

Value-Based Reinforcement Learning for Digital Twins in Cloud Computing

Bui, Van Phuc; Pandey, Shashi Raj; De Sant Ana, Pedro M.; Popovski, Petar

Published in:
IEEE International Conference on Communications

DOI (link to publication from Publisher):
[10.1109/ICC51166.2024.10622403](https://doi.org/10.1109/ICC51166.2024.10622403)

Publication date:
2024

Document Version
Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Bui, V. P., Pandey, S. R., De Sant Ana, P. M., & Popovski, P. (2024). Value-Based Reinforcement Learning for Digital Twins in Cloud Computing. *IEEE International Conference on Communications*, 1413-1418. Article 10622403. <https://doi.org/10.1109/ICC51166.2024.10622403>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Value-Based Reinforcement Learning for Digital Twins in Cloud Computing

Van-Phuc Bui, *IEEE Student Member*, Shashi Raj Pandey, *IEEE Member*, Pedro M. de Sant Ana, Petar Popovski, *IEEE Fellow*

Abstract—The setup considered in the paper consists of sensors in a Networked Control System that are used to build a digital twin (DT) model of the system dynamics. The focus is on control, scheduling, and resource allocation for sensory observation to ensure timely delivery to the DT model deployed in the cloud. Low latency and communication timeliness are instrumental in ensuring that the DT model can accurately estimate and predict system states. However, acquiring data for efficient state estimation and control computing poses a non-trivial problem given the limited network resources, partial state vector information, and measurement errors encountered at distributed sensors. We propose the REinforcement learning and Variational Extended Kalman filter with Robust Belief (REVERB), which leverages a reinforcement learning solution combined with a Value of Information-based algorithm for performing optimal control and selecting the most informative sensors to satisfy the prediction accuracy of DT. Numerical results demonstrate that the DT platform can offer satisfactory performance while reducing the communication overhead up to five times.

Index Terms—Digital twin, Reinforcement Learning, Internet of Things, Dynamic Systems, Sensor Networks.

I. INTRODUCTION

The Industry 4.0 smart manufacturing paradigm necessitates the acquisition of substantial volumes of real-time data, which emanate from a diverse array of wireless sensors [1]. In contrast to conventional simulation tools or optimization methodologies, digital twin (DT) models undergo a transformation of these extensive datasets into predictive models. The utilization of these models allows for the emulation of potential control strategies, which in turn supports real-time interactions and decision-making for system operators [2].

A network control system (NCS) representing the physical world, alongside a DT is situated either in the cloud, catering to extensive physical systems, or at the edge, tailored to local physical systems. The network comprises sensor devices and/or central/distributed units integral to a 5G system or beyond [3]. The acquired knowledge from the DT model serves twofold purposes: controlling the physical world; and providing monitoring and forecasting future states. Given the intricate interplay between communication and computation systems, devising a joint design strategy poses challenges in preserving the predictive efficacy of the DT model, and performing accurate control signals while simultaneously extending the operational lifespan of the network [4], [5]. There has been a substantial research on the use of Reinforcement

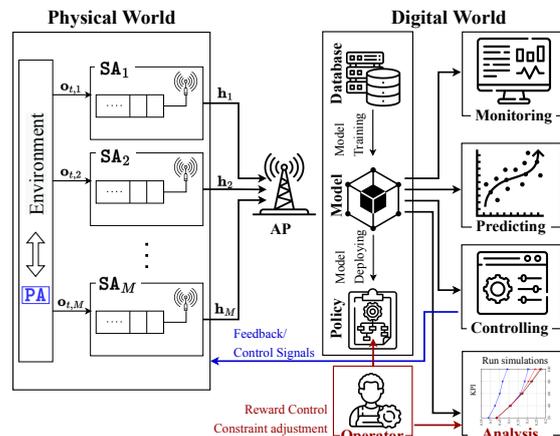


Fig. 1: The considered architecture with a Digital Twin (DT).

Learning (RL) at the DT to perform various tasks [4]. In [5], the issue of scheduling IoT sensors is examined through the use of Value of Information (VoI) while taking into account the limitations of communication and reliability. The ultimate goal is to reduce the Mean Squared Error (MSE) of state estimation, which is often hindered by imprecise measurements. Additionally, authors in [6] offers a possible resolution for scheduling sensing agents through the utilization of VoI, ultimately leading to enhanced accuracy in a variety of summary statistics for state estimation. These aforementioned works [4]–[6] and the reference therein, however, do not consider the influence of control performance in the physical world and the strategy of selecting effective sensing agents based on the reliability of estimates and latency requirements.

This paper addresses an optimization challenge that involves managing both the optimal control, accuracy of state estimation and power consumption in an NCS, illustrated in Fig. 1. Alongside this, a scheduling algorithm is introduced for SAs, which takes into account the quality of their observations as well as the communication constraints of the DT model. Our contributions are listed as follows: (i) we introduce a DT architecture tracking dynamic changes of the system parameters and controlling system dynamics; (ii) we propose an uncertainty control reinforcement learning framework that learn to perform actions while controlling the state’ uncertainty estimation; (iii) we formulate a novel optimization problem to efficiently schedule sensing agents for maintaining the confidence of DT’s system estimate while minimizing the energy cost under latency requirements. We then propose a VoI-based algorithm resulting in a practical and efficient solution in polynomial time. (iv) Numerical simulations were

V.-P Bui, S.R. Pandey, and P. Popovski (emails: {vpb, srp, fchi, petarp}@es.aau.dk) are all with the Department of Electronic Systems, Aalborg University, Denmark. P. M. de Sant Ana is with the Corporate Research, Robert Bosch GmbH, 71272 Renningen, Germany (email: Pedro.MaiadeSantAna@de.bosch.com). This work was supported by the Villum Investigator Grant “WATER” from the Velux Foundation, Denmark.

conducted to evaluate the algorithm's performance, which confirms that it surpasses other benchmarks in both controlling and power consumption while improving DT estimation error.

