



Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Relativistic Digital Twin: Bringing the IoT to the future

Luca Sciuillo^{a,b,*}, Alberto De Marchi^c, Angelo Trotta^a, Federico Montori^{a,b}, Luciano Bononi^{a,b}, Marco Di Felice^{a,b}

^a Department of Computer Science and Engineering, University of Bologna, Italy

^b Advanced Research Center on Electronic Systems, University of Bologna, Italy

^c Institute of Applied Mathematics and Scientific Computing, Department of Aerospace Engineering, University of the Bundeswehr Munich, Germany

ARTICLE INFO

Keywords:

Internet of Things (IoT)
Digital Twin (DT)
Web of Things (WoT)
Machine learning

ABSTRACT

Complex IoT ecosystems often require the usage of Digital Twins (DTs) of their physical assets in order to perform predictive analytics and simulate what-if scenarios. DTs are able to replicate IoT devices and adapt over time to their behavioral changes. However, DTs in IoT are typically tailored to a specific use case, without the possibility to seamlessly adapt to different scenarios. Further, the fragmentation of IoT poses additional challenges on how to deploy DTs in heterogeneous scenarios characterized by the usage of multiple data formats and IoT network protocols. In this paper, we propose the Relativistic Digital Twin (RDT) framework, through which we automatically generate general-purpose DTs of IoT entities and tune their behavioral models over time by constantly observing their real counterparts. The framework relies on the object representation via the Web of Things (WoT), to offer a standardized interface to each of the IoT devices as well as to their DTs. To this purpose, we extended the W3C WoT standard in order to encompass the concept of behavioral model and define it in the Thing Description (TD) through a new vocabulary. Finally, we evaluated the RDT framework over two disjoint use cases to assess its correctness and learning performance, *i.e.*, the DT of a simulated smart home scenario with the capability of forecasting the indoor temperature, and the DT of a real-world drone with the capability of forecasting its trajectory in an outdoor scenario. Experiments show that the generated DT can estimate the behavior of its real counterpart after an observation stage, regardless of the considered scenario.

1. Introduction

The near future of the Internet of Things (IoT) paradigm is intertwined with the novel concept of Digital Twins (DTs). DTs are digital replicas of physical assets that evolve over time in order to mimic their behavior, their reaction to events and their internal structure as closely as possible through computational models [1]. In previous years, other entities, like simulators, emulators, digital shadows, mirror systems, or avatars, assumed this role. DTs are different from them in that (i) they constantly receive inputs from the mirrored appliance(s) to adapt to its behavioral changes, as well as (ii) they feed back their outputs into the real environment to provide guidance [2]. The capability of DTs to reproduce real devices and entities makes them extremely important for a variety of IoT scenarios that take advantage of predictive and what-if analytics. In fact, DTs find their natural applicability in manufacturing and Industry 4.0 scenarios, where putting the physical asset at risk results in a massive loss in terms of resources [3]. Because of their crucial role, DTs are, most of the time, tailored to the appliance that

they reproduce. This is obviously a safe choice, however, it results in additional engineering cost and effort when it is necessary to adapt the computational model to a new use case or a new asset [4]. As a matter of fact, nowadays there is a profound lack of general-purpose solutions that can offer adaptable DTs or automatically generate computational models able to synthesize heterogeneous scenarios. It is clear that a zero-knowledge solution is nearly impossible to obtain, but we can envision an automatic process that is able to fit a variety of scenarios given very little initial information. Moreover, the age-old problem of IoT fragmentation still persists in most of the IoT ecosystems and exacerbates the difficulties of introducing such automation for DTs, as devices may use different technologies and protocols to expose their data and services in a highly dynamic environment [5]. To overcome this issue, several solutions have been proposed over the last decade; one of the most supported is the Web of Things (WoT), recently standardized by the World Wide Web Consortium (W3C) [6]. The WoT abstracts each IoT device into a Web Thing (WT), a digital entity

* Corresponding author at: Department of Computer Science and Engineering, University of Bologna, Italy.

E-mail addresses: luca.sciuillo@unibo.it (L. Sciuillo), alberto.demarchi@unibw.de (A. De Marchi), angelo.trotta5@unibo.it (A. Trotta), federico.montori2@unibo.it (F. Montori), luciano.bononi@unibo.it (L. Bononi), marco.difelice3@unibo.it (M. Di Felice).

<https://doi.org/10.1016/j.future.2023.12.016>

Received 13 April 2023; Received in revised form 11 December 2023; Accepted 16 December 2023

Available online 19 December 2023

0167-739X/© 2023 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

that exposes its functionalities (defined as *affordances* in the official documentation [6]) via standardized Web technologies and semantics that vaguely resemble the REST concepts. WTs are described via a Thing Description (TD), a machine-readable JSON-LD document that can be shared with other WTs in order to automate interactions in a WoT ecosystem. A number of Industrial IoT frameworks are now dedicated to embracing the concept of DTs, at the same time leveraging integration with the WoT. An example is Eclipse Ditto [7] or Eclipse Arrowhead [8], which presents a one-to-one bi-directional translation between service records and WTs [9]. These solutions exploit the WoT for abstracting heterogeneous DTs of physical assets and integrating them within a single application ecosystem, however, they do not take into account a homogenized representation of their behavioral models, leaving this aspect to the third parties.

In this paper we propose a solution to the two above-mentioned shortcomings of the current landscape of DTs in IoT: either (i) DTs are too specialized for the use case they operate upon, or (ii) interoperability-oriented IoT frameworks do not have the necessary expressiveness to encompass their model description. Specifically, we propose a framework that is based on the WoT, where all entities are abstracted into a single WT, and all interactions between entities are encoded as affordances. However, as an extension to the W3C WoT standard and as the main contribution of our work, each of the TDs is enriched with a parametric behavioral model that characterizes the dynamic of the target IoT appliance over time. Based on it, our framework is able to generate a DT for each of the represented WTs that presents the exact same affordances; in addition, by periodically monitoring the behavior of the real appliance, the framework fits the model parameters to mimic the physical asset over time. Eventually, the generated DTs will be indistinguishable from the real assets, as they both expose the exact same interface to other actors in the WoT ecosystem. A fully operational scenario aims to advocate a paradigm shift, where virtual entities can be queried — and acted upon — with a predictive connotation, by running them at future instants in time, as the DTs will be able to foretell the behavior of the asset. This vision is inspired by the “Twin Paradox” [10] in which one of the twins is sent forward in time and is eventually older than the other twin; in the same way here, the DT can interact with the other elements in the system without having to stick to the current moment. This enables the construction of what-if scenarios, where we can spawn multiple DTs of the same WT, each of them with its own timeline. Each DT is then acted upon predictively, allowing us to “bring the IoT to the future”. Despite the shared birth between the DT and its physical counterpart, and because of the independent timeline of virtual entities, we name our framework the Relativistic Digital Twin (RDT) framework.

More precisely, this paper makes the following three contributions:

- We propose an extension to the WoT TD syntax, in order to introduce the concept of behavioral model as well as its interactions; because certain entities in an ecosystem may influence each other, we also take into account that certain behavioral models may be dependent on the properties of other IoT objects in the same environment.
- We propose the RDT framework, which is able to generate DTs of each WT that includes a behavioral model in the ecosystem, by learning over time the model parameters through constant observation of the WT properties. At the end of the learning phase, the RDT framework is able to generate the Digital Twin Web Thing (DTWT), *i.e.*, new WTs which are functionally equivalent to the real-world ones but are attached to the behavioral model rather than to the IoT physical asset.
- We evaluate our proposal over two conceptually different use cases: a DT of a simulated scenario including the representation of a room with adjustable temperature and multiple sensors/actuators, and a DT of a real quadcopter drone able to reproduce its kinematic.

The rest of the paper is organized as follows: Section 2 reviews the state-of-the-art concerning DT frameworks and the WoT, Section 3 introduces the system model used in this work along with the scenario description and assumptions declaration, Section 4 outlines in detail how we extend the TD syntax to include behavioral models and their interactions, Section 5 illustrates the framework architecture and how we implemented all the components, Section 6 describes the learning module and how each behavioral model is fitted to the real observations, Section 7 introduces the two use cases on which we tested the framework and their results, and, finally, Section 8 concludes the paper, outlining future works.

