and externally by others in the federation. Define C^i as the subset of indices denoting the digital twins to which the i-th digital twin is directly coupled within a federation of N digital twins. To mathematically represent the dependency of the computational models of each digital twin on the information shared by others in the federation, we consider a scalar-valued quantity of interest (QoI) associated with the i-th digital twin dynamics, denoted as

$$y^i = f^i\left(t; \boldsymbol{\theta}_i^i, \left\{\boldsymbol{\theta}_j^i\right\}_{j \in C^i}\right), \quad t \in [0, T], \tag{1}$$

where f^i is a function of time t, internal quantities $\boldsymbol{\theta}_i^i$, external quantities $\left\{\boldsymbol{\theta}_j^i\right\}_{j \in C^i}$, and T is the length of the prediction horizon. Consider the illustrative federation shown in Fig. 2a composed of two coupled digital twins, namely DT1 and DT2, which share information $\boldsymbol{\theta}_1^2$ and $\boldsymbol{\theta}_2^1$ with each other, as shown in the figure. In this example, $C^1 = \{2\}$, $C^2 = \{1\}$, $y^1 = f^1\left(t; \boldsymbol{\theta}_1^1, \boldsymbol{\theta}_2^1\right)$, and $y^2 = f^2\left(t; \boldsymbol{\theta}_2^2, \boldsymbol{\theta}_1^2\right)$

We proceed by defining a baseline communication frequency $\mathcal{F}$ at which all subsystem-level digital twins exchange information and synchronize. A communication window is defined as the period in between synchronization updates prescribed by the baseline communication frequency, i.e., a time window of length $1/\mathcal{F}$. Due to resource constraints, our desire is to have this all-system communication take place infrequently, resulting in long communication windows during which only some of the digital twins interact depending on their coupling needs. Within one communication window, we consider three coupling scenarios between subsystem-level digital twins, as follows.

I. *Feedback coupled*: During the communication window, DT1 and DT2 continually exchange information (Fig. 2a). In this coupling scenario, the communication is frequent and bidirectional. The evolution of the digital twins must happen synchronously, requiring a coupled solution approach to satisfy the internal models of both digital twins simultaneously. This could be achieved, for example, via a simultaneous joint solution or a fixed-point iterative approach. This is the most computationally intensive scenario.

II. *Feedforward coupled*: During the communication window, the communication is one-directional and DT1 continually broadcasts information to DT2, but DT2 does not broadcast information to DT1 (Fig. 2b). The quantities $\boldsymbol{\theta}_2^1$ associated with DT2 are fixed at a nominal value $\bar{\boldsymbol{\theta}}_2^1$ for use in DT1.

III. *Decoupled*: During the communication window, DT1 and DT2 do not share information. The parameters are set at nominal values $\bar{\boldsymbol{\theta}}_1^2$, $\bar{\boldsymbol{\theta}}_2^1$, and only updated at the end of each communication window (Fig. 2c).

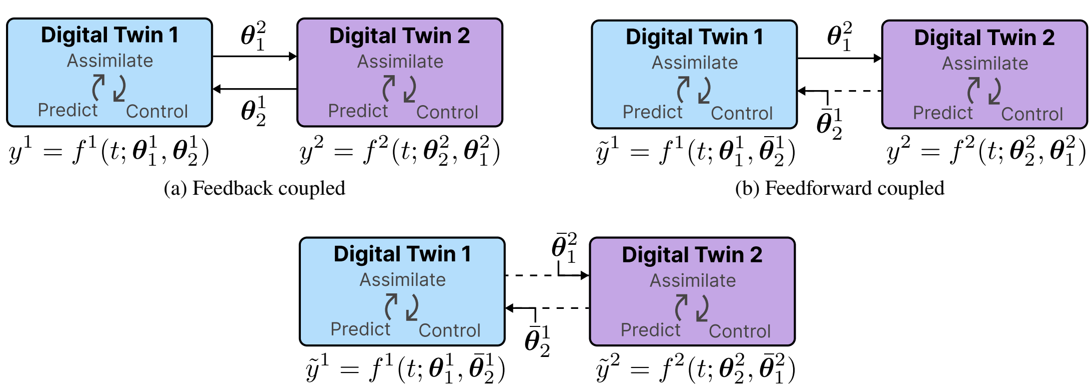

(a) Feedback coupled

(b) Feedforward coupled

(c) Decoupled

Fig. 2 Coupling scenarios between federated digital twins. Dashed arrows represent communication as prescribed by the baseline communication frequency. Solid arrows represent more frequent communication between digital twins. Model inputs fixed to nominal values are denoted as $\bar{\theta}_j^i$.

In practice, each subsystem's QoI is sensitive to only a subset of statistically influential model inputs, which are defined as the quantities whose uncertainty contributes the most to the uncertainty in the QoI predictions. Let $\mathcal{I}^i \subseteq \mathcal{C}^i$ denote the set of indices of the digital twins providing the quantities that are the most statistically influential to the i-th digital twin. The indices of the digital twins providing the otherwise statistically noninfluential model inputs are contained in the set $\mathcal{C}^i \setminus \mathcal{I}^i$. Decoupling is introduced by fixing the noninfluential model inputs to nominal values, simplifying communication within the federated architecture during each communication window. Introducing this decoupling results in an approximation on the internal models of each digital twin, and the QoI becomes

$$\tilde{y}^i = f^i\left(t; \boldsymbol{\theta}_j^i, \{\boldsymbol{\theta}_j^i\}_{j\in\mathcal{I}^i}, \{\bar{\boldsymbol{\theta}}_j^i\}_{j\in\mathcal{C}^i\setminus\mathcal{I}^i}\right), \tag{2}$$

where $\tilde{y}^i$ is an approximation to the QoI y^i. Given the introduced decoupling, rather than requiring continual updates for the statistically noninfluential quantities, the i-th digital twin only needs access to these updates at the end of the communication window prescribed by the baseline communication frequency.

B. Designing the Communication Strategy for Federated Digital Twins

The design of the federated architecture begins with introducing a partition of the system-wide states, observations, and control inputs. The partition can be done by subsystem or discipline comprising the physical system of interest. Components of this partition are associated with each subsystem-level digital twin, which is equipped with computational models of its subsystem-specific dynamics and interactions with other subsystems. Each subsystem-level digital twin is in charge of dynamically updating its computational models and exposing communication interfaces for sharing information with others in the federation. To fully define these communication interfaces, we need to specify what information needs to be shared, how it is accessed by others, and when.