II. DIGITAL TWIN ARCHITECTURE

A. System Model

We adopt a DT architecture, as illustrated in Fig. 1, including a single primary agent (PA) and a set of sensing agents (SAs) denoted by $\mathcal{M} = \{1, 2, \dots, M\}$. These SAs are responsible for observing the environment and communicate with the *access point* (AP) through a wireless channel, which facilitates the construction of the DT model for the PA and operates in Frequency Division Duplexing (FDD) mode. We assume the communication link between AP and the cloud is perfect. The communication diagram is illustrated in Fig. 2, where T_{config} accounts for the configurable and computing time the DT. At the beginning of each *query interval* (QI), the DT model is required to update the active state of SAs, after which it estimates the full system state, updates policy, computes optimal control signal, schedules at most C SAs and applies a fusion algorithm at the AP. Once a control command is generated, the controller promptly transmits it through a downlink channel to the physical world. The application output for actuators control, such as motor drives, retrieves the most recently stored command values from memory and applies them to drive the system dynamics.

Each QI takes place at time instances $t \in \mathcal{T} = \{1, 2, \dots, T\}$. To maintain synchronization and consistency, these SAs periodically synchronize their DTs, including locations and power budget status, with the Cloud platform, ensuring a high degree of reliability and managing the overhead of synchronization. The PA engages in interactions with the environment, operating within a K -dimensional process as $\mathcal{K} = \{1, 2, \dots, K\}$. The state at the t -th QI is $\mathbf{s}_t = [s_t^1, s_t^2, \dots, s_t^K]^T$, and its evolution is described as:

$$\mathbf{s}_t = f(\mathbf{s}_{t-1}) + \mathbf{B}\mathbf{a}_{t-1} + \mathbf{u}_t, \quad \forall t \in \mathcal{T}, \quad (1)$$

Here, $f: \mathbb{R}^K \rightarrow \mathbb{R}^K$ denotes the state update function, \mathbf{a}_{t-1} is the control signal, and the matrix \mathbf{B} describes the how the control impacts the dynamics. $\mathbf{u}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_u)$ is the process noise. At QI t , a D -dimensional observation is received at SA_m ($m \in \mathcal{M}$) as $\mathbf{o}_{t,m} = g(\mathbf{s}_t) + \mathbf{w}_{t,m} \in \mathbb{R}^D$ ($D \leq K$) corresponding to the PA's state. To simplify the analysis, the observation $\mathbf{o}_{t,m}$ is assumed to be linearly dependent on the system state, which is $\mathbf{o}_{t,m} = \mathbf{H}_m \mathbf{s}_t + \mathbf{w}_{t,m}, \forall m \in \mathcal{M}$, with the observation matrix $\mathbf{H}_m \in \mathbb{R}^{D \times K}$ and the measurement noise $\mathbf{w}_{t,m} \sim \mathcal{N}(\mathbf{0}, \mathbf{C}_{w_m})$. The covariance matrices \mathbf{C}_u and \mathbf{C}_{w_m} are generally not diagonal. Let τ^{max} be the maximum tolerable latency for transmitting SA_m 's information, the system fulfills the application reliability at a QI t if

$$\mathbb{P}[\tau_{t,m} > \tau^{\text{max}}] \leq \varepsilon, \forall m \in \mathcal{M}, t \in \mathcal{T}, \quad (2)$$

with ε is a outage probability parameter depending on system characteristics [7].

B. Problem Formulation

The DT model's objective is to uphold a precise estimate of PA's state and offer the optimal sequence of actions to

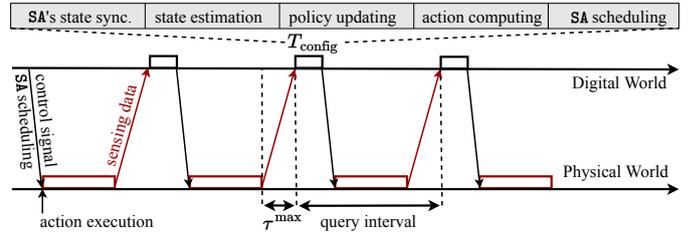


Fig. 2: The communication diagram.

be executed in the physical realm based on its beliefs about states. Herein, the predicted estimator $\hat{\mathbf{s}}_t$ of \mathbf{s}_t is modeled with $p(\mathbf{s}_t) \sim \mathcal{N}(\hat{\mathbf{s}}_t, \Psi_t), n \in \mathcal{N}$. The MSE of the estimator is

$$\text{MSE} = \mathbb{E}[\|\mathbf{s}_t - \hat{\mathbf{s}}_t\|_2^2], t \in \mathcal{T}. \quad (3)$$

Remark 1. We define the maximum acceptable standard deviation for feature $k \in \mathcal{K}$ as ξ_k . This corresponds to the following condition:

$$\sqrt{[\Psi_t]_k} \leq \xi_k, \forall k \in \mathcal{K}, \quad (4)$$

where $[\Psi_t]_k$ is the k -th element of the diagonal of Ψ_t .

Defining $V^\pi(s_0)$ as value function of controlling PA in (1) under control policy π and \mathbf{p}_t^{tx} as the transmitted power consumption to forward observations $\mathbf{o}_t \triangleq \{\mathbf{o}_{t,m}\}$ from SAs to AP at QI t , we interested in joint minimizing the sum power consumption and delivering optimal control signals. We introduce the ultimate goal $h(\{\mathbf{p}_t^{\text{tx}}\}, \{\mathbf{a}_t\}) = [V^\pi(\hat{\mathbf{s}}_0), -\sum_{t=0}^\infty \mathbf{p}_t^{\text{tx}}]^T$, then formulate the optimization problem as

$$\text{maximize}_{\{\mathbf{p}_t^{\text{tx}}\}, \{\mathbf{a}_t\}} h(\{\mathbf{p}_t^{\text{tx}}\}, \{\mathbf{a}_t\}) \quad (5a)$$

$$\text{subject to (1), (2), (4)}. \quad (5b)$$

It is noted that we consider a scenario where the belief vector can be enhanced through estimation techniques facilitated by DT prior to being utilized by RL to suggest the optimal action as an output. By employing this approach, the accuracy of noisy observations can be enhanced, enabling the agent to make more precise decisions.