2. Related work

2.1. Digital twin and learning

Adapting to behavioral changes is a prerogative of DTs, however, in a variety of scenarios, this is subject to multi-objective optimization problems that are mathematically intractable and for which only learning-based approaches and heuristics are viable solutions [11]. In some of these cases, the behavior of a physical asset is discrete and can be modeled as a state machine. When this happens, most solutions in the literature opt for Reinforcement Learning (RL) techniques applied over Markov Decision Processes (MDP), so that global KPIs can be autonomously reached by tuning model parameters over time. Examples can be found in [12] and [13], however, these works are solely designed for a single use case and cannot generalize. On the other hand, though, we note that the need for a generalized technique for generating DTs is extremely concrete and seen as a fundamental building block of the path ahead in this field [14]. In a first proposal in such a direction, we modeled the state space of discrete DTs as MDPs and provided a generic “generator” of DTs via the WoT paradigm [15]. We used a combination of frequency-based past observation and deep neural networks to model the physical asset’s behavior even in the presence of unseen states. Similarly, the work in [16] is a recent proposal that aims to move from one-off implementation of DTs to a generative paradigm, where formal models represent mathematically the physical object and learn its behavior according to its state space and what is being observed in the real world by its sensors. In detail, they propose a general framework based on a graphical model that estimates the state variables probabilistically according to partially observable Markov Chains. They use then the UAV context to evaluate the framework in the real world. This work however focuses on the mathematical models and gives little space to data representation, which is instead a core part of the present work. In the same vein, we can find other proposals such as the work in [17], based on Probabilistic Digital Twins (PDTs) that are built through Bayesian probabilistic frameworks. In this case, the key concept is that DTs must rely not only on *observations* coming from sensing experience but also on *assumptions*, as in fact it is never possible to observe every single property that affects the physical asset, therefore, some of the observations must be replaced by mathematical models in order to capture uncertainties and to provide a minimum starting knowledge to the framework. In that work, such uncertainties are modeled as epistemic and aleatory within the proposed framework, which uses deep reinforcement learning for making predictions. Similarly, in this work, DTs are modeled through both assumptions (which define the base models for properties) and observations (which contribute to parameter tuning). However, we detach from state-based discrete properties, rather we take into account continuous properties that can be synthesized via regressive models. This applied to generalized DTs has not been tackled yet in literature, furthermore, we provide an accurate vocabulary for DTs representation via the WoT, introduced in the following section.

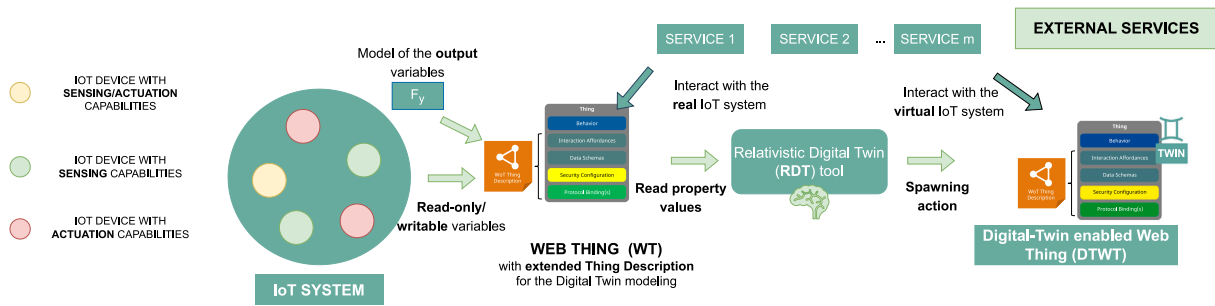


Fig. 1. An IoT/WoT generic scenario and its DTWT concept..

2.2. Web of things

The world of IoT is known to be affected by technology and standard fragmentation, due to the rise of newer and newer technologies, to which current standards often cannot keep up the pace. This causes interoperability problems that need to be tackled at design time, at the same time tolerating the coexistence of heterogeneous systems. The WoT is a paradigm that proposes to extend IoT entities by using standard Web technologies [18], which are solid and coherent throughout history. In this paper we will make use of the W3C WoT standard [19], however, we will only use the term WoT to indicate it. In detail, the WoT standard abstracts a physical asset into a virtual counterpart called Web Thing (WT), which acts as a proxy between the Web world and the physical one. A WT exposes the capabilities and properties of the physical asset just like a Web Server, by providing a set of endpoints that can be queried, furthermore, it can perform actions actively against other WTs. For this reason, the software component embedding the WT is called “servient”. The metadata describing all the interactions that a WT can handle is stored within a Thing Description (TD), a JSON-LD document that acts as the manifest of the WT by expressing its affordances. These are:

- *Properties*, which correspond to state variables that can be fetched and changed at runtime.
- *Actions*, which denote function calls to be performed by the WT, typically triggering a defined behavior.
- *Events*, which are asynchronous and cause the WT to react to their occurrence.

By nature, the WoT paradigm provides an environment that fits very well with the implementation of DTs, which can easily be embedded into existing servients. However, besides our past work that we mentioned [15], these concepts have rarely been paired in literature. An example is [20], where a “Digital Twin WT” is envisioned as a virtual entity derived from the aggregation of multiple physical ones. Differently, in [21], the WoT is seen as the scaffolding to better grasp the intertwining of single components within structurally complex Systems-of-Systems, whilst in [22] DTs are WTs that are automatically generated and managed within a container middleware, however, authors do not specify whether the behavior is also reproduced, as the focus is semantic interactions. DT is an extremely novel concept and very few research efforts have been directed toward automatic generation of DTs in an IoT generic context. Out of them, none implies or proposes a standard like the WoT to describe how to interact with them. A glossary of the terms used alongside the paper is presented in Table 1.

3. System model

We now give a formal presentation of the generic setting for our framework, before moving to a more concrete and illustrative setting. Let us consider a generic IoT scenario depicted in Fig. 1 and composed of $N \geq 1$ devices, $D = \{d_1, \dots, d_N\}$, with sensing and/or actuation capabilities. Each IoT device d_i is associated to a list $S_i = \{s_{i,1}, s_{i,2}, \dots\}$




Fig. 2. The room simulated scenario.

of time-dependent variables determining its state over time. Let $s_{i,j}(t)$ be the value of variable $s_{i,j}$ of device d_i at time t , where the time $t \in \mathbb{R}$ is continuous. More specifically, we assume that for each device d_i , $S_i = R_i \cup W_i$, where:

- $R_i = \{r_{i,1}, r_{i,2}, \dots\}$ indicates the set of *read-only* variables. An instance can be the variable storing the measurement produced by a sensor unit, e.g., the temperature value.
- $W_i = \{w_{i,1}, w_{i,2}, \dots\}$ indicates the set of *writable* variables. An instance can be a tunable variable triggering an action: e.g., in the case of a led actuator d_1 having a single variable $w_{1,1}$ that defines the led luminescence, setting $w_{1,1}(t)$ equal to $\{\text{True} \mid \text{False}\}$ corresponds to the action at time t of turning on or off the light, respectively. Another instance can be represented by a configuration parameter of the device, e.g., the sampling frequency of a sensor unit.

This distinction allows to capture the difference between properties whose value directly depends on external actions, or is the result of some dynamics or relationships within the system of interest. While writable variables are essentially inputs applied to the system, read-only variables can be merely monitored, if observable at all. Note that this categorization is not unique and may arise from, and then reflect, the intended purpose of the DT. Moreover, we consider that all IoT devices in our environment can have causal relationships with each other, i.e., the change of a writable variable for one device may affect the values of the variables of other devices. Given such premises, the overall goal of our study is how to support the automatic deployment of a DT of the whole IoT scenario, considered as a single entity. To this purpose, let $R = \bigcup_i R_i$ be the set of the read-only system variables, $W = \bigcup_i W_i$ be the set of the writable system variables, and $S = R \cup W$. Moreover, let $R(t)$, $W(t)$, and $S(t)$ be the value at time t for all the state variables in R , W , and S , respectively. We call $Z = \{z_1, z_2, \dots\}$, with $Z \subseteq R$, the set of *output* system variables that the DT is in charge of prediction. Our goal is to automatize the estimation of the (partially unreadable) system state based on previous actions and measurements.

It is worth mentioning that this model formalization is independent of a specific use case, and hence it can be easily applied to any IoT

Table 1
Glossary of acronyms used in the paper.