The information to be shared is composed of quantities maintained by a digital twin that others in the federation may need as model inputs, defined mathematically in Sec. II.A as $\boldsymbol{\theta}_j^i$. The way the information is accessed by others is an implementation choice and may follow, for instance, a publisher/subscriber framework. Lastly, to specify how often the digital twins communicate, we define a baseline communication frequency $\mathcal{F}$ that prescribes a default timeline for synchronizing updates. We consider the scenario where the corresponding communication window $1/\mathcal{F}$ is long. The need for exchanging information more frequently within a communication window is defined based on the coupling scenarios presented in Fig. 2. In practice, not all coupling scenarios feature the same strength, and instead come in a spectrum between weak and strong, which in turn raises different communication needs. In the absence of any information about the coupling strength, updates to $\boldsymbol{\theta}_j^i$ would be received by the i-th digital twin infrequently. If the coupling between digital twins is weak, such an infrequent communication may only have a small effect on the predictions issued by this digital twin. However, if the coupling between digital twins is strong, the errors in the predictions will be large, requiring to communicate updates more frequently. We propose to quantify the coupling strength between digital twins and use this metric to define whether some digital twins in the federation need to communicate more often during each communication window.

Let $S_{j,\text{total}}^i(T)$ denote a metric that quantifies the coupling strength from the j-th to the i-th digital twins evaluated at the end of the prediction horizon $T > 1/\mathcal{F}$. The proposed metric is defined mathematically in Sec. II.C. Consider a threshold for the coupling strength denoted by $\varepsilon \geq 0$. We define the spectrum of coupling strength by comparing $S_{j,\text{total}}^i(T)$ with ε. If the coupling strength is equal or below the threshold, i.e., $S_{j,\text{total}}^i(T) \leq \varepsilon$, we introduce a decoupling from the j-th to the i-th digital twins following the definition in Eq. 2 and updates are communicated as prescribed by the baseline communication frequency. Alternatively, if the coupling strength is above the threshold, the coupling falls on a spectrum from weak to strong. When the coupling strength is below a small factor $\kappa > 1$ of the threshold, i.e., $\varepsilon < S_{j,\text{total}}^i(T) \leq \kappa\varepsilon$, we define the i-th digital twin to be *weakly coupled* to the j-th digital twin. When the coupling strength exceeds the small factor of the threshold, i.e., $S_{j,\text{total}}^i(T) > \kappa\varepsilon$, we define the i-th digital twin to be *strongly coupled* to the j-th digital twin. The stronger the coupling, the more frequently the digital twins need to communicate, requiring a boosted communication frequency to share updates originating from the j-th to the i-th digital twin.

C. Quantifying the Coupling Strength Between Digital Twins

We propose to quantify the strength of the coupling between digital twins through the sensitivity of their computational models to lack of knowledge about quantities maintained by other digital twins in the federation that are needed as model inputs. We represent this lack of knowledge as uncertainty in those model inputs that is propagated through each digital twin's computational models into subsystem-specific QoIs. Contributions to the output uncertainty are apportioned to the model inputs following a variance-based sensitivity analysis by computing Sobol' indices generalized in time to account for the system's evolution [23]. These sensitivity indices, denoted $S^{i}_{j,\text{total}}(t)$, rank the model inputs by statistical influence and provide a principled framework to assess the coupling strength between digital twins in the federation. A larger total generalized Sobol' index $S^{i}_{j,\text{total}}(t)$ implies a higher cumulative contribution to the cumulative output uncertainty from the quantities $\boldsymbol{\theta}^{i}_{j}$ maintained by the j-th digital twin up to time t. We define this scenario to represent *strong coupling* from the j-th to the i-th digital twins at time t. Conversely, a smaller total generalized Sobol' index $S^{i}_{j,\text{total}}(t)$ implies a lower cumulative contribution to the cumulative output uncertainty by $\boldsymbol{\theta}^{i}_{j}$, and represents *weak coupling* from the j-th to the i-th digital twins. Since the generalized Sobol' sensitivity indices are a function of time, we propose using their values at the end of the prediction horizon, i.e., $S^{i}_{j,\text{total}}(T)$, as a representative metric of the overall coupling strength between digital twins.

We introduce this problem mathematically following the presentation of the generalized Sobol' sensitivity analysis by Alexanderian et al. [23]. Consider the scalar-valued QoI representing the i-th digital twin and its model inputs consisting of internal $\boldsymbol{\theta}^{i}_{i}$, and external quantities $\{\boldsymbol{\theta}^{i}_{j}\}_{j\in C^{i}}$ as defined in Eq. 1. Nominal values for these quantities typically come from the system's design parameters or expected values during operation and are denoted as

$$\bar{\boldsymbol{\theta}}^{i}_{j} = \left[\bar{\theta}^{i}_{j,1}, \bar{\theta}^{i}_{j,2}, \dots, \bar{\theta}^{i}_{j,P_j}\right]^{\top} \in \mathbb{R}^{P_j}, \quad i, j \in \{1, \dots, N\}. \tag{3}$$

Define random perturbations of these quantities about the nominal values as

$$\boldsymbol{\Theta}^{i}_{j} = \left[\Theta^{i}_{j,1}, \Theta^{i}_{j,2}, \dots, \Theta^{i}_{j,P_j}\right]^{\top}, \quad \Theta^{i}_{j,p} = \bar{\theta}^{i}_{j,p}\left(1 + \alpha^{i}_{j,p}\,\xi\right), \quad \alpha^{i}_{j,p} \geq 0, \quad \xi \sim \mathcal{U}(-1,1), \quad p \in \{1, \dots, P_j\}, \tag{4}$$

which are independent and uniformly distributed within intervals defined by each perturbation factor $\alpha^{i}_{j,p}$.

After propagating the uncertainty from the random variables through each digital twin's models (assumed to be deterministic functions), the output QoI becomes a random variable too

$$Y^{i} = f^{i}\left(t; \boldsymbol{\Theta}^{i}_{i}, \{\boldsymbol{\Theta}^{i}_{j}\}_{j\in C^{i}}\right), \tag{5}$$

The variance of Eq. 5 is computed through a second-order analysis of variance (ANOVA) decomposition [24] as

$$D^{i}(t) = \operatorname{var}[\, Y^{i} \,|\, t \,] = \sum_{j=1}^{N} D^{i}_{j}(t) + \sum_{1\leq j<k\leq N} D^{i}_{jk}(t)\,, \tag{6}$$

with contributions

$$D^{i}_{j}(t) = \operatorname{var}\left[\, \mathbb{E}\left[\, Y^{i} \,\middle|\, t, \boldsymbol{\Theta}^{i}_{j} \,\right] \,\right], \tag{7}$$

$$D^{i}_{jk}(t) = \operatorname{var}\left[\, \mathbb{E}\left[\, Y^{i} \,\middle|\, t, \boldsymbol{\Theta}^{i}_{j}, \boldsymbol{\Theta}^{i}_{k} \,\right] \,\right]. \tag{8}$$

The decomposition of variance in Eq. 6 apportions the first-order contributions that each quantity has to the total variance individually (Eq. 7), and the second-order contributions resulting from the interactions between two quantities (Eq. 8). The total contribution to the variance associated with the quantities shared by the j-th digital twin is defined as

$$D^{i}_{j,\text{total}}(t) = D^{i}_{j}(t) + \sum_{\substack{k=1\\k\neq j}}^{N} D^{i}_{jk}(t)\,. \tag{9}$$