In the following, we propose the REVERB (REinforcement learning and Variational Extended Kalman filter with Robust Belief) framework including two-step approach to address the problem (5): (i) we employ an uncertainty control RL algorithm to devise control actions for the physical world and effectively managing the state estimation errors; (ii) the VoI-based SA scheduling with optimal power control algorithm is utilized to identify the most significant SAs for observing its sensing signals, guided by the requirements from the RL model and DT. To further enhance the accuracy of estimated states and forecast the forthcoming system state, the Extended Kalman Filter (EKF) technique is revised and integrated.

III. UNCERTAINTY CONTROL POMDP

The control problem in (5) is considered as Partially Observable Markov Decision Process (POMDP) that expands upon the MDP by incorporating the sets of observations and observation probabilities to actual states because the provided observations only offer partial and potentially inaccurate information. In particular, a POMDP is presented by a 7-tuple $\langle \mathcal{S}, \mathcal{A}, \mathcal{O}, \mathcal{P}, \mathbf{0}, r, \gamma \rangle$. In particular, \mathcal{S} is a the finite set

of possible states, \mathcal{A} is a set of control primitives and \mathcal{O} denotes a set of possible observations. At a time instant t , the agent makes an action \mathbf{a}_t to move from state \mathbf{s}_t to \mathbf{s}_{t+1} with the transition probability $P = \mathbb{P}[\mathbf{s}_{t+1}|\mathbf{s}_t, \mathbf{a}_t]$. An observation \mathbf{o}_{t+1} received from SAs tracking system's state occurred with a probability $O = \mathbb{P}[\mathbf{o}_{t+1}|\mathbf{s}_{t+1}, \mathbf{a}_t]$. Also, as soon as the transition, the agent receives a numerical reward $r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1})$ verifying $r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{o}_{t+1}) \leq r^{\max}$. An agent do not know exactly its state at QI t and it maintains a estimate-vector $\hat{\mathbf{s}}_t$ describing the probability of being in a particular state $\mathbf{s}_t \in \mathcal{X}$. We define π as a policy of the agent that specifies an action \mathbf{a}_t based on its policy $\pi(\hat{\mathbf{s}}, \mathbf{a})$. With initial belief $\hat{\mathbf{s}}_0$, the expected future discounted reward for policy $\pi(\hat{\mathbf{s}}, \mathbf{a})$ is given as $V^\pi(\hat{\mathbf{s}}_0) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}) | \hat{\mathbf{s}}_0, \pi]$, where $0 < \gamma_t < 1$ is the discount factor.

At QI t , the estimated state vector is $\hat{\mathbf{s}}_t = [\hat{s}_t^1, \hat{s}_t^2, \dots, \hat{s}_t^K]^T \in \mathbb{R}^K$, which is governed by $\hat{\mathbf{s}}_t \sim p(\hat{\mathbf{s}}_t | \mathbf{s}_t; \boldsymbol{\eta}_t)$ as the conditional probability distribution function (pdf) of $\hat{\mathbf{s}}_t$ given \mathbf{s}_t . $\boldsymbol{\eta}_t = [\eta_{t,1}, \eta_{t,2}, \dots, \eta_{t,K}]^T \in \mathbb{R}^K$ is parameterized by the accuracy vector, whose each element $\eta_{t,k}$ indicates the accuracy of k -th process of estimated state $\hat{\mathbf{s}}_t$. In this work, we define $\eta_{t,k} = \frac{1}{\sqrt{[\Psi_t]_k}}, \forall k \in \mathcal{K}$. As $\eta_{t,k}$ increases, the confidence level of $\hat{s}_{t,k}$ also increases, facilitating the RL agent in making accurate decisions. Nevertheless, achieving high reliability of $\hat{s}_{t,k}$ necessitates low error of measurement SA or multiple observations, consequently amplifying both communication and processing expenses. Given \mathbf{s}_t and $\boldsymbol{\eta}_t$, the $\hat{s}_{t,k}$ is assumed exhibit statistical independence, meaning that we can express $p(\hat{\mathbf{s}}_t | \mathbf{s}_t; \boldsymbol{\eta}_t) = \prod_{k \in \mathcal{K}} p(\hat{s}_{t,k} | \mathbf{s}_t; \boldsymbol{\eta}_t)$ in terms of their factorization.

1) *Actor-critic DRL Algorithm*: For training the policy $\pi(\hat{\mathbf{s}}, \mathbf{a})$, we employ Proximal Policy Optimization (PPO) with the actor-critic structure at RL agents that involves dividing the model into two distinct components, thus harnessing the strengths of both value-based and policy-based methods [8]. Specifically, the actor is primarily responsible for estimating the policy, which dictates the agent's actions in a given state, and the critic is dedicated to estimating the value function predicting the expected future reward for a particular state or state-action pair. The actor's policy undergoes refinement based on the feedback provided by the critic. To address the joint design problem with the objective of optimizing actions while minimizing communication energy, we implement the Deep RL (DRL) approach within the DT cloud environment, where the action and reward are needed to be redefined.

2) *Action Space Reformulation*: We focus on the issue of optimizing the accuracy of the estimated state by the RL agent, enabling it to select $\eta_{t,k}$ on a continuous scale. The underlying motivation for the proposed framework is to unveil the inherent characteristics of the observation space in terms of the informational value that the observations offer for the given task. The formulation of the action vector structure within the RL agent implemented in DT is represented as

$$\mathbf{a}_t = [a_{t,1}, \dots, a_{t,Z}, \eta_{t,1}, \dots, \eta_{t,K}] \in \mathbb{R}^{ZK}, \quad (6)$$

where $\mathcal{Z}(|\mathcal{Z}| = Z)$ is the action space and $\{a_{t,z}\}_{z \in \mathcal{Z}}$ correspond to the control signals that exert an influence on the physical environment, enabling the agent to advance towards

its objective. Additionally, $\eta_{t,k} \in [0, \infty]$ denotes the accuracy selection pertaining to the estimated state.