Acronym	Meaning
WoT	W3C Web of Things: it refers to the set of standards for the Web of Things
WT	Web Thing: a single Thing that was mapped with the W3C WoT
DT	Digital Twin
TD	Thing Description: semantically-enriched description of a WT according to the WoT
DTWT	Digital Twin Web Thing: DT of a WT
RDT	Relativistic Digital Twin: the framework proposed in the paper

scenario that involves devices that interact with each other and that can be described through *read-only*, *writable* variables. This can be considered the only requirement for the target IoT System, since these variables must be mapped to affordances and models in a TD (as better described in Section 4) and enhanced with specific domain properties, that can be expressed through a semantic description. This is better depicted in the left part of Figure 1, where the overall methodology is presented: a generic IoT system composed of heterogeneous IoT devices with sensing and/or actuation capabilities is represented through a TD that is composed of some *read-only* and *writable* variables and the set of models for the *output* variables and of the devices, encoded as affordances in the TD. This means that — for example — a *smart vehicle* with its engine sensors, a *smart room* with domotics devices, or a *smart field* with its irrigation actuators share the same methodology of being mapped to a TD; the only difference is represented by the domain knowledge that enriches the TD affordances and that can be manually added to semantically describe the IoT System. To give concreteness to all these concepts, we considered the scenario depicted in Fig. 2, developed in A and then evaluated in Section 7. The IoT system in the exam is a closed smart home environment composed of two rooms, R_A and R_B . The full system is mapped to a single WT, exposing as properties the variables of the sensors/actuators. In our case, each room is hosting a temperature sensor, T_A and T_B , respectively. Each of the temperature sensors has a read-only property that outputs their sensor value. Both rooms host a heater (H_A and H_B in the figure) while the first room also hosts a cooler (C). Both heaters and the cooler have a single writable property that represents their power, which can be set manually to a value from 0 to 9 for the cooler, and on-off for the heaters. Temperature values of both rooms are clearly affected by a number of factors: the power of the heaters, the power of the cooler, the temperature of the other room, and, finally, the temperature of the external environment. In the scenario described so far, we have $R = \{T_A, T_B, H_A, H_B, C\}$ and $W = \{H_A, H_B, C\}$, where $\{H_A, H_B, C\}$ are both readable and writable. The goal of the DTWT of such a complex system is to model the way in which these components affect each other, in order to output temperature values that resemble the real ones (i.e., $Z = \{T_A, T_B\}$).

Generally speaking, DTs are anticipated not only to model the devices' capabilities but also the communication occurring between each device or object within the network. In our case, this translates into the fact that the DT of a single entity should be able to consider and simulate the network communications involved among the IoT objects belonging to the single entity. Nevertheless, we highlight that the DT model created through the framework is strongly bound to the set of *read-only* and *writable* variables considered for each device. This means that the communications involved in the single entity can be modeled in the DT if and only if the communication capabilities of the devices are described through some *read-only* and/or *writable* properties, that is if the DT's intended use requires considering the communication as part of the Thing. In all the other cases, the communications among the IoT objects are considered granted and can be abstracted by the DT of the single entity, since they do not affect the other variables.

Clearly, it is not possible to predict the behavior of any IoT system without taking into account domain-specific information, since — for instance — the DT of a smart home may use ML techniques that require different information than the DT of an industrial appliance. For this

reason, in this paper, we make use of ML with explicit models: each readable variable $r_j \in R$ is associated with a function that describes its behavior. We consider two types of functions in this study. By relating some unknowns and their rate of change, *differential equations* offer a powerful modeling tool to describe the evolution of physical quantities over time. Vice versa, *algebraic expressions* allow to model the instantaneous value of specific quantities. Once, combined, these two types of expressions allow to model a broad spectrum of cyber-physical systems and operations, capturing both physical and computational elements.

We highlight that the choice of using ML with explicit models rather than other popular approaches such as artificial neural networks is quite realistic since in several practical use cases (including the ones evaluated in Section 7, and also industrial environments such as [23]) the system behavior can be described with well-known models derived from the physics and the engineering. As suggested in [24], a DT should impart some form of agency toward a *desired function*, so the underlying model could be intended for a specific application and goal-oriented, capturing only some (cyber)physical aspects. Since all models are wrong, the question is whether the model is *illuminating and useful* for the user [25]. More general models could be considered as well, e.g., by including delayed or stochastic differential equations, or difference equations for discrete-time models. Physics-informed neural networks also give the opportunity to combine data-driven and model-based approaches. All these components can be readily incorporated into our framework, as such models build upon (known) equations and (partially unknown) parameters. Then, one can capture features such as unobservable dynamics, response delays, and jitter in the communication network. A detailed discussion of these models is beyond the scope of this paper.

We split the set of readable variables R into two disjoint sets $A = \{a_1, a_2, \dots\}$ and $B = \{b_1, b_2, \dots\}$, i.e., $R = A \cup B$, $A \cap B = \emptyset$, where A contains the variables $r_j \in R$ that are modeled by an algebraic function and B contains the variables modeled by a differential equation. In case the value of the variable $a_j \in A$ at time t can be modeled by an algebraic expression g_j , then the latter is defined as follows:

$$a_j(t) = g_j(t, B(t), W(t), P_j) \quad (1)$$

where $B(t)$ defines the value at time t for all the state variables in B . Here, t is the time, $B(t)$ and $W(t)$ are the set of variables that a_j could depend on, P_j is the set of coefficients. We define $G = \{g_j \mid a_j \in A\}$ as the set of the algebraic functions.

On the other side, in case the value of variable r_j depends on its evolution over time, a differential equation is defined as follows:

$$\dot{b}_j(t) = f_j(t, B(t), A(t), W(t), P_j) \quad (2)$$

where $A(t)$ defines the value at time t for all the state variables in A . Similar to the previous case, t is the time, $A(t)$, $B(t)$, and $W(t)$ are the set of variables that b_j could depend on, and P_j is the set of coefficients. Let $F = \{f_j \mid b_j \in B\}$ be the set of all the differential equations.

Both the algebraic and the differential equations defined for the variables r_j depend on a set of model parameters P_j . We assume that $|P_j \cap P_{j'}| \geq 0$, $\forall r_j, r_{j'} \in R$, i.e., some model parameters can be shared among different state variables. We indicate with $P_{\text{share}} = \bigcup_{r_j, r_{j'} \in R} (P_j \cap P_{j'})$ the shared set of parameters.

The goal of the DT is to predict the system behavior by forecasting the values of the Z variables at time $t' > t$, where t is the actual time. To do that, for each output variable in Z , the functions F and G are assumed to be known in the pre-deployment phase, except for the exact values of the model parameters P . The RDT framework allows to tune the value of the parameters $P = \bigcup_{r_j \in R} P_j$ in order to properly estimate the values of the Z variables, through the methods defined in Section 6. In the following, we indicate with F^0 and G^0 the *behavioral models* where the parameters are initialized with the values defined by P^0 , and with \hat{F} and \hat{G} the *trained behavioral models*, i.e., the functions \hat{F} and \hat{G} in which all the parameters have been properly instantiated with the values \hat{P} computed during the training phase.

In the context of the rooms scenario depicted in Fig. 2, we partition the set of readable variables, denoted as R , into two subsets: $A = \{H_A, H_B\}$ and $B = \{T_A, T_B, C\}$. Variables in set A require an algebraic expression for representation, whereas those in set B necessitate a differential equation. We emphasize the fact that the modeling decision is in charge of the user that is building the DT. Various configurations of sets A and B are feasible, and the subsequent definitions of algebraic and differential equations can vary accordingly. For instance, the heater H_A which belongs to set A , can be represented as an instantaneous value that characterizes the heater's power. This can be expressed as $H_A(t) = p_{H_A}$, where $p_{H_A} \in P_{H_A}$ represents the trainable parameter for H_A . Conversely, the cooler C , a member of set B , can be described using a differential equation to capture the effect of its increasing power post activation. This can be given by $\dot{C}(t) = p_{C,1} \cdot (p_{C,2} \cdot C_{\text{ref}}(t) - C(t))$, where C_{ref} is the cooler set-point power and $p_{C,1}, p_{C,2} \in P_C$ serve as the trainable parameters for C . A comprehensive definition and further details of the rooms scenario will be elaborated upon in Section 7.1.

Given these premises, the key contributions of our work consist in how to:

- (i) allow the user to describe the characteristics of the IoT systems, such as the variables in R and W , the IoT protocols to access them, and the model behavior G and F , in a domain-independent way;
- (ii) based on such inputs, fully automatize the data acquisition process from the IoT system in order to learn the values of the model parameters P that best fit the observed data, and hence produce the trained behavioral model (\hat{F} and \hat{G});
- (iii) at the end of the learning phase, automatize the spawning of the DT which will consist in a software module enabling the user to interact with the virtual counterpart like with the original system.

Regarding the first contribution, our approach is to extend the WoT standard reviewed in the previous Section, and specifically the TD, with a new vocabulary that allows to model the DT of a generic IoT system and more specifically the functions F and G previously introduced. As shown in Fig. 1, the IoT system — considered as a single entity — is mapped to one WT, exposing as properties the list of read-only and writable variables (S). This is important, as canonical WoT deployments may also choose to map every IoT device to a WT, ending up with several WTs within a single environment. We instead propose to have a single WT where all IoT devices are mapped as affordances, as, otherwise, capturing their interactions would be far more challenging. To this purpose, we map the read-only variables R of the IoT devices to readable WoT Properties (rProp), and the writable variables W to writable WoT Properties (wProp). Notice that the choice of having a single WT to model the whole IoT environment is for readability; our proposal can support also scenarios with multiple interacting WTs, each associated with its own DT.