Ratios between the total contributions to the variance and the total variance, i.e., $D^{i}_{j,\text{total}}(t)/D^{i}(t)$, define the total Sobol' sensitivity indices pointwise in time [25]. However, these pointwise Sobol' indices are not aware of the temporal evolution and correlation of the QoI. Additionally, the variance of the QoI itself changes in time (Eq. 6), which raises the challenge of comparing a small contribution to a large variance with a large contribution to a small variance and limits the direct comparison of pointwise indices between two different times. To address these issues, the Sobol' sensitivity indices have been generalized in time [23] by considering the ratios between the cumulative contributions to the variance and the cumulative total variance up to a given time. The total generalized Sobol' sensitivity index is defined as

$$S^{i}_{j,\text{total}}(t) = \frac{\int_0^t D^{i}_{j,\text{total}}(\tau)\,d\tau}{\int_0^t D^{i}(\tau)\,d\tau}. \tag{10}$$

To compute the generalized Sobol' indices above, we need to compute (1) the total variance and contributions as a function of time, and (2) the time integrals. In a general setting, analytical expressions for the forward models defining the QoIs may not be available, instead requiring efficient and accurate numerical approximation techniques to estimate the generalized Sobol' indices. Integration in the time interval $[0, t]$ can be approximated numerically through a quadrature rule with nodes $\{t_m\}_{m=1}^{M}$ and weights $\{\omega_m\}_{m=1}^{M}$. Computing the variances, however, represents a more involved computational task given the high-dimensional integration in probability space. If a large computational budget is available, Monte Carlo integration using fixing methods can be used to approximate the total variance and contributions. These methods use two sets of N_{MC} independent samples that are combined strategically to produce estimates for the total variance $\tilde{D}^{\,i}(t)$, the first-order $\tilde{D}^{\,i}_{j}(t)$, and total $\tilde{D}^{\,i}_{j,\text{total}}(t)$ contributions requiring $N_{\text{MC}}(2 + \sum_{j=1}^{N} P_j)$ forward model evaluations. Multiple estimators are available in the literature, including biased [26] and unbiased [27] versions. The resulting estimate for the total generalized Sobol' sensitivity index becomes

$$S^{\,i}_{j,\text{total}}(t) \approx \frac{\sum_{m=1}^{M} \omega_m \tilde{D}^{\,i}_{j,\text{total}}(t_m)}{\sum_{m=1}^{M} \omega_m \tilde{D}^{\,i}(t_m)} \,. \tag{11}$$

The set $\mathcal{I}^{\,i}$ containing the indices of the digital twins providing the most statistically influential model inputs can be identified by ranking the total generalized Sobol' indices and selecting the ones exceeding a given threshold $\varepsilon \geq 0$. Fixing the otherwise noninfluential model inputs to nominal values introduces an approximation in the QoI predictions. A priori error estimates for this approximation are provided in [23], which are proportional to sum of the total generalized Sobol' sensitivity indices associated with the fixed variables.

III. Application: Federated Digital Twins Representing an Orbiting Spacecraft

We illustrate the capabilities of the proposed methodology through the design of a federated digital twin architecture for representing an academic example of an orbiting satellite via subsystem-level digital twins that are coupled through a federated architecture. Section III.A provides the details of the multidisciplinary model that enables the simulation of the satellite in operation by representing the multiple subsystems comprising the spacecraft. The abstraction as subsystem-level digital twins is presented in Sec. III.B, and the proposed design of the federated digital twin architecture is presented in Sec. III.C. Finally, a simple numerical example is set up in Sec. III.D to illustrate the communication strategy between digital twins featuring weak and strong coupling in the federation.

A. Multidisciplinary Model of a Satellite's Operation

The subsystems comprising a spacecraft are often designed and developed by different teams using heterogeneous technologies. In the field of multidisciplinary design optimization, the design problem of such a modular system is setup as large-scale optimization problem integrating computational models and constraints representative of the interacting disciplines and subsystems [28]. Our chosen application problem uses the model of a satellite's operation developed by Hwang et al. [29] for the multidisciplinary design optimization of a communications satellite.

The mission of this communications satellite is to continuously collect data and transmit as much of that data as possible to the ground stations. The satellite is composed of a cubic body and four fins. Solar panels cover each of the 12 exposed surfaces. The full satellite design problem was considered simultaneously, including seven major disciplines, multiple time scales, and design variables parameterizing the variation of several quantities over time. The goal was to maximize the total data downloaded subject to constraints in the power and energy available over a year of operation of the satellite. Outputs of the multidisciplinary design optimization problem include the design values for the fin and antenna angles. Additionally, control inputs are also scheduled as part of the solution to the optimization problem, including roll angle profile, antenna power profile, and for each solar panel, the output current for maximum power point tracking is commanded. We refer the reader to [29] for more details about the computational models underlying each of the interacting disciplines and the solution of the multidisciplinary optimization problem.

In this work, we consider the multidisciplinary model representing the spacecraft and the obtained optimal design as inputs. Figure 3 shows the forward model that computes the QoIs representative of the operation of the satellite by evaluating the multidisciplinary model. Feedforward coupling between disciplines is enabled by sharing internal state variables sequentially. Among the input parameters are the fin and antenna angles, which are set to the design values of 64.4° and −45°, respectively. We consider the operation of the satellite during a window of $T = 12$ hours set 11 months after launch. Initial conditions include the initial position, velocity, and temperature of the spacecraft, and the initial battery's state of charge (SOC). The scheduled control inputs determined by the optimizer are shown in Fig. 4.

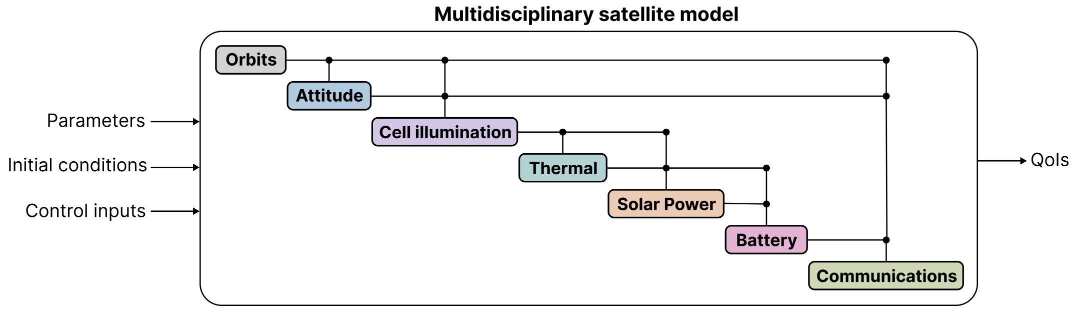

Fig. 3 Forward model to obtain the QoIs from a multidisciplinary simulation of an orbiting satellite. Adapted from [29].