3) *Reward Function Reformulation*: It is imperative for the RL agent to not only navigate towards the primary objective defined for the problem but also acquire the ability to regulate the acceptable level of accuracy $\{\eta_{t,k}\}_{k \in \mathcal{K}}$. Consequently, the goal-based reward r_t is transformed into an uncertainty-based reward as $\tilde{r}_t = f(r_t, \boldsymbol{\eta}_t)$, wherein $f(\cdot)$ is a monotonically non-decreasing function of r_t and $\boldsymbol{\eta}_t$. In scenarios where a direct cost function, denoted as $c_k(\cdot)$, exhibits an upward trend with the accuracy of the observation $\eta_{t,k}$, a suitable additive formulation can be employed. Specifically, the modified reward, \tilde{r}_t , can be expressed by

$$\tilde{r}_t = r_t + \kappa \sum_{k=1}^K c_k(\eta_{t,k}). \quad (7)$$

Here, $c_k(\eta_{t,k})$ represents a non-increasing function of $\eta_{t,k}$, and $\kappa \geq 0$ serves as a weighting parameter. Therefore, the primary objective of the agent is two-fold: to maximize the original reward while simultaneously minimizing the cost associated with the observations.

IV. VOI-BASED SAS SCHEDULING AND POWER CONTROL

In this section, our focus is to involve the scheduling of SAs based on three factors: (i) the acceptable level of accuracy $\boldsymbol{\eta}_t$ for the estimated state, as determined by the RL agent; (ii) the requirement pertaining to the accuracy of the DT model as in (4); and (iii) the communication resources, determined by the system capacity and latency requirements (2).

A. Sensing Agent Scheduling Problem

We formulate a combined SAs scheduling and power control problem, where our objective is to determine the SAs that should engage in transmission. We introduce arbitrary variables $\bar{\xi}_{t,k}^2$ indicates desired DT's error level of system's state at QI t . Given the reliability for the DT in (4) and the requisite level of accuracy $\boldsymbol{\eta}_t$ to uphold the precision of the RL model, the DT should meet the error constraints at QI t as

$$[\Psi_t]_k \leq \bar{\xi}_{t,k}^2 \triangleq \min \left\{ \xi_k^2, \frac{1}{\eta_{t,k}} \right\}, \forall t \in \mathcal{T}, k \in \mathcal{K}. \quad (8)$$

Initially, we establish the available SA set \mathcal{P}_t and the scheduling set \mathcal{Q}_t at QI t as $\mathcal{P}_t = \mathcal{M}$, $\mathcal{Q}_t = \emptyset$. Denoting the power allocation vector $\mathbf{p}_t^{\text{tx}} = \{p_{t,m}^{\text{tx}}\}_{m \in \mathcal{M}}$ as the variable, we formulate the optimization problem with specifications as

$$\mathbf{p}_t^{\text{tx}*} = \underset{\mathbf{p}_t^{\text{tx}}}{\operatorname{argmin}} (1 - \alpha) \sum_{k \in \mathcal{K}} \max \left\{ \frac{[\Psi(n)]_k}{\bar{\xi}_k^2} - 1, 0 \right\} + \alpha \sum_{m \in \mathcal{P}(n)} p_m^{\text{tx}}(n) \quad (9a)$$

$$\text{subject to } |\mathcal{Q}_t| \leq C, \quad (9b)$$

$$\mathbb{P}[\tau_{t,m} > \tau^{\max}] \leq \varepsilon, \forall m \in \mathcal{M}, \quad (9c)$$

wherein the non-negative parameter $\alpha \in [0, 1]$ represents the relative weight to accuracy and energy efficiency within the underlying objective function. It is observed that objective (9a) represents a relaxation of constraint (8) due to its dependence

on practical conditions, i.e., in situations where the error surpasses a certain threshold, even querying all sensor data fails to guarantee the desired level of reliability $\bar{\xi}_{t,k}^2$. We note that obtaining observations from additional sources SAs leads to an enhancement in estimation accuracy; however, this improvement comes at the cost of compromising energy efficiency. For those particular SAs that exhibit considerable errors in their measurements or possess features that do not significantly contribute to meeting the confidence requirements of the PA (i.e., those with a low VoI), measuring and transmitting observations leads to an unnecessary expenditure of energy. The constraints given by (9b) and (9c) are required in ensuring communication capacity and the reliable execution of latency satisfaction within a FDMA uplink slot. It is important to note that the optimization problem presented as (9) is inherently non-convex due to the non-convex nature of the objective function (9a), as well as the constraints (9b), (9c). Furthermore, the node selection aspect renders the problem analogous to the classic NP-hard knapsack problem. To derive an efficient suboptimal solution, a heuristic algorithm based on EKF is employed.

B. Sensing Agent Selection with Extended Kalman Filter

Due to the complexity of sensor selection based on VoI, we adopt a heuristic approach with primary concept guiding the resolution of (8) is to ensure that, during each QI t the minimum necessary number of SAs is selected for transmission. This selection aims to sustain the desired level of certainty in estimating the state s_t while saving communication resource for other purposes. The expression for the Minimum Mean Square Error (MMSE) estimator applied to a KF is provided in [9, Eq. (1)]. In this work, for aiming to minimize the (weighted) variance of state components, we employ the Extended Kalman estimator, which is common in the IoT literature [10]. We also assume that the virtual environment possessing complete awareness regarding the process statistics, encompassing the update function $f(s)$ as well as the noise covariance matrices. The primary strategy to solve (9) involves selecting the minimum number of SAs to transmit at each QI t in order to maintain the required level of estimation certainty for state s_t as specified by $\{\bar{\xi}_{t,k}^2\}$. Our proposed algorithm are summarized in Algorithm 1, effectively addresses problem (9).