We remark that the WoT is not a pre-requisite of our solution and that any IoT scenario that is not native WoT-based can be properly extended with WoT interfaces. However, the choice of the WoT is motivated by the possibility of describing — in a uniform and domain-agnostic way — the input/output and the behavior of the DT, thanks to our TD extension described in Section 4. The WT is then passed

as input to the RDT framework, which is in charge of the next two contributions presented in Sections 5 and 6, i.e., the generation of the trained behavioral model and then the spawning of the Digital Twin-enabled Web Thing (DTWT). The latter has the same TD as the original WT, hence external services can interact with the DTWT in the same way as the original WT, by accessing the same list of affordances. However, any property change action generates a consequence only on the virtual environment without affecting the original system, hence enabling the what-if analysis. In addition, the user can modify a time property allowing to monitor the future behavior of the system, hence justifying the parallelism with the relativistic concept. Additionally, the Thing Description of the W3C WoT provides a selection of well-established security mechanisms, along with metadata describing the configuration of these security mechanisms. These mechanisms and their configurations are directly built into the protocols eligible as Protocol Bindings for W3C WoT. This integration of security within the WoT framework ensures that our RDT framework can leverage these established protocols, thereby inheriting a robust security layer.

4. W3C TD extension

In this Section, we discuss how to create the WoT TD of the environment, including its read-only and writable variables and the description of its behavioral model. Notice that the system could have been represented by multiple WTs, one per IoT object, however, that would not allow us to use our framework seamlessly, because capturing dependencies among different WTs brings in additional challenges and needs more research in the future. In this paper, we use a single WT to represent a full IoT ecosystem, thus assuming that two different WTs do not explicitly influence each other, and a single WT can represent multiple IoT objects, with their affordances merged in a single pool.

Regarding the structure of an ordinary rProp in the legacy W3C WoT standard, it includes the terms of the Thing Description (TD) Ontology¹: `type`, `description`, `observable`, `readOnly`, `writeOnly` and `uriVariables`. We present below a brief recall:

- `type` is a JSON-LD keyword to label the object with semantic tags about its type (e.g., “number”).
- `description` is a human-readable piece of text that describes what the property is all about.
- `observable` is a boolean value that indicates whether the property is observable and thus if the servient hosting the WT needs to support the relative protocol binding.
- `readOnly` is a boolean value that indicates whether the property can only be read and not modified by a consumer.
- `writeOnly` is a boolean value that indicates whether the property can only be modified by a consumer and cannot be read.
- `uriVariables` is a data schema that indicates the list of parameters passed onto the request URI.
- `forms` is a list of hypermedia controls that represent how the interaction can be performed; for instance, in the case of an HTTP protocol binding, forms contain the endpoint and the port at which the property is accessible.

The behavioral model cannot be described in the current specifications of the W3C WoT standard. For this reason, in this paper, we propose to extend the existing TD with a new vocabulary allowing the general description of behavioral models for any domain, in order to support the next steps of automatic deployment of the DTWT. Specifically, we added the terms `model`, `modelInput`, and `valueFrom` as shown in Table 2 and explained in the following. Since the vocabulary is not reviewed by the W3C WG yet, we use the example namespace <http://example.org/2022/wot/dtwot>, so the recommended prefix to use in TDs is `dtwt:`.

¹ <https://www.w3.org/2019/wot/td>

Table 2

DTWT terms defined for a Property. Their assignment is optional for TDs, but they are required if the Thing shall be turned into a DTWT.

Term (dtwt:)	Description	Class	Type
model	Mathematical model of the property	PropertyAffordance	string
modelInput	Container for inputs of the model	PropertyAffordance	Class
ValueFrom	Indication from where to take the value of the property, <i>i.e.</i> , from the physical sensor or from the computed model	PropertyAffordance	string

4.1. Model

The field `model` is an arithmetical expression, encoded in a Python-like syntax, that describes the behavior of the property. This is basically represented via a function that is provided by the manufacturer of the appliance. As described in Section 3, this field allows to describe the mathematical function that models the state of a readable variable $r_j \in R$. Indeed, we assume that the manufacturer knows such a function to some extent, obviously, the actual behavior of the appliance is determined by a number of factors that are dependent on the environment and are outside the control of the manufacturer a priori (*i.e.*, the model parameters). In order to accurately represent this, the `model` is defined as a continuous function of a number of parameters that will be subsequently learned. More formally, a model is defined in the following way:

```
{behavior} = {function} | {constraints} | {
  guess}
```

Let us dissect the above syntax:

4.1.1. Behavior

The behavior is one of `self` and `dot(self)`. In the first case, it means that the behavior of the property can be described by an algebraic expression, that is a fixed value that has no dynamics over time and depends on top of some static properties. An example might be the intrinsic power dissipated by a heater, which is dependent solely on its power level, manually set by the user. Here, $r_j \in A$ will be described by a function in G in the form of (1).

In the second case, it means that the behavior of the property depends on time through a differential equation, therefore the function outputs its *integration* over time. This includes properties that may change depending on the time passed since a certain event occurred, even if all the other properties of the DT remain constant. An example can be the temperature of a room after the heater has been turned on. In fact, after such an event, the values do not change instantly, rather it takes a significant transient to reach a new steady state. In this case, r_j belongs to B and the function is in the form of (2).

4.1.2. Function

The function is an arithmetical expression, encoded in python-like syntax. Apart from common numbers, operands, and parentheses, we have added a few keywords that point to definite values in our TD. First of all, the function is defined on top of the set P of parameters that, as said, are tuned by the optimizer to fit the model to the real behavior of the property. These parameters are organized in two arrays:

- `params[n]` is an array of n parameters that are specific to this model (*i.e.*, their scope is limited to this function). Any parameter can be accessed using `params[i]` where $0 \leq i < n$.
- `global[m]` is an array of m parameters that are global and defined only once (*i.e.*, their scope is extended to the whole thing description). Any parameter can be accessed using `global[i]` where $0 \leq i < m$.

The distinction between `global` and `local` parameter is needed when this latter is shared among multiple state variables. For this reason, the `global[m]` array will contain the whole set P_{share} . Furthermore, the function can use other properties of the same DTWT, which are called “model inputs” and are listed in the `modelInput` field of

the `rProp`. An example may be our temperature property that takes as input the power of the heater, which is another `rProp` of the DTWT. The detailed interaction between a model input and the model that imports it is described in the next subsection. Finally, the function can use the keyword `self`, which represents the non-integrated value of the property, and the keyword `value()`, which represents the actual value read from the real sensor, when available.

4.1.3. Constraints

The `constraints` is a set of preconditions regarding the `params[n]` and the `global[m]` of the model. They are read by the optimizer and respected when performing the parameter optimization. We assume that each parameter $p_j \in P$ is bounded by a minimum and a maximum value, *i.e.*, $p_j^{\min} \leq p_j \leq p_j^{\max}$. The set of all the parameter constraints forms the hyperbox \mathcal{P} . For instance a valid syntax can be `params[0] >= 0.0, params[1] <= 0.0, global[0] <= 1.0`. The global parameters’ constraints are to be set only once throughout the TD (clashing constraints will lead to errors).

4.1.4. Guess

The `guess` is a set of initial value assignment regarding the `params[n]` and the `global[m]` of the model, defining the set P^0 . They are the default values before the optimization takes place. For instance, a valid syntax can be `params[0] = 2.5, params[1] = 0.0, global[0] = 0.01`. The global parameters’ `guess`, as for the constraints, is to be set only once throughout the TD (clashing assignments will lead to errors).

4.2. Model inputs

The field `modelInputs` is a list of other `rProps` of the WT that are used within the model. Each of these has five fields:

- `title`, which is a friendly name,
- `propertyName`, which is the exact name of the `rProp` of the WT that this input is referring to,
- `type`, which is the JSON-LD type,
- `model`, which is again a model (using the same syntax as described above) that defines how this input affects the model of the property, and
- `modelType`, which is an optional field denoting a class assigned to this input that is only valid within the scope of the property and is used to aggregate all inputs belonging to the same `modelType` (if there is more than one).