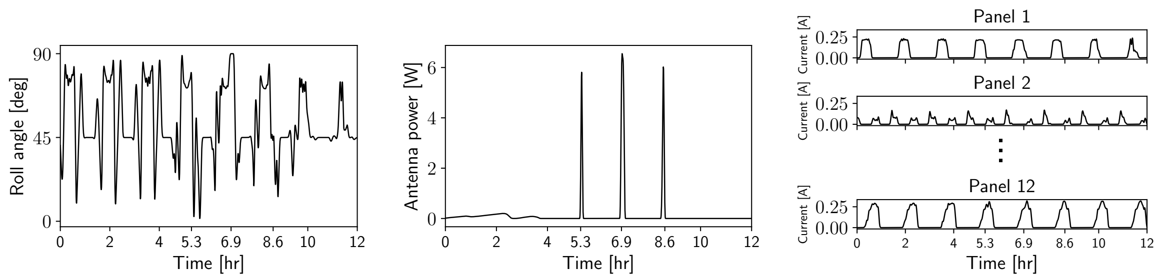

(a) Attitude profile

(b) Antenna power

(c) Solar panel current

Fig. 4 Control inputs to the forward model obtained from the solution to the satellite multidisciplinary design optimization problem [29].

We choose seven QoIs that are representative of the evolution of each subsystem during the operation of the satellite. Figure 5 shows those QoIs evaluated using the nominal set of parameters, initial conditions, and control inputs. For the orbits discipline (Fig. 5a), we consider the magnitude of the position vector, i.e., range, of the satellite's center of mass (CoM) described in the Earth Centered Inertial (ECI) frame of reference. For the attitude discipline (Fig. 5b), we consider the total power required by the set of reaction wheels to maintain the commanded pointing attitude prescribed by the roll angle profile. For the cell illumination discipline (Fig. 5c), we consider the total surface area that is exposed to the sun when the satellite and the sun are within line of sight (LOS). For the thermal discipline (Fig. 5d), we consider the lumped temperature of the satellite's body. For the solar power discipline (Fig. 5e), we consider the total output power of the 12 solar panels, which depends on the panels' I-V curve and the commanded output current for maximum power point tracking for each panel. For the battery (Fig. 5f), we consider the evolution of the battery's state of charge (SOC) as a measure of the total power available. Finally, for the communications discipline (Fig. 5g), we consider the total data downloaded, which depends on both the line of sight between the spacecraft's antenna and the ground station, and the commanded antenna power profile

B. Abstraction as Subsystem-level Digital Twins

Consider a monolithic digital twin representing the system-wide state of the satellite in operation by integrating all disciplines presented in Fig. 3 simultaneously. We seek to introduce a partition to such a monolithic digital twin and instead represent the satellite through a set of coupled subsystem-level digital twins that collaborate through a federated architecture. In this section, we introduce such a partition to the monolithic digital twin based on the disciplines comprising the satellite multidisciplinary model. We proceed to define the subsystem-level digital twins, their roles, and the information they will share within the federation. The orbits digital twin receives spacecraft localization measurements and is in charge of tracking the position and velocity of the satellite's center of mass as the satellite orbits the Earth. The attitude digital twin is in charge of guaranteeing the satellite maintains a forward-facing orientation along its orbit and executing the commanded roll angle profile is satellite. This digital twin estimates the attitude of the spacecraft from measurements coming from the on-board inertial measurement unit (IMU), and commands torques and angular velocities for the reaction wheels. Additionally, the current and voltage of the reaction wheels are measured

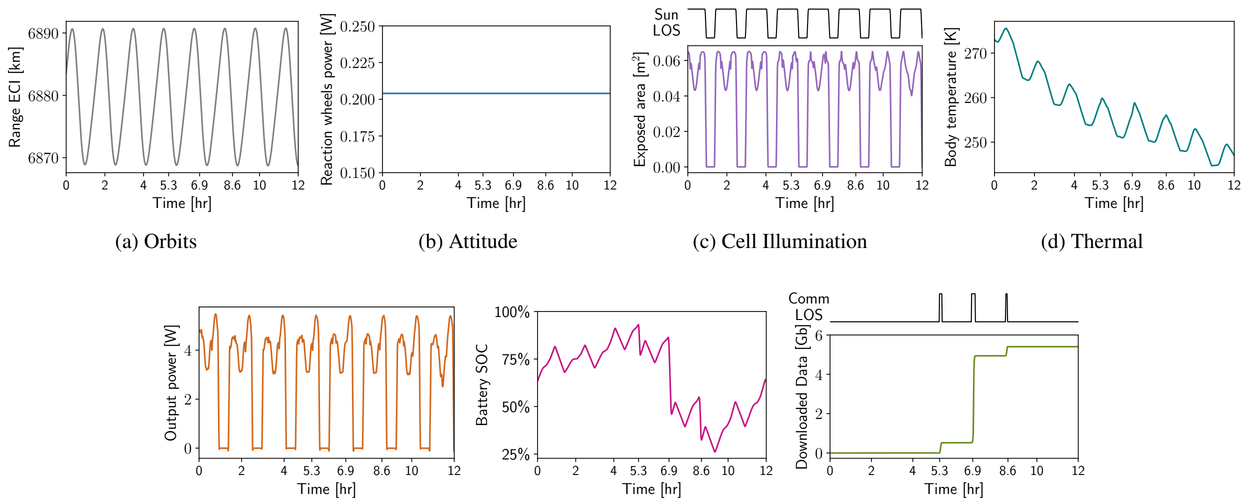

(a) Orbits (b) Attitude (c) Cell Illumination (d) Thermal

(e) Solar Power (f) Battery (g) Communications

Fig. 5 Multidisciplinary simulation of an orbiting satellite during a 12-hour operational window. Line of sight with the sun (Sun LOS) and the ground station (Comm LOS) for data downlink dictate the thermal, power, and communication dynamics, which combined affect the evolution of the battery state of charge (SOC).

to estimate their power consumption. The cell illumination digital twin is in charge of tracking the LOS between the satellite and the sun and estimating how much area is exposed to the sun based on measurements of the sun's azimuth and elevation angles. Additionally, once the spacecraft is deployed in orbit and the fins are extended, this digital twin is in charge of estimating any deviations of the fin angle with respect to the design value. The thermal digital twin assimilates noisy temperature measurements collected at the satellite body and fins and is in charge of predicting the evolution of those temperatures along the spacecraft's orbit. Additionally, the thermal properties of the solar panels on the surfaces of the satellite may degrade over time, and this digital twin is in charge of continually estimating any changes on the absorptivity of the panels. The solar power digital twin is in charge of tracking the power output from the solar panels from measurements of the output current and voltage. The underlying models of this subsystem depend on the diode ideality factor of the solar panels, which needs to be continually calibrated since it conditions the predicted solar power output. By leveraging these dynamically updated models, this digital twin is in charge of commanding the output current for each solar panel that maximizes power point tracking. The battery digital twin is in charge of tracking the evolution of the battery's SOC and cycle count as an indicator of the battery's health. Additionally, the power load from scientific instruments on-board and other instruments needs to be continually estimated to refine the estimates of the available power in the battery. Finally, the communications digital twin tracks the LOS between the antenna and the ground station. By assimilating telemetry data from the antenna, the data download rate and antenna power conversion efficiency are continually calibrated. Additionally, this digital twin commands the power profile for the antenna and triggers the decision to uplink/downlink data.