In the dynamic setup, the initial state s_t is considered a random vector characterized by a specific mean $\mathbb{E}[s_t] = \mu_{s_0}$ and covariance matrix $\text{Cov}[s_0] = \mathbf{C}_{s_0}$. \mathcal{Q}_t is initialized as an empty set due to the absence of any prior information. The EKF then calculates the estimation errors for the belief

$$\hat{s}_t \sim \mathcal{CN}(\mu_{\hat{s}_t}, \Psi_{\hat{s}_t}^{\text{pr}}) \quad (10)$$

at the PA based on prior updates \hat{s}_{t-1} as described by

$$\Psi_{\hat{s}_t}^{\text{pr}} = \mathbf{P}\Psi_{\hat{s}_{t-1}}\mathbf{P}^T + \mathbf{C}_{\mathbf{u}_t}, \quad (11)$$

where the Jacobian matrix $\mathbf{P} = \mathcal{J}\{f(s_{t-1})\}$ linearizes the nonlinear model of $f(s_{t-1})$. Subsequently, for a given set of error variance qualities $\{\bar{\xi}_{t,k}^2\}_{t \in \mathcal{T}, k \in \mathcal{K}}$, the conditions specified in (8) lead to two potential cases: (i) If (8) are satisfied for all $k \in \mathcal{K}$, the DT model achieves bounds without requiring

SAs' observations. The prior update suffices to establish the necessary confidence in the estimate, resulting in an empty set for $\mathcal{Q}_t^* = \emptyset$. (ii) If any of conditions is violated, at least one process feature lacks sufficient accuracy, the acquisition of the corresponding observations becomes essential to enhance the estimation process, as dictated by the scheduling approach implemented in our proposed heuristic.

In the first scenario, the the belief \hat{s}_t can be achieved through the EKF blind update as

$$\hat{s}_t = \hat{s}_t^{\text{pr}} = f(\hat{s}_{t-1}) + \mathbf{B}\mathbf{a}_{t-1} + \mathbf{u}_t, \forall t \in \mathcal{T}. \quad (12)$$

In the second case, Algorithm 1 is utilized to identify the SAs with the highest VoI for querying their observations. In order to identify the most suitable candidate feature $s_{t,k}^*$, where $k \in \mathcal{K}$, an optimization problem is formulated at the i -th iteration as

$$s_{t,k}^* = \underset{s_{t,k} \in s_t}{\text{argmax}} \quad [\Psi_t^{(i)}]_k / \bar{\xi}_{t,k}^2 \quad (13a)$$

$$\text{subject to} \quad \mathbf{S}\mathbf{A}_m \in \mathcal{P}_t, \forall m \in \mathcal{M}, \quad (13b)$$

$$\mathbf{S}\mathbf{A}_m \rightarrow s_{t,k}, \forall m \in \mathcal{M}, \quad (13c)$$

where the notation $\mathbf{S}\mathbf{A}_m \rightarrow s_{t,k}$ signifies that the $\mathbf{S}\mathbf{A}_m$ measures feature $s_{t,k}$. We note that at the i -th iteration, if any constraint in (8) is still unmet and $|\mathcal{Q}_t| < C$, $|\mathcal{P}_t| > 0$, there is room for scheduling new SAs to join \mathcal{Q}_t . According to constraint (13), feature $s_{t,k}^*$ is selected only if at least one $\mathbf{S}\mathbf{A}_m \in \mathcal{P}_t$ can provide coordinating observations. Then, $\mathbf{S}\mathbf{A}_m^* \in \mathcal{P}_t$ measuring feature $s_{t,k}^*$ with the minimum error covariance is chosen to send its measurement. The scheduled and available SA sets \mathcal{Q}_t and \mathcal{P}_t are updated as

$$\mathcal{Q}_t \leftarrow \mathcal{Q}_t \cup \{\mathbf{S}\mathbf{A}_m^*\}; \mathcal{P}_t \leftarrow \mathcal{P}_t \setminus \{\mathbf{S}\mathbf{A}_m^*\}. \quad (14)$$

\mathbf{H}_t and $\mathbf{C}_{\mathbf{w}_t}$ are the combination observation and covariance matrices, which are respectively formulated as

$$\mathbf{H}_t = [\mathbf{H}_1; \mathbf{H}_2; \dots; \mathbf{H}_{|\mathcal{Q}_t|}], \quad (15)$$

$$\mathbf{C}_{\mathbf{w}_t} = \text{diag}[\mathbf{C}_{\mathbf{w}_1}, \mathbf{C}_{\mathbf{w}_2}, \dots, \mathbf{C}_{\mathbf{w}_{|\mathcal{Q}_t|}}], \quad (16)$$

where \mathbf{H}_m is the observation matrix of the $\mathbf{S}\mathbf{A}_m (m \in \mathcal{Q}_t)$. From here, we can compute the EFK gain by

$$\mathbf{K}_t = \Psi_{\hat{s}_t}^{\text{pr}} \mathbf{H}_t^T (\mathbf{C}_{\mathbf{w}_t} + \mathbf{H}_t \Psi_{\hat{s}_t}^{\text{pr}} \mathbf{H}_t^T)^{-1}, \quad (17)$$

The posterior error covariance matrix is derived by

$$\Psi_{\hat{s}_t}^{\text{pos}} = (\mathbf{I} - \mathbf{K}_t \mathbf{H}_t) \Psi_{\hat{s}_t}^{\text{pr}}, \quad (18)$$

The iterative loop continues while all three conditions are true:

$$\left\{ \begin{array}{l} \{|\mathcal{Q}_t^*|\} < C; \\ \exists [\Psi_{\hat{s}_t}^{\text{pos}}]_k \geq \bar{\xi}_k^2; \\ \exists s_{t,k}^* \text{ as a solution in (13)}. \end{array} \right. \quad (19)$$

It makes intuitive see that the loop repeats for at maximum C iterations before terminating. The posterior update is then

$$\hat{s}_t^{\text{pos}} = \hat{s}_t^{\text{pr}} + \mathbf{K}_t(\mathbf{o}_t - \mathbf{H}_t \hat{s}_t^{\text{pr}}), \quad (20)$$

where \mathbf{o}_t represents the combination of received SA observations. Accordingly, we update $\hat{s}_t = \hat{s}_t^{\text{pos}}$. Our approach ensures long-term balance between state certainty and communication cost with respect to the PA's requirements, despite

the local scheduling solution provided by SA for different QIs.