The property `model` can use the `modelInputs` in two ways. The first is by using the construct “`input({title})`”; this imports the exact model of the `modelInput` identified by “`title`”. The model input may have its own `params[]`: in this case, `params[i]` of the property model is not equal to `params[i]` of the model input, as they have different scoping. The model input may refer to the value outputted by the actual property identified by `propertyName` though the keyword `self`. Alternatively, the property `model` can import a group of model inputs through the construct `inputType({modelType})`; in this case, the property model imports all the model inputs with the type `modelType` and aggregates them. The aggregation takes place through an aggregation function which must enclose the `inputType()` call (for instance `sum(inputType(modelType))`).

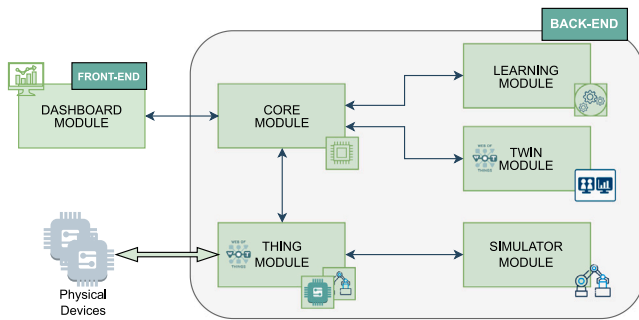


Fig. 3. The RDT framework architecture.

4.3. ValueFrom

The field `valueFrom` describes whether there is a real property that can be read (e.g., a real sensor), in such case the value of this field will be set to “readProperty”. Alternatively, the field can be set to “model”, which instead triggers the simulation of the property value (depending on the behavior, via algebraic evaluation or integration over time) as described by the `model` field of the `rProp`. This indicates where the data comes from when querying directly the `rProp` via its URI, however, when the latter gets imported into a model, it is its model that always gets imported.

In the Appendix, we show the JSON code of the TD of the smart home environment, restricting our attention to the new fields which are used to describe the behavioral model of the IoT system.

5. Framework architecture and implementation

Based on the TD extension described in the previous Section, we designed the RDT framework to support the automatic generation of a DTWT. The framework allows to monitor the properties of the WT, to store their values, to fit the behavioral model till creating the trained model, and, finally, to generate the DTWT. The RDT framework architecture is micro-services oriented, with each macro-functionality mapped to a specific module. As shown in Fig. 3, the architecture is composed of six main modules: a single one for the *front-end* functionalities, while the other five run on the *back-end*. More in detail, the *Dashboard module* provides an intuitive Graphical User Interface (GUI) through which users can easily interact with the framework.

We detail a typical flow of operations in Fig. 4. Three main interaction modes are supported: simulating the IoT system operations and producing simulated sensory data, uploading traces containing the measurements, or attaching a physical device to the platform via the API of the RDT framework. Using any of the previous options, the user provides the training data for the parameters’ identification and then configures the training hyperparameters, such as the tolerance for termination of the training algorithm, or specifies the algorithm itself. Once the model has been trained, a new DTWT sharing the same TD as the original WT can be spawned; the DTWT is completely detached from the physical IoT device and hence can be used for what-if analysis (i.e., the user can observe how the twin would react to future events). In addition, the framework supports re-sync operations when the properties of the DTWT can be re-initialized with the current values of the original WT. The back-end functionalities are implemented within (i) the *core module*, (ii) the *learning module*, (iii) the *twin module*, (iv) the *Thing module*, and (v) the *simulator module*. The *core module* enables the communication between the *dashboard* and the rest of the services through a dedicated API, by translating the user’s requests into proper commands to be invoked on the services. Additionally, the *core module* monitors the *Things*, collecting all the measurements and data needed for the *training phase* from the *learning module*. The latter, as

better detailed in Section 6, allows to parse the TD and to instantiate the behavioral model described by our TD extension. Furthermore, it uses the data provided by the *core module* to derive the optimal value of the model parameters. Within the RDT framework, the data storage is primarily dependent on the specific application and it is not directly related to the framework itself. It depends on the amount of data collected for training the behavioral models. Different applications may generate varying volumes of data, which in turn influences the storage requirements. Our framework is designed to be adaptable to these varying needs, ensuring that the data storage solutions can be tailored according to the specific demands of each application. The *twin module* takes in input the trained model and generates the DTWT, which forecasts the values of the properties through the trained behavioral model. The *Thing module* contains all the mechanisms to interact with the IoT counterpart. This can be a physical device — in this case, a dedicated channel is created for the data exchange — or can be simulated through the *simulator module*. The latter is just a wrapper for domain-specific tools that allow reproducing the behavior of a device, by generating the values of its properties over time and allowing to receive actions in input. As a result, any custom simulator can be connected to the RDT framework by implementing a proper connector. Clearly, the process of training, learning, spawning, and simulating may be repeated over time, as the behavior of the physical devices may change or there may be aspects that had never manifested during the previous training iteration. In such case, the newly generated DTWTs will replace the old ones. The full process designed within the RDT framework primarily introduces overhead during the automatic generation and deployment of digital twins. However, it is crucial to emphasize that instantiating new digital twins within our framework requires constant resources. This efficiency in spawning new digital twins ensures that the framework can scale effectively, accommodating the creation and management of multiple digital twins without significant increases in computational load or complexity. Such scalability is a key feature of our framework, enabling it to adapt and respond to varying demands in diverse IoT environments. The most time-consuming aspect of the RDT framework is the training phase. The real-time efficiency in the re-spawning phase is influenced by the specific problem and the complexity of the training model. While our framework is designed for efficient re-training process, the actual speed and responsiveness in real-time scenarios depend on the training model’s intricacies and the data nature. Therefore, the RDT framework’s effectiveness in real-time operations varies with each application’s specific requirements and complexity.

5.1. Implementation

The front-end functionalities are included in the *Dashboard module* which consists of a Web application written using the Angular² framework, as shown in Fig. 5; the Apache Echarts³ library is used to plot the results of the what-if analysis, i.e., showing the real measurements and the predictions made by the DT. The back-end functionalities are implemented within the other five modules of the architecture: each of them is implemented in Typescript and runs over NodeJS as runtime system. The *Thing* and the *twin* modules have been implemented as WTs using the Eclipse node-wot tool⁴. More in detail, each WT runs a *Servient*, as indicated in the official guidelines of the W3C WoT Scripting API⁵. The *learning module* is a Web server that collects requests from the core module through a REST API. The optimization tasks are implemented as Python scripts using the SciPy library, as better detailed in Section 6. The *simulator module* is written in Typescript.

² <https://angular.io/>

³ <https://echarts.apache.org/en/index.html>

⁴ <https://github.com/eclipse/thingweb.node-wot>

⁵ <https://www.w3.org/TR/wot-scripting-api/>

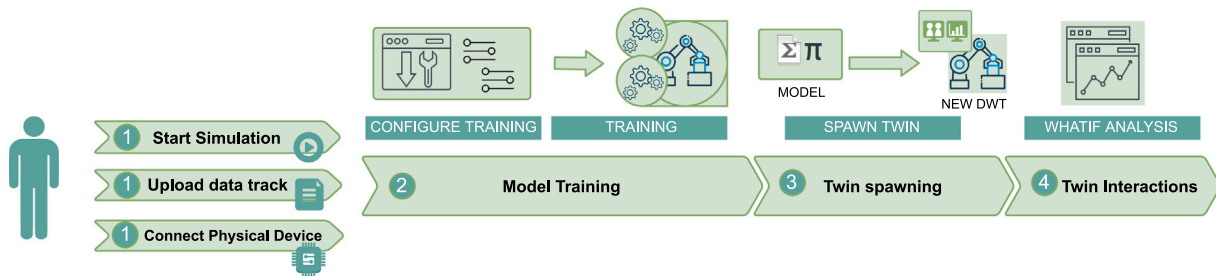


Fig. 4. The sequence of operations of the RDT framework. The user can interact with it in three ways: connecting it to a simulator, uploading tracks with property values, or connecting it to a physical device.

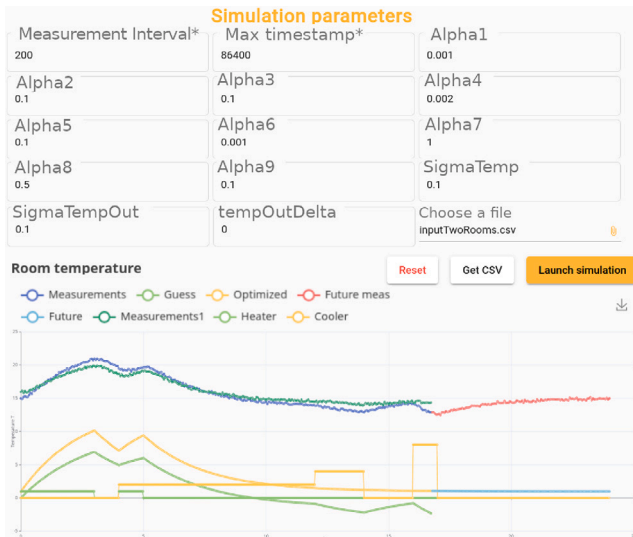


Fig. 5. A screenshot of the dashboard module. Using such GUI, the user can simulate, train, and visualize the guess of the model of the room scenario described in Section 4.