The computational model underlying each subsystem-level the digital twin is the forward model shown in Fig. 3. This model is evaluated by each subsystem-level digital twin to obtain the QoI from Fig. 5 associated with the corresponding subsystem. Within each digital twin, the model inputs that are used to evaluate the forward model depend on quantities maintained both internally and externally by others in the federation. Therefore, despite featuring the same computational model, two digital twins may produce different outputs for the same QoI if their realizations of the model inputs differ.

Note that the subsystem-level digital twins are so far decoupled. To enable collaboration within the federation, each digital twin needs to expose a communication interface. As a first step towards defining these communication interfaces, Table 1 summarizes the quantities that are shared by each digital twin with others in the federation. For each quantity, we define nominal values from the results reported in [29], which were either specified as model parameters or were obtained through the solution to the multidisciplinary optimization problem. This table also lists perturbation factors chosen to represent the lack of knowledge over these quantities probabilistically as defined in Eq. 4. Perturbation factors were below 20% for all quantities. The perturbation factor for the satellite's CoM position was chosen to be 0.05%, smaller than for the rest of quantities, since small perturbations of this quantity have large effects on the satellite's orbit.

Table 1 Information that each digital twin shares with others in the federation. Each quantity is shared by the digital twin corresponding to the subsystem indicated by the row header to the left.

	Quantity	Description	Nominal Value	Perturbation
Orbits	$\boldsymbol{\theta}_{\text{orb}}$	Satellite's CoM position	$[-0.69, -1.08, -6.76] \times 10^3$ km	0.05%
Attitude	$\boldsymbol{\theta}_{\text{att}}$	Roll angle profile	Curve on Fig. 4a	20%
		Reaction wheel bias current	0.017 A	10%
Cell Illumination	$\boldsymbol{\theta}_{\text{cell}}$	Fin angle estimate	64.4°	10%
Thermal	$\boldsymbol{\theta}_{\text{therm}}$	Body and fins temperatures	273 K	1%
		Solar panels' absorptivity	0.9	5%
Solar Power	$\boldsymbol{\theta}_{\text{solar}}$	Diode ideality factor	1.35	10%
Battery	$\boldsymbol{\theta}_{\text{batt}}$	State of charge	63%	10%
		Instruments power load	2.0 W	20%
Communications	$\boldsymbol{\theta}_{\text{comm}}$	Antenna efficiency	20%	10%

C. Federated Digital Twin Architecture

Next step is to perform a generalized Sobol' sensitivity analysis as described in Sec. II.C to assess the coupling strength between the digital twins introduced in Sec. III.B. Inputs for this analysis are the forward model presented in Fig. 5 and the shared quantities listed in Table 1, including their nominal values and chosen perturbation factors. The results of the generalized Sobol' sensitivity analysis are presented in Fig. 6. For each subfigure, the left panel shows uncertainty in the output QoI obtained from sampling the random perturbations to the nominal model inputs. The center panel shows the evolution of the total generalized Sobol' sensitivity indices as a function of time. Each colored line corresponds to the sensitivity index of the quantity shared by the digital twin with same color in the legend at the top. The right panel summarizes the sensitivity indices evaluated at the end of the prediction horizon, i.e., $S^i_{j,\text{total}}(T)$.

From the generalized Sobol' sensitivity indices shown in Fig. 6, we observe that there are no feedback coupled digital twins, i.e., there is no pair of indices $i \neq j$ such that $S^i_{j,\text{total}}(T) \approx S^j_{i,\text{total}}(T)$. This is a consequence of the fact that the coupling between disciplines in the multidisciplinary model in Fig. 3 is feedforward, hence uncertainty is only propagated downstream in this model. Additionally, the generalized Sobol' sensitivity analysis captures indirect coupling through the multidisciplinary model, revealing which sources of uncertainty are the most statistically influential beyond what is depicted by the direct connections between disciplines shown in Fig. 3. Take, for instance, the results for the solar power digital twin shown in Fig. 6e. Uncertainty in the quantities shared by the orbits and attitude digital twins have large contributions to the uncertainty in the predicted output power, yet the solar power discipline is not directly coupled to either the orbits or attitude disciplines in the multidisciplinary model. Instead, the coupling between these disciplines happens indirectly through the cell illumination discipline, and uncertainty from the quantities shared by the orbits and attitude digital twins is propagated downstream through the multidisciplinary model.

We proceed to rank the sensitivity indices to assess the strength of the coupling between digital twins following the methodology described in Sec. II.B. In this analysis, we set the sensitivity threshold to $\varepsilon = 0.03$, and $\kappa = 3.0$. The resulting ranking of the coupling strength between digital twins, which in turn defines the coupling structure within the federated architecture, is shown in Fig. 7. We use weighted arrows to represent the coupling strength and indicate the shared information near the arrow heads. If two digital twins are decoupled within one communication window, no arrow is used to connect them. Instead, an arrow connecting two digital twins implies that they need to communicate more often during each communication window. Note how the coupling between the digital twins in the federation comes in a spectrum between weak and strong. The solar power digital twin, for instance, is strongly coupled to the orbits, attitude and cell illumination digital twins. Updates originating from these three digital twins need to be continually communicated to the solar power digital twin. The battery digital twin, on the other hand, is strongly coupled to the solar power digital twin, but only weakly coupled to the cell illumination digital twin. In this case, updates originating from the cell illumination digital twin do not need to be communicated to the battery digital twin as often as updates originating from the solar power digital twin. Recall that, as defined in Sec. II.B, all digital twins synchronize updates at the end of every communication window prescribed the baseline communication frequency. However, the coupling structure shown in Fig. 7 informs a boosted communication strategy, highlighting which digital twins have higher priority to receive updates originating from others in the federation during one communication window.