C. Optimal Power Scheduling

Both DL and UL transmissions are subject to potential latency during data delivery, which is influenced by the channel quality and the scheduled transmission resource. The channel between SA_{*m*} and AP is modeled as Rician channel with strong LoS link and small-scale fading with rich scattering, where the instantaneous Signal-to-Noise (SNR) ratio is modeled as

$$\gamma_{t,m} = \frac{\Gamma p_{t,m}^{\text{tx}}}{d_m^\alpha W N_0} \mathcal{G}_m, \forall m \in \mathcal{M}, \quad (21)$$

with Γ is constant depending on the system parameter (operating frequency and antenna gain); $p_{t,m}^{\text{tx}}$ is the transmitted power, d_m represents the distance between device m and AP. α is the path loss exponent; W is the allocated bandwidth and N_0 stands for noise power. \mathcal{G}_m represents the fading power with expected value $\bar{\mathcal{G}} \triangleq \mathbb{E}[|\mathcal{G}_m|^2] = 1$. Herein, we adopt \mathcal{G}_m as Rician distribution with a non-central chi-square probability distribution function (PDF) which is expressed as [11]

$$f_{\mathcal{G}}(w) = \frac{(G+1)e^{-G}}{\bar{\mathcal{G}}} e^{-\frac{(G+1)w}{\bar{\mathcal{G}}}} I_0\left(2\sqrt{\frac{G(G+1)w}{\bar{\mathcal{G}}}}\right), \quad (22)$$

where $w \geq 0$ and $I_0(\cdot)$ denotes the zero-order modified Bessel function of the first kind, while G represents the Rician factor, signifying the ratio between the power within the Line-of-Sight (LoS) component and the power distributed among the non-LoS multipath scatters. We use Shannon's bound to achieve the channel capacity $R_{t,m}$ of every link, which is $R_{t,m} = W \log_2(1 + \gamma_{t,m})$. The optimal transmission power for each scheduled SA is computed through the following lemma.

Lemma 1. *Given the Rician factor G , the specific transmission bandwidth W and the distance d_m , the optimal allocated power $p_{t,m}^*$ of SA_{*m*} to satisfy (2) is lower bounded by*

$$p_{t,m}^* = \frac{2WN_0(1+G)(2^{D/(\tau^{\max}W)} - 1)}{y_Q^2 d_m^{-\alpha} \Gamma}, \quad (23)$$

where $y_Q = \sqrt{2G} + \frac{1}{2Q^{-1}(\varepsilon)} \log\left(\frac{\sqrt{2G}}{\sqrt{2G} - Q^{-1}(\varepsilon)}\right) - Q^{-1}(\varepsilon)$, D is the length of data packet, and $Q^{-1}(\cdot)$ is the inverse Q -function.

Proof. The outage probability (2) is formulated as

$$\mathbb{P}[\Delta_{t,m} > \tau^{\max}] = \mathbb{P}\left[W \log_2(1 + \gamma_{t,m}) < \frac{D}{\tau^{\max}}\right] \quad (24)$$

$$= \mathbb{P}[\gamma_{t,m} < 2^{\frac{D}{W\tau^{\max}}} - 1] \quad (25)$$

$$\triangleq 1 - Q(x_\tau, y_\tau) \leq \varepsilon, \quad (26)$$

where $x_\tau = \sqrt{2G}$ and

$$y_\tau = \sqrt{2(G+1)(2^{\frac{D}{W\tau^{\max}}} - 1)\bar{\mathcal{G}}/\gamma_{t,m}}, \quad (27)$$

with $Q(x_\tau, y_\tau)$ as the the first order Marcum Q-function. At the maximum tolerable value of ε , then $y_Q = Q^{-1}(x_Q, 1 - \varepsilon)$ formed as the inverse Marcum Q-function respecting to second argument. In this study, we consider a Rician channel with strong LoS component, i.e., $G > G_0^2/2$ and $Q^{-1}(\varepsilon) \neq 0$, which yields the approximated form of y_τ as [7, Lemma 1]:

Algorithm 1 SA scheduling algorithm for problem (9)

Input: $\mathbf{b}_{0,\mathbf{o}_0}, \boldsymbol{\mu}_{\mathbf{u}_0}, C_{\mathbf{u}_0}$ Available uplink slots C , The state and requirement certainty $(\mathbf{s}, \{\xi_k^2\})$

Output: The scheduled user set $\{Q_t^*\}$; their belief $\{\hat{\mathbf{s}}_t^*, \Psi_t^*\}$, and the associated transmit power $\{p_{t,m}^{\text{tx}}\}$