6. DT learning

The *learning module* is a crucial component of the RDT framework. It is in charge of parsing the TD of the WT, extracting its behavioral model, and generating the trained behavioral model from the data. The operations of the module can be split into three major tasks: (i) model import, (ii) model inference, and (iii) model learning. In the following, a detailed description of each task is provided.

6.1. Model import

A *parser* module has been implemented in order to translate the arithmetical expressions that are encoded in a Python-like syntax within the TD of the WT, *i.e.*, the algebraic functions G and the differential equations F , into an executable Python code. Listing 1 shows an example of the translation of the temperature property for the smart home example (Listing 2).

```

y[0]=readProperty("heater",timestamp,data)
dxdt[0]=params[1]*(params[2]*(params[3]-x[0])
+((params[0]*y[0])+(-x[2]))+params[4]*(x[1]
-x[0]))
dxdt[1]=params[6]*(params[7]*(params[3]-x[1])
+((params[5]*y[0]))+params[4]*(x[0]-x[1]))
dxdt[2]=params[8]*(params[9]*max(0,min(round(
readProperty("cooler",timestamp,data)),9))-
x[2])

```

Listing 1: DTWT TD, parsed room model

In Listing 1 we can notice the differences between the algebraic variables (the y) and the differential properties (the x and the $dxdt$). In this specific example, the variables $x[0]$ and $x[1]$ correspond to the temperature values of R_A and R_B at time `timestamp`, while $y[0]$ and $x[2]$ correspond to the heating and cooling power, respectively. The rationale of the equations is discussed in Section 7.1.

6.2. Model inference

The *inference* task returns the history of states in a specific time interval, given the initial state, the model parameters and the control actions; *i.e.*, it is used as a mathematical simulation engine to build past or future states. More in detail, given t in a time interval $\Gamma_T := [t_0, t_T]$ from an initial time t_0 to an end time t_T , the actions executed on the writable properties performed during this time interval W_T , some model parameters \hat{P} , and an initial state $\hat{R}(t_0)$, the models defined by the sets F and G (in the form of the Eqs. (1)–(2)) allow the prediction of the states $R(t)$ over Γ_T . To this aim, we introduce a parametric and non-autonomous ordinary differential equation (ODE) that is defined for each $b_j \in B$ from the executed actions and the model parameters:

$$\begin{aligned} \dot{b}_j(t) &= F_{j,W_T,\hat{P}}(t, B(t)) \\ &= f_j(t, B(t), A(t), W_T(t), \hat{P}), \quad t \in \Gamma_T, \\ r_j(t_0) &= \hat{r}_j(t_0). \end{aligned} \quad (3)$$

Solving the Cauchy — or initial value — problem results in a prediction of the DTWT behavior according to the parametrized model. In this work, we employed ODE solvers based on Runge–Kutta schemes [26], [27, Ch. 5], which provide a numerical approximation $\hat{b}_j(t)$ of a solution to (3). Then, one can readily compute the approximation of $\hat{a}_k \in A$ of the algebraic states using (1) pointwise. For our purposes⁶, the ODE solver provides $B_{W_T,\hat{P},\hat{R}(t_0)}(t)$ and $A_{W_T,\hat{P},\hat{R}(t_0)}(t)$, namely it returns the trajectories of any $b_j \in B$ and $a_k \in A$ on a time interval, given a history of actions W_T , a set of parameters \hat{P} and an initial state $\hat{R}(t_0)$.

6.3. Model learning

The inference task assumes the knowledge of the model parameters P (called \hat{P} in Section 6.2). However, these values are not known a priori and need to be estimated. For this reason, this task is in charge of estimating P with \hat{P} starting from the observed behavior of the WT. Let $\{\Omega_1, \dots, \Omega_O\}$ be the set of O observations of the WT at the time instants $\{t_1, \dots, t_O\} \subset \mathbb{R}$. Let the function $h(\cdot)$ be an observation model, such that

$$\Omega_k \approx h(t_k, B(t_k), A(t_k), W(t_k), \hat{P})$$

for $k = 1, \dots, O$. Effectively, given some past actions, we seek the optimal choice of parameters \hat{P} that minimizes the deviation between

⁶ In our specific implementation, we use the function `solve_ivp` from the SciPy library `scipy.integrate` [28]. A modern ecosystem for solving differential equations is described in [29].

observations and retrospective predictions based on the parametric model.

Given a set of actions W_Q executed in $\Gamma_Q := [t_0, t_Q]$, with $t_Q \geq t_0$, and defining the predictor

$$H_{W_Q, \hat{P}, \hat{R}(t_0)}(t) := h(t, B_{W_Q, \hat{P}, \hat{R}(t_0)}(t), A_{W_Q, \hat{P}, \hat{R}(t_0)}(t), W_Q(t), \hat{P}),$$

we consider the following optimization problem to fit the timestamped observations (t_k, Ω_k) , $k = 1, \dots, O$:

$$\underset{\hat{P} \in \mathcal{P}, R(t_0) \in \mathcal{R}_0}{\text{minimize}} \sum_{k=1}^O \|H_{W_Q, \hat{P}, \hat{R}(t_0)}(t_k) - \Omega_k\|^2. \quad (4)$$

Here, \mathcal{P} is the hyperbox containing the lower and upper bounds of each parameter, and \mathcal{R}_0 is the set of the possible initial states. In the smart room example presented so far, \mathcal{R}_0 is composed of the constraints: $-20 \leq T_A(t_0), T_B(t_0) \leq 50$, $0 \leq H_A(t_0), H_B(t_0) \leq 1$, and $0 \leq C(t_0) \leq 9$.

The resulting mathematical task (4) is a nonlinear least squares problem with bound constraints [30, Ch. 4].⁷ Although in general nonconvex and possibly with many local minimizers, the complexity of algorithms for tackling problem (4) is addressed in [31], discussing how many function evaluations might be required to obtain an approximate solution.

Starting from some guesses P^{guess} and $R^{\text{guess}}(t_0)$, and seeking some optimal values P^* and $R^*(t_0)$, the numerical solver needs to compute $H_{W_Q, \hat{P}, \hat{R}(t_0)}(t_k)$ for some tentative values $\hat{P} \in \mathcal{P}$ and $\hat{R}_0(t_0) \in \mathcal{R}_0$. In turn, such evaluation requires to solve an initial value problem as in (3), obtaining $B_{W_Q, \hat{P}, \hat{R}(t_0)}(\cdot)$ and $A_{W_Q, \hat{P}, \hat{R}(t_0)}(\cdot)$. Then, these predicted trajectories are sampled at the time instants t_k , $k = 1, \dots, O$, where the measurement model h is evaluated, yielding $H_{W_Q, \hat{P}, \hat{R}(t_0)}(t_k)$.

Given the estimates, \hat{P} and $\hat{R}(t_0)$ of model parameters and (differential) state at time t_0 , the *learning module* is able to predict the DTWT state at any time for some given history of actions on the writable variables W_Q , effectively simulating the system dynamics and behavior according to the identified model. In this case, the W_Q function is needed to simulate the action executed in future times t_k , i.e., for $t_0 < t_k \leq t_Q$.

7. Validation

In this Section, we present a comprehensive analysis and validation of the proposed RDT framework. We assess the performance on two different case studies: (i) a simulation test where we present the results of the smart home scenario introduced firstly in Section 4, and (ii) a real test-bed consisting of a drone quadcopter which can be employed e.g., for video-surveillance applications. In the following two subsections, we describe the setup of the experiments and show the evaluation results.

7.1. Simulated testbed: smart home

In Section 4 we introduced the smart home environment as a showcase of the WT TD extension. In this Section, we detail the mathematical model characterizing the system behavior and we evaluate the performance of the learning process.

The proposed scenario is composed of two closed rooms: room R_A and room R_B . At time t the temperature of R_A is denoted by $T_A(t)$, the temperature of R_B by $T_B(t)$, and the outdoor temperature by $T_{\text{out}}(t)$. Inside each room (differently from our initial example) we have a heater that we assume to be connected to the same heating system, i.e., they are simultaneously either *ON* or *OFF*. The heaters' status at time t is defined by the binary variable $b_H(t)$:

$$b_H(t) = \begin{cases} 1 & \text{if the heating system is ON} \\ 0 & \text{otherwise} \end{cases}$$

⁷ In our specific implementation, we use the function `least_squares` from the SciPy library `scipy.optimize` [28].