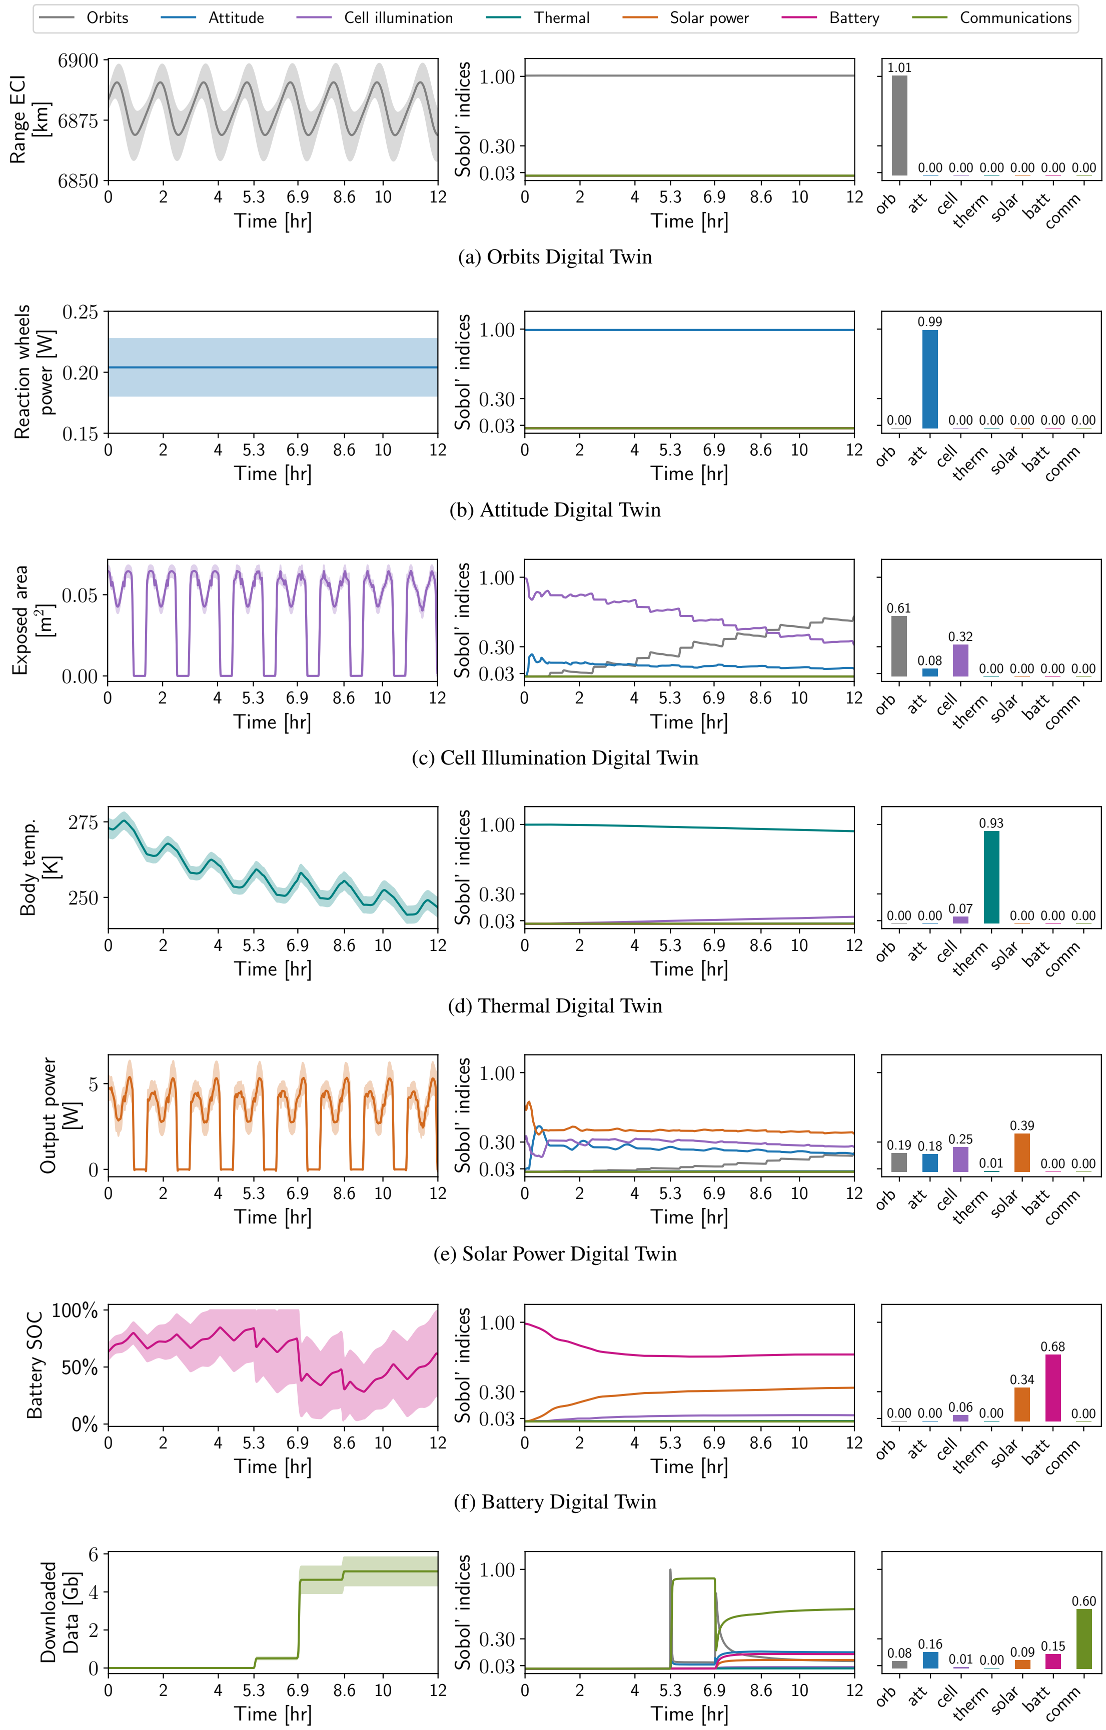

(a) Orbits Digital Twin

(b) Attitude Digital Twin

(c) Cell Illumination Digital Twin

(d) Thermal Digital Twin

(e) Solar Power Digital Twin

(f) Battery Digital Twin

(g) Communications Digital Twin

Fig. 6 Forward uncertainty propagation and generalized Sobol' sensitivity analysis. Leftmost panels show the mean and 2σ uncertainty bands in the subsystem-level QoIs. Center and rightmost panels show the generalized Sobol' sensitivities over time and at the last timestep, respectively.

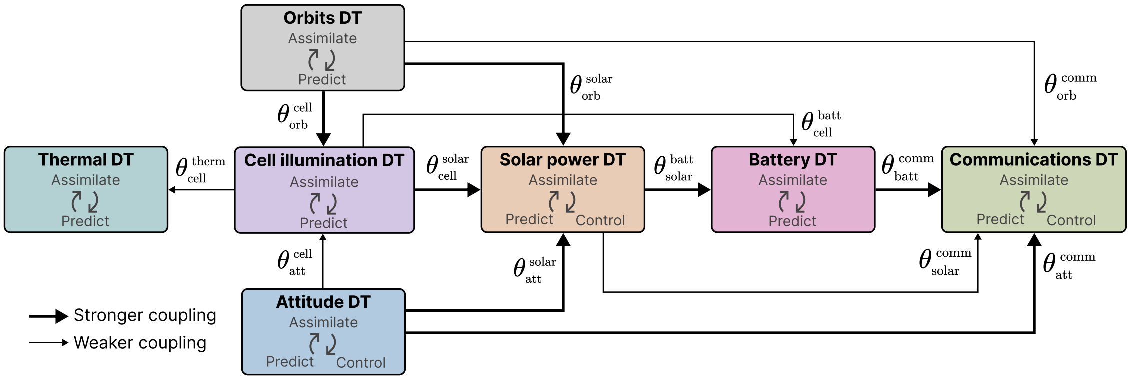

Fig. 7 Federated digital twin (DT) architecture design resulting from the generalized Sobol' sensitivity analysis.