- 1: Initial $Q_t = \emptyset$
- 2: Compute the prior errors Ψ_t^{pr} as in (11)
- 3: **if** $[\Psi_t^{\text{pr}}]_k \leq \xi_k^2, \forall k$ **then**
- 4: Compute $\hat{\mathbf{s}}_t^{\text{pr}}$ as in (12)
- 5: Update $\hat{\mathbf{s}}_t^* = \hat{\mathbf{s}}_t^{\text{pr}}$ and $\Psi_{\hat{\mathbf{s}}_t}^* = \Psi_{\hat{\mathbf{s}}_t}^{\text{pr}}$
- 6: **else**
- 7: Set $i = 1$
- 8: **while** conditions (19) hold **do**
- 9: Update $s_{t,k}^*$ by solving (13)
- 10: Update Q_t and \mathcal{P}_t as in (14)
- 11: Update the $\mathbf{K}_t, \mathbf{H}_t$ and $\mathbf{C}_{\mathbf{w}_t}$ as in (15), and (16)
- 12: Compute $\Psi_{\hat{\mathbf{s}}_t}^{\text{pos}}$ as in (18)
- 13: Set $i = i + 1$
- 14: **end while**
- 15: Update $Q_t^* = Q_t$
- 16: Compute $\Psi_{\hat{\mathbf{s}}_t}^* = \Psi_{\hat{\mathbf{s}}_t}^{\text{pos}}$ and $\hat{\mathbf{s}}_t^* = \hat{\mathbf{s}}_t^{\text{pos}}$ as in (18) and (20)
- 17: Update η_t^* , and power transmission $\mathbf{p}_t^{\text{tx}*} = \{p_{t,m}^{\text{tx}}\}$ as in (23)
- 18: **end if**

TABLE I: Simulation Parameters

Parameter	Value	Parameter	Value
Carrier frequency (f_c)	2.4 GHz	Maximum latency (τ^{\max})	5 ms
Bandwidth (W)	5 MHz	Channel noise power	-11.5 dBm
DT required error variances ($\xi_{\text{pos}}^2, \xi_{\text{vel}}^2$)	(0.01, 0.001)	Max. distance (d^{\max})	20 m
Max. connection (C)	10	Learning rate (both actor and critic)	$1e^{-3}$
Rician factor (G)	15 dB	Outage probability factor (ε)	$1e^{-5}$
System parameter (Γ)	1	Data packet size (D)	1024 bits

$y_\tau = \sqrt{2G} + \frac{1}{2Q^{-1}(\varepsilon)} \log\left(\frac{\sqrt{2G}}{\sqrt{2G} - Q^{-1}(\varepsilon)}\right) - Q^{-1}(\varepsilon)$, with G_0 is the intersection of the sub-functions at $x_\tau > \max[0, Q^{-1}(\varepsilon)]$. Plug this into (27), we obtain (23) and complete the proof. \square

V. NUMERICAL RESULTS

We examine our proposed REVERB algorithm with the MountainCarContinuousv0 environment from the OpenAI Gym [12]. The state vector $\mathbf{s}_t = [x_t, \dot{x}_t]^T$ includes position x_t and velocity \dot{x}_t . The observation matrices are given by $\mathbf{H}_{\text{pos}} = [1, 0]$; and $\mathbf{H}_{\text{vel}} = [0, 1]$, respectively for the position and velocity. Other important parameters are included in Table I. In the DT system, the agent makes decisions concerning the applied force $a_t \in [-1, 1]$ on the car and selects an accuracy level denoted by $\boldsymbol{\eta}_t = [\eta_{t,1}, \eta_{t,2}]^T$. The original reward for each QI in the environment is $r_t = -0.1 \times a_t^2$. In our system, we adopt a modified reward indicated in (7) as

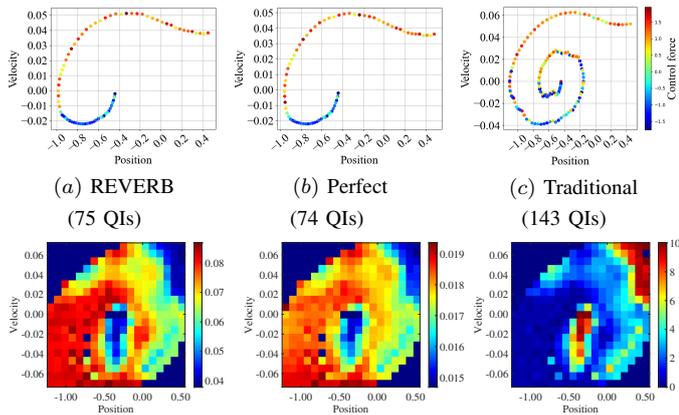
$$\hat{r}_t = r_t + \kappa \times \left(\frac{1}{2} \sum_{i=1}^2 \eta_{t,i}\right), \quad (28)$$

where the positive weighted parameter $\kappa = 5 \times 10^{-6}$. Additionally, the original environment includes a termination reward when the car successfully reaches the target position at 0.45. The evolution of the state in (1) is defined with

$$\mathbf{f}(\mathbf{s}_{t-1}) = \begin{bmatrix} \dot{x}_{t-1} \\ -\varphi \cos(3x_{t-1}) \end{bmatrix}, \quad \mathbf{B} = [0, \vartheta]^T, \quad (29)$$

where the constants $\varphi = 0.0025$ and $\vartheta = 0.0015$. M SAs are assuming placed randomly within an area where the maximum distance d^{\max} from SAs to AP.

We compare our proposed REVERB and four other benchmarks: *Perfect* allows DT to get the observation of the next state without any error; *Cost-based* greedy indicates that the



(d) Position uncertainty (e) Velocity uncertainty (f) No. selected SAs
 Fig. 3: Performance evaluation: (a), (b) and (c) compare the trajectory with coordinated force of different schemes; (d) and (e) are position and velocity uncertainty level ($\{\bar{\xi}_{t,k}\}$) of REVERB; and (f) represents number of selected SAs under position-velocity coordinate.

AP queries all nearest SAs based on ascending order of distance from AP to SA at each QI. *Error-based* greedy [5] is similar to *Cost-based* but it relies on the decreasing confidence levels of SAs. In the greedy benchmarks is $|Q_t^*| = C$. We also consider the *Traditional* scheme [13] where the RL agent gets noise observation from one SA every QI without Algorithm 1. All schemes are conducted under 1000 Monte Carlo simulations.