Table 3

Parameters used in the simulated room environment.

β_1	0.001 K/J	β_4	0.002 K/J	β_7	1.0 W
β_2	0.1 W/K	β_5	0.1 W/K	β_8	0.5 W
β_3	0.1 W/K	β_6	0.001 s ⁻¹	β_9	0.1 W

The heaters' power at time t are defined by $H_A(t)$ and $H_B(t)$ for the heater in R_A and R_B , respectively. In R_A we suppose to have also a cooler whose power at time t is denoted by $C(t)$; this depends on a user-defined setpoint $C_{\text{ref}}(t)$ which has 10 degrees of active powers, namely $0 \leq C_{\text{ref}}(t) \leq 9$, $C_{\text{ref}}(t) \in \mathbb{N}$ for all times t . Based on our previous notation, we have the set of readable variables $R = \{T_A, T_B, H_A, H_B, C, C_{\text{ref}}, b_H\}$ and writable variables $W = \{C_{\text{ref}}, b_H\}$.

The following differential equations describe the system dynamics over time t ; they will be used in the TD of the WT:

$$\begin{aligned} \dot{T}_A(t) &= \alpha_1^R [\alpha_2^R (\alpha_6^R - T_A(t)) + H_A(t) - C(t) + \\ &\quad + \alpha_3^R (T_B(t) - T_A(t))] \end{aligned} \quad (5a)$$

$$\begin{aligned} \dot{T}_B(t) &= \alpha_4^R [\alpha_5^R (\alpha_6^R - T_B(t)) + H_B(t) + \\ &\quad + \alpha_3^R (T_A(t) - T_B(t))] \end{aligned} \quad (5b)$$

$$H_A(t) = \alpha_7^R b_H(t) \quad (5c)$$

$$H_B(t) = \alpha_8^R b_H(t) \quad (5d)$$

$$\dot{C}(t) = \alpha_9^R [\alpha_{10}^R C_{\text{ref}}(t) - C(t)] \quad (5e)$$

Herein, the linear model for the temperature dynamics is provided by Newton's law of heat transfer, whereas considering thermal inertia in the cooling system yields the first-order linear model in (5e). The parameters set P is defined by $P = \{\alpha_1^R, \dots, \alpha_{10}^R\}$, where α_*^R are free variables and their values will be estimated by the learning module of the RDT framework, based on observations from the simulated/real environment.

Given such a model, we generated the TD for the rooms as previously introduced in Section 4. In this description, we note that some parameters are shared among different properties. For example, α_3^R is connected to both R_A and R_B . The parameter α_6^R is shared between the two heaters but it coincides with the outdoor temperature, assumed to be constant over time, i.e., $T_{\text{out}}(t) = T_{\text{out}} = \alpha_6^R, \forall t$. Finally, we defined a `global` variable for the outdoor temperature, called `global` [2] in Listing 2.

We set up a simulation environment to generate the temperature data acquired by the virtual sensors positioned inside the two rooms. To avoid a straightforward learning system where the proposed DTWT model mimics exactly the simulated environment, we defined a different model to simulate the rooms' temperature. Considering a thermal inertia, and thus first-order dynamics, in the heating system and not in the cooling system, the simulation model reads as follows:

$$\begin{aligned} \dot{T}_A(t) &= \beta_1 [\beta_2 (T_{\text{out}}(t) - T_A(t)) + H_A(t) - C(t) + \\ &\quad + \beta_3 (T_B(t) - T_A(t))] \end{aligned} \quad (6a)$$

$$\begin{aligned} \dot{T}_B(t) &= \beta_4 [\beta_5 (T_{\text{out}}(t) - T_B(t)) + H_B(t) + \\ &\quad + \beta_3 (T_A(t) - T_B(t))] \end{aligned} \quad (6b)$$

$$\dot{H}_A(t) = \beta_6 [\beta_7 b_H(t) - H_A(t)] \quad (6c)$$

$$\dot{H}_B(t) = \beta_6 [\beta_8 b_H(t) - H_B(t)] \quad (6d)$$

$$C(t) = \beta_9 C_{\text{ref}}(t) \quad (6e)$$

Parameters β_* in (6) are used to customize the rooms' simulations and are unknown to the user. Where not specified otherwise, their default values are indicated in Table 3.

The observed parameter values at time t_k are represented by the set $\Omega_k = \{T_A^{\text{obs}}(t_k), T_B^{\text{obs}}(t_k), H_A^{\text{obs}}(t_k), H_B^{\text{obs}}(t_k), C^{\text{obs}}(t_k), b_H^{\text{obs}}(t_k)\}$. To make

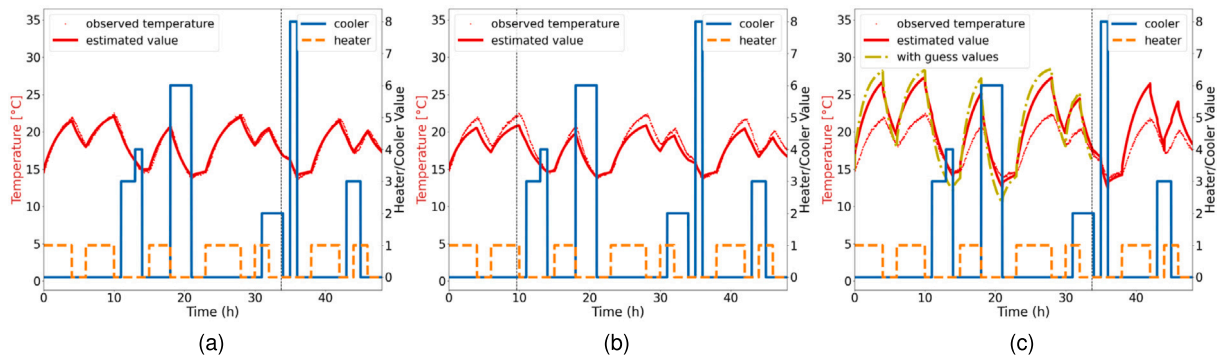


Fig. 6. The evaluation of the accuracy of $T_A(t)$ estimated by the spawned DTWTs. Fig. 6(a) shows the observed and estimated values when the DTWT is spawned at time 34h; Fig. 6(b) shows the same curves when the DTWT is spawned at time 10h; Fig. 6(c) shows the same results when the DTWT is spawned at time 34h but without any initial calibration of the parameter guesses p^{guess} .

the simulation more realistic, we added a Gaussian noise to the observed temperature during the learning process, *i.e.*,

$$T_A^{\text{obs}}(t_k) = T_A(t_k) + \varepsilon \tag{7}$$

$$T_B^{\text{obs}}(t_k) = T_B(t_k) + \varepsilon \tag{8}$$

where $\varepsilon \sim \mathcal{N}(\mu, \sigma)$ simulates the error introduced by the devices. Finally, also in the simulation part, we keep the T_{out} constant with an additive Gaussian noise, *i.e.*, $T_{\text{out}} = 15^\circ\text{Celsius} + \mathcal{N}(\mu, \sigma)$. In the following tests, we used $\mu = 0\text{ K}$ and $\sigma = 0.1\text{ K}$.

In the experiments, we evaluated the ability of the trained behavioral model to forecast indoor temperature. We generated two DTWTs for the two rooms, and we set up the framework in order to learn the value of T_A . As described in Fig. 4, we uploaded the tracks containing the temperature measurements from time $t = 0$ to the time of the DTWT generation t_{spawn} . These tracks are generated using the rooms simulator described by Eqs. (6a)–(8) generating the observations Ω_k , for $0 < k < t_{\text{spawn}}$, using a predefined sequence of actions $W_{t_{\text{spawn}}}$ for the writable variables (shown in blue and orange lines in Figs. 6(a)–6(c)). Figs. 6(a)–6(c) show the observed and forecast values of T_A as produced by the DTWT, for different t_{spawn} values, *i.e.*, for different training lengths. The observed values are depicted with a dotted red line, and the forecast values with a solid red line; similarly, we highlight the event of training end at time t_{spawn} with a vertical dashed black line. We set $t_{\text{spawn}} = 34\text{ h}$ for Figs. 6(a) and 6(c), and $t_{\text{spawn}} = 10\text{ h}$ for Fig. 6(b). In all the Figures, we show the operational state of the heater and cooler over time, *i.e.*, their power values, tuned as a consequence of the user’s actions. For the experiment in Fig. 6(c), the parameters p^{guess} were initialized with random values.