D. Numerical Example: Communication Between Federated Digital Twins

We set up a simple numerical example to illustrate the effects of the choices made about the communication between digital twins and the design of the federated architecture. This example, summarized in Fig. 8, considers the initial calibration step of the solar power digital twin once the federated architecture is deployed. First, measurements of the panels' current and voltage are received by the solar power digital twin. This digital twin proceeds to assimilate the measurements into a new estimate of the diode ideality factor, which we represent probabilistically. The old estimate of this parameter was defined by sampling uniform perturbations about its nominal design value. The next step is to predict the output solar power given the new estimate for the diode ideality factor. Samples for both the old and new estimates are pushed through the forward model and we compare the mean of the predictions. The relative error between the predicted means is used to assess the error of the predictions if the estimate of the diode ideality factor had not been updated. In this scenario, the errors in the output power prediction would have been close to 7% when in LOS with the sun. Next, the updates about the diode ideality factor estimates are shared with the battery and communications digital twins. This is where the effects of receiving the updates infrequently during one communication window becomes clear.

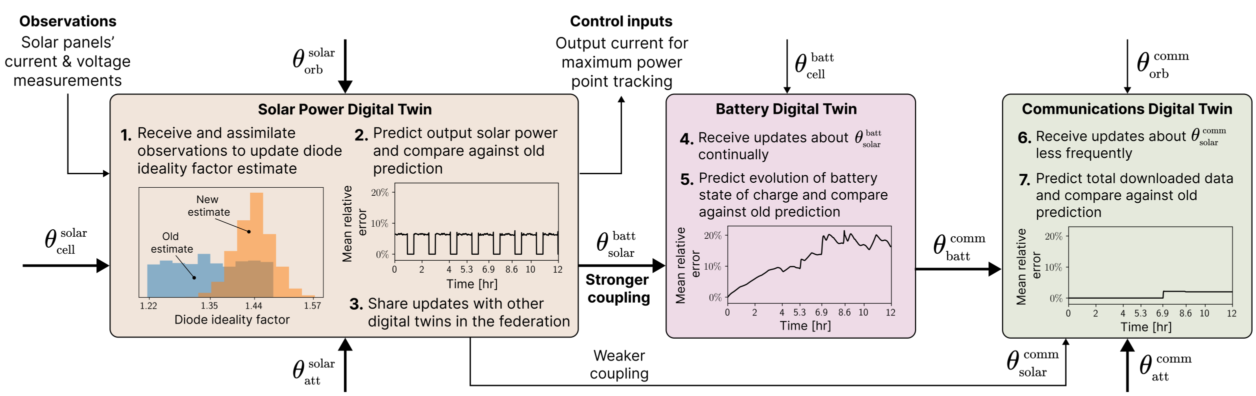

Fig. 8 Simple example illustrating the effects of the choices made about the coupling between digital twins in the federated architecture. The mean relative error shown in the figure accounts for the error in the prediction of each QoI if the estimate of the diode ideality factor had not been updated within each subsystem-level digital twin.

If the diode ideality factor had not been updated to the new estimate within the computational models underlying the battery digital twin, Fig. 8 shows that the mean relative error in the prediction of the evolution of the battery SOC is as high as 20% , confirming that the battery digital twin is strongly coupled to the solar power digital twin, as our analysis summarized in Fig. 7 suggested. If the communication budget is limited or the network is congested, the battery digital twin digital twin should have higher priority to receive information shared by the solar power digital twin. On the other hand, the communications digital twin is weakly coupled to the solar power digital twin. Therefore, the errors in the predictions of the downloaded data due to not having access to the new estimate of the diode ideality factor are not expected to be as high as the ones observed in the prediction of the evolution of the battery SOC. In fact, for the entire prediction horizon, the mean relative error in the prediction of this QoI remains below 3% if the diode ideality factor had not been updated within the communications digital twin.

Note that the results of the generalized Sobol' sensitivity analysis presented in Fig. 6 may change based the choice of QoIs, nominal values for the model inputs, perturbation factors, and other hyperparameters. Since we are extrapolating decisions from this analysis to the design of the federated digital twin architecture, results need to be verified once the architecture is deployed in operation. If newer estimates for the model inputs fall outside of the assumed level of uncertainty in their perturbations, the generalized sensitivity analysis needs to be conducted again and the communication strategy within the federation may change.

IV. Conclusions and Future Work

This paper proposed a federated architecture for building digital twins of complex engineering systems, such as an orbiting spacecraft. In contrast to a monolithic formulation, the proposed federated architecture reflects a divide-and-conquer approach. Modular digital twins are defined at the subsystem-level and coupled through well-defined communication interfaces to represent an entire system. Collaboration between digital twins is enabled by coordinated information exchange, which in turn leads to improved operational health monitoring and control capabilities. The federated architecture breaks down the complexity and computational cost of the monolithic formulation and provides built-in privacy and security constraints for the models and data underlying each digital twin.

We focused on establishing a principled strategy to design the communication strategy within a federated architecture by quantifying the coupling strength between digital twins through a generalized Sobol' sensitivity analysis. This analysis apportions the contributions that model inputs have to the output uncertainty of a scalar-valued QoI. For each subsystem-level digital twin, some of the model inputs to its underlying computational models are quantities maintained externally by others in the federation. We proposed to assess the coupling strength between digital twins by ranking the contributions to the uncertainty from each of these external quantities, defining a spectrum from weak to strong coupling. The computational cost of Monte Carlo-based approaches for conducting such a sensitivity analysis, however, may be prohibitive if the number of digital twins in the federation is large and the computational models of each digital twin depend on a large number of external quantities. Therefore, further work is needed to provide computationally tractable and accurate estimators for the generalized Sobol' sensitivity indices.

We illustrated how the proposed strategy can be applied to the design of a federated digital twin architecture representing a communications satellite in operation. We built upon a multidisciplinary model representing the satellite during a 12-hour operation window and integrating seven feedforward coupled disciplines. We proposed an abstraction as subsystem-level digital twins corresponding to each of the disciplines in the model, establishing the roles of the digital twins, their QoIs, and the information they will share with others in the federation. A generalized Sobol' sensitivity analysis was conducted and a design for the federated digital twin architecture was proposed, providing details about the coupling strength and the flow of information between digital twins. Finally, a simple numerical example illustrated the effects of delayed communication between strongly and weakly coupled digital twins. Future work should focus on expanding this example by integrating data assimilation and control decisions for the subsystem-level digital twins. The approach proposed in this paper is based on an offline analysis of the coupling strength and is limited to small variations of the shared quantities about nominal values. Once the federated architecture is deployed, the values of these shared quantities may fall outside the assumed range of variation. Future work should consider how to estimate the coupling strength online and adapt the architecture to fulfill the evolving communications needs.