Figure 3(a), (b) and (c) compare the trajectory evolution with coordinated force of REVERB, *Perfect*, and *Traditional* schemes. It is observed that REVERB’s trajectory is close to the *Perfect* when using only 1 step more to reach the goal. On the contrary, the trained network with the *Traditional* must use up to $2\times$ number of steps to reach the goal. These results confirm REVERB’s reliability regarding force adjustment under a noisy environment. Fig. 3(d)-(f) depict the noise levels and number of selected SAs over the whole observation space. When comparing Fig. 3(d)-(e) to Fig. 3(f), we notice that DT gathers more data as uncertainty increases and the agent approaches the target (position ≥ 0.45). In scenarios where the noise level is high but the agent is still far from the target (position < -0.5), requiring additional observations to improve control efficiency is considered ineffective. It is worth noting that this ineffectiveness does not have an adverse impact on the original task’s performance.

Figure 4 illustrates different schemes’ power consumption and Mean-root-mean-square-error (MRMSE). REVERB proves to be the most efficient, consuming the least power and achieving the lowest MRMSE compared to other baselines. On average, REVERB consumes 1.09 [W] to reach the goal, compared to 5.98 [W] of the *Error-based* algorithm. Even though *Traditional* one only queries 2 SAs per QI, the power consumption is up to 5.4 [W] because this scheme takes more QIs to reach the goal, as in Fig. 3(c). At the same time, REVERB’s MRMSE is always maintained at approximately the same low level as the *Cost-based* and *Error-based* methods.

The results achieved in Fig. 4 are explained by a snapshot of the uncertainty evolution and the strategic selection of SAs in

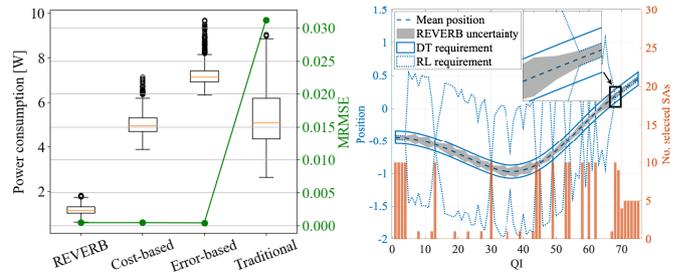


Fig. 4: Power consumption & MRMSE of different schemes.

Fig. 5: Snapshot of uncertainty evolution and no. selected SAs vs. QI.

Fig. 5, based on their contributions to the DT’s performance. It is noteworthy that the management of REVERB’s uncertainty is subject to the control of both DT and RL requirements as flowing (8). The requisites imposed upon the RL agent’s behavior are primarily administered through the reward function delineated in (28). Hence, DT only requests more observations when the DT threshold is surpassed or when the RL agent necessitates high precision, typically when the agent is nearing its goal and precise force control is imperative.

VI. CONCLUSIONS

This work introduced the DT framework for reliability monitoring, predicting and controlling of a communication system. Under the latency constraint, the DT platform was shown to obtain more reliable control and trajectory tracking than conventional methods while significant saving communication cost. This result is achieved thanks to the proposed uncertainty control POMDP scheme combined with an efficient algorithm selecting subsets of partial SAs to meet the requirements in the confidence of state estimation. Future study could explore the long-term impacts of scheduling decisions in complex systems and incorporate deep learning-based estimators.

REFERENCES

- [1] Y. Tang *et al.*, “Tracking control of networked multi-agent systems under new characterizations of impulses and its applications in robotic systems,” *IEEE Trans. Ind. Electron.*, vol. 63, no. 2, pp. 1299–1307, 2015.
- [2] S. Mihai *et al.*, “Digital twins: A survey on enabling technologies, challenges, trends and future prospects,” *IEEE Commun. Surveys Tut.*, vol. 24, no. 4, pp. 2255–2291, 2022.
- [3] D. Tse and P. Viswanath, *Fundamentals of wireless communication*. Cambridge university press, 2005.
- [4] C. Ruah, O. Simeone, and B. Al-Hashimi, “A bayesian framework for digital twin-based control, monitoring, and data collection in wireless systems,” *IEEE J. Sel. Areas Commun.*, 2023.
- [5] V.-P. Bui, S. R. Pandey, F. Chiariotti, and P. Popovski, “Scheduling policy for value-of-information (voi) in trajectory estimation for digital twins,” *IEEE Commun. Lett.*, vol. 27, no. 6, pp. 1654–1658, 2023.
- [6] F. Chiariotti, A. E. Kalør, J. Holm, B. Soret, and P. Popovski, “Scheduling of sensor transmissions based on value of information for summary statistics,” *IEEE Networking Lett.*, vol. 4, no. 2, pp. 92–96, 2022.
- [7] M. M. Azari *et al.*, “Ultra reliable uav communication using altitude and cooperation diversity,” *IEEE Trans. Commun.*, vol. 66, no. 1, pp. 330–344, 2018.
- [8] V. Konda and J. Tsitsiklis, “Actor-critic algorithms,” *Proc. NIPS*, vol. 12, 1999.
- [9] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *J. Basic Eng.*, vol. 82, no. 1, pp. 35–45, 03 1960.
- [10] Y. Huang, W. Yu, E. Ding, and A. Garcia-Ortiz, “EPKF: Energy efficient communication schemes based on Kalman filter for IoT,” *IEEE Internet Things J.*, vol. 6, no. 4, pp. 6201–6211, 2019.
- [11] M. K. Simon and M.-S. Alouini, “Digital communications over fading channels (mk simon and ms alouini: 2005)[book review],” *EEE Trans. Inf. Theory*, vol. 54, no. 7, pp. 3369–3370, 2008.
- [12] A. W. Moore, “Efficient memory-based learning for robot control,” University of Cambridge, Computer Laboratory, Tech. Rep., 1990.
- [13] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.