Acknowledgments

This work was supported by the Air Force Office of Scientific Research (AFOSR) grants FA9550-23-1-0678 under the Space University Research Initiative (SURI), and FA9550-24-1-0327 under the Multidisciplinary University Research Initiative (MURI).

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] 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.

[3] 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.

[4] Torzoni, M., Tezzele, M., Mariani, S., Manzoni, A., and Willcox, K. E., "A digital twin framework for civil engineering structures," *Computer Methods in Applied Mechanics and Engineering*, Vol. 418, 2024, p. 116584. https://doi.org/10.1016/j.cma.2023.116584.

[5] 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. https://doi.org/10.48550/arXiv.2407.20490.

[6] Hernandez-Boussard, T., Macklin, P., Greenspan, E. J., Gryshuk, A. L., Stahlberg, E., Syeda-Mahmood, T., and Shmulevich, I., "Digital twins for predictive oncology will be a paradigm shift for precision cancer care," *Nat. Med.*, Vol. 27, No. 12, 2021, pp. 2065–2066. https://doi.org/10.1038/s41591-021-01558-5.

[7] Chaudhuri, A., Pash, G., Hormuth, D. A., Lorenzo, G., Kapteyn, M., Wu, C., Lima, E. A. B. F., Yankeelov, T. E., and Willcox, K., "Predictive digital twin for optimizing patient-specific radiotherapy regimens under uncertainty in high-grade gliomas," *Frontiers in Artificial Intelligence*, Vol. Volume 6 - 2023, 2023. https://doi.org/10.3389/frai.2023.1222612.

[8] Laubenbacher, R., Mehrad, B., Shmulevich, I., and Trayanova, N., "Digital twins in medicine," *Nat. Comput. Sci.*, Vol. 4, No. 3, 2024, pp. 184–191. https://doi.org/10.1038/s43588-024-00607-6.

[9] 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.

[10] Henao-Garcia, S., Kapteyn, M., Willcox, K. E., Tezzele, M., Castroviejo-Fernandez, M., Kim, T., Ambrosino, M., Kolmanovsky, I., Basu, H., Jirwankar, P., and Sanfelice, R., "Digital-Twin-Enabled Multi-Spacecraft On-Orbit Operations," *AIAA SciTech 2025 Forum*, 2025. https://doi.org/10.2514/6.2025-1432.

[11] Sisson, W., Karve, P., and Mahadevan, S., "Probabilistic Digital Twin for Damage-Aware Rotorcraft Path-Tracking Optimization Under Uncertainty," *AIAA Journal*, Vol. 0, No. 0, 2025, pp. 1–13. https://doi.org/10.2514/1.J064643.

[12] Braun, R., Gage, P., Kroo, I., and Sobieski, I., "Implementation and performance issues in collaborative optimization," *6th Symposium on Multidisciplinary Analysis and Optimization*, 1996. https://doi.org/10.2514/6.1996-4017.

[13] Bloebaum, C. L., Hajela, P., and Sobieszczanski-Sobieski, J., "Non-hierarchic system decomposition in structural optimization," *Eng. Optim.*, Vol. 19, No. 3, 1992, pp. 171–186. https://doi.org/10.1080/03052159208941227.

[14] Sobieszczanski-Sobieski, J., Agte, J. S., and Sandusky, R. R., "Bilevel Integrated System Synthesis," *AIAA Journal*, Vol. 38, No. 1, 2000, pp. 164–172. https://doi.org/10.2514/2.937.

[15] Sobieszczanski-Sobieski, J., Altus, T. D., Phillips, M., and Sandusky, R., "Bilevel Integrated System Synthesis for Concurrent and Distributed Processing," *AIAA Journal*, Vol. 41, No. 10, 2003, pp. 1996–2003. https://doi.org/10.2514/2.1889.

[16] Kim, H. M., Michelena, N. F., Papalambros, P. Y., and Jiang, T., "Target Cascading in Optimal System Design," *Journal of Mechanical Design*, Vol. 125, No. 3, 2003, pp. 474–480. https://doi.org/10.1115/1.1582501.

[17] Feddema, J., Lewis, C., and Schoenwald, D., "Decentralized control of cooperative robotic vehicles: theory and application," *IEEE Transactions on Robotics and Automation*, Vol. 18, No. 5, 2002, pp. 852–864. https://doi.org/10.1109/TRA.2002.803466.

[18] 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.

[19] Heimbigner, D., and McLeod, D., "A federated architecture for information management," *ACM Transactions on Information Systems*, Vol. 3, No. 3, 1985, p. 253–278. https://doi.org/10.1145/4229.4233.

[20] McMahan, B., Moore, E., Ramage, D., Hampson, S., and Arcas, B. A. y., "Communication-Efficient Learning of Deep Networks from Decentralized Data," *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research, Vol. 54, edited by A. Singh and J. Zhu, PMLR, 2017, pp. 1273–1282.

[21] San, O., Pawar, S., and Rasheed, A., "Decentralized digital twins of complex dynamical systems," *Scientific Reports*, Vol. 13, 2023. https://doi.org/10.1038/s41598-023-47078-9.

[22] Hua, Z., Djemame, K., Tziritas, N., and Theodoropoulos, G., "A Framework for Digital Twin Collaboration," *2024 Winter Simulation Conference (WSC)*, IEEE, 2024, pp. 3046–3057. https://doi.org/10.5555/3712729.3712982.

[23] Alexanderian, A., Gremaud, P. A., and Smith, R. C., "Variance-based sensitivity analysis for time-dependent processes," *Reliability Engineering & System Safety*, Vol. 196, 2020, p. 106722. https://doi.org/10.1016/j.ress.2019.106722.

[24] Smith, R. C., *Uncertainty Quantification: Theory, Implementation, and Applications, Second Edition*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2024. https://doi.org/10.1137/1.9781611977844.

[25] Sobol', I., "Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates," *Mathematics and Computers in Simulation*, Vol. 55, No. 1–3, 2001, p. 271–280. https://doi.org/10.1016/s0378-4754(00)00270-6.

[26] Saltelli, A., "Making best use of model evaluations to compute sensitivity indices," *Computer Physics Communications*, Vol. 145, No. 2, 2002, pp. 280–297. https://doi.org/10.1016/S0010-4655(02)00280-1.

[27] Owen, A. B., "Variance Components and Generalized Sobol' Indices," *SIAM/ASA Journal on Uncertainty Quantification*, Vol. 1, No. 1, 2013, pp. 19–41. https://doi.org/10.1137/120876782.

[28] Martins, J. R. R. A., and Ning, A., *Engineering Design Optimization*, Cambridge University Press, Cambridge, UK, 2022. https://doi.org/10.1017/9781108980647.

[29] Hwang, J. T., Lee, D. Y., Cutler, J. W., and Martins, J. R. R. A., "Large-Scale Multidisciplinary Optimization of a Small Satellite's Design and Operation," *Journal of Spacecraft and Rockets*, Vol. 51, No. 5, 2014, pp. 1648–1663. https://doi.org/10.2514/1.A32751.