Fig. 6(a) shows that the proposed model is able to predict the room temperature in a very accurate way. Indeed, after the training interval (equal to 34 h), the line of the estimated values overlaps with the line of the observed values. In Fig. 6(b) we reduced the training length compared to the previous analysis ($t_{\text{spawn}} = 10\text{ h}$). As expected, the amount of training data impacts the forecast error, which is higher than Fig. 6(a). However, the curve of estimated values is still quite close to the curve of the observed values, thanks to the effective initialization of the model parameters p^{guess} . Indeed, through the *guess* field explained in Section 4, the user can provide the initial values of the model parameters within the TD of the WT, based on prior contextual knowledge of the use case; such assignments are then automatically loaded by the RDT framework. Vice versa, Fig. 6(c) shows the impact of parameter calibration and training on the performance of the DT. More specifically, we plot with a yellow dashed line the estimated temperature values when no learning has been performed at all: in addition, the model parameters are initialized with random p^{guess} values, and not changed during the training phase. We can notice that the expected values diverge significantly from the observed ones. In the same Figure, the red solid line depicts the case where the training

Table 4

Mean squared error for experiments plotted in Fig. 6. The initial guess for each experiment can be either designed by the modelers or randomly assigned.

% Training	Initial guess	MSE
20 (10 h)	Designed	0.822
70 (34 h)	Designed	0.16
0 (0 h)	Random	17.649
70 (34 h)	Random	6.376

is executed, but without any initial calibration of the p^{guess} parameters. Comparing Figs. 6(a)–6(c), we can notice that the calibration has a positive impact on the forecast accuracy; however, even without it, the module is able to learn from the incoming data and to improve its performance over time.

Finally, in Table 4 we summarize the performance for all the configurations considered in Figs. 6(a)–6(c), by reporting the mean squared error (MSE) during the test phase only (*i.e.*, $t > t_{\text{spawn}}$). The third row (0% training, Random initial guess) refers to the baseline where no parameter initial calibration and no learning phase have been performed. Vice versa, it is easy to notice that both the training time and the tuning of the initial guesses are of paramount importance to minimize the MSE metric (row 2).

7.2. Real testbed: quadcopter

Differently from the previous experiment, in this Section, we evaluate the proposed framework in a real testbed scenario. The case study is the DT of a quadcopter (in our case, a *DJI Mini 2 drone*⁸), flying in an outdoor environment. A quadcopter is a small-scale helicopter with four rotors directed upwards and placed in a square formation from the quadcopter’s center of mass. The quadcopter movements can be controlled by setting the angular velocities of the rotors which are spun by electric motors. In the literature, several models have been proposed in order to characterize the quadcopter’s kinematics, at different granularity levels. Interested readers can refer to [32] for a detailed review. In this work, we used a second-order linear dynamics model in the body reference frame [32]. Specifically, let q be the target quadcopter. The latter is defined by its position in global reference frame with the triple $(q_x, q_y, q_z) \in \mathbb{R}^3$ and by its attitude with $(q_\phi, q_\theta, q_\psi) \in \mathbb{R}^3$. Here, $(q_\phi, q_\theta, q_\psi)$ are the three Euler angles modeling the *yaw* ($-\pi \leq q_\psi \leq \pi$), the *pitch* ($-\pi/2 \leq q_\theta \leq \pi/2$), and the *roll* ($-\pi/2 \leq q_\phi \leq \pi/2$), which represent the orientation of the quadcopter. In the experiments, latitude and longitude are remapped to local q_x, q_y coordinates. The velocity of q in global reference frame is defined by $q_v = (q_{vx}, q_{vy}, q_{vz}) \in \mathbb{R}^3$ and

⁸ <https://www.dji.com/mini-2>




Fig. 7. Coordinate system for the quadcopter. The red elements represent the joystick inputs.

in the body frame reference by $q_v^B = (q_{vx}^B, q_{vy}^B, q_{vz}^B) \in \mathbb{R}^3$ (see Fig. 7 for a visual description of the variables). Here, we assume a small angle approximation that yields to $q_\phi = q_\theta = 0$. The angular velocity for the yaw is defined by $q_{\psi\omega}$. Finally, the following rotation matrix $RT(q_\psi)$ is used to convert the copter velocity from the body frame reference to the inertial frame one:

$$\begin{bmatrix} q_{vx} \\ q_{vy} \end{bmatrix} = RT(q_\psi) \cdot \begin{bmatrix} q_{vx}^B \\ q_{vy}^B \end{bmatrix} = \begin{bmatrix} \cos q_\psi & -\sin q_\psi \\ \sin q_\psi & \cos q_\psi \end{bmatrix} \cdot \begin{bmatrix} q_{vx}^B \\ q_{vy}^B \end{bmatrix}$$

The model described so far defines the actual state of the quadcopter q , i.e., its position, altitude, and velocity. The quadcopter is controlled by a remote controller whose modeled inputs are:

- *throttle* Th , that commands climb/descent movement of the copter. Positive values of Th are for climbing request, descent otherwise.
- *rudder* Ru , that commands the yaw movements. Positive values of Ru are for the clockwise rotation of the copter.
- *elevator* El , that commands the pitch, i.e., the forward and backward movements. Here $El > 0$ means forward, backward otherwise.
- *aileron* Ai , that commands the roll, i.e., the left and right movements, where $Ai > 0$ means left movements, right direction otherwise.

The aforementioned user inputs constitute the set of writable variables W of the device. The full quadcopter model description and its translation in TD is described in Appendix B. Using the resulting TD, we deployed the WT of the DJI quadcopter and connected it to the proposed RDT framework. During the flight, we collected the drone's data thanks to the DJI APIs that allow downloading the raw data of the onboard sensors. More specifically, we measured the following quantities: *latitude*, *longitude*, and *altitude* for the drone position (that are remapped to local (x, y, z) coordinates); *compass* for the drone's attitude (where *compass* defines the *yaw*), and *elevator*, *aileron*, *throttle*, *rudder* as the joystick user's inputs. The experiments aim to validate the capability of the RDT framework related to the following aspects: (i) the overall ability of the proposed TD extension to model complex IoT systems like the proposed one, (ii) the accuracy of the learning phase and the impact of different parameters such as the training length; (iii) the effectiveness of the DT to support what-if analysis and hence to avoid harmful operations, thanks to the capability of simulating the system reactions to user's commands.

After the take-off of the quadcopter, the RDT framework starts acquiring raw data from the drone and activates the *model training* phase. The take-off and the landing procedures are considered special movements hence they are not taken into account during the training phase. We tested two DTWT spawning events at a distance of 10 s one from the other in order to evaluate the impact of the training length. In Figs. 8(a), 8(b), 8(c) we plot the real (q_x, q_y, q_z) values (got from the onboard sensors) with a dashed red line. In the same Figures, we plot

the estimated position when the DTWTs are spawned at $t_{\text{spawn}} = 5$ s (blue solid line) and at $t_{\text{spawn}} = 15$ s (green solid line). It is easy to notice that the three curves almost overlap during the training phase, i.e., when $t < t_{\text{spawn}}$. Vice versa, the curves of observed values and of predicted values with $t_{\text{spawn}} = 5$ s diverge significantly during the test phase when the DT is used as a standalone component: this is a clear symptom that the training phase was too short. Indeed, with $t_{\text{spawn}} = 15$ s, the DT is able to forecast the trajectory of the quadcopter in a quite accurate way for all three axes.

In the next experiment, we investigate the possibility to deploy the DT in parallel with the actual IoT system and to use it for real-time what-if analysis. More in detail, in Fig. 9 we envision the usage of the RDT framework by a drone pilot to safely control his/her device. We consider a scenario with geofence areas, i.e., flight zones from which the drone is not allowed to leave. After the definition of such areas, the user generates a DTWT of the drone through our RDT framework, and he/she tests any pilot command on the DTWT before executing them on the real drone. In Fig. 9, the drone is in position (q_{x0}, q_{y0}) ; the user provides a sequence of throttle/rudder/elevator/aileron commands, with the relative time intervals, to the DTWT. The latter returns an *alert* warning since, based on the positions estimated by the trained behavioral model, the drone would be located in position (q_{x1}, q_{y1}) outside the geofence area. As a result, the user inputs another sequence of commands which are validated by the DTWT (returning position (q_{x2}, q_{y2})). Finally, he/she issues them to the real drone which moves to the position (q_{x3}, q_{y3}) , remaining inside the geofence area.

Let d_{thr} be the radius of a circular geofence area centered on the actual position of the drone. Given a set of control commands, the DT must estimate the future position of the drone, and, in addition, detect the events in which the drone will move outside the geofence area. Let t_{la} be the look-ahead time that specifies the time offset in the future when the drone position must be predicted. As the evaluation metric for this experiment, we used the *precision* index. The latter is the ratio of correctly predicted positive observations (true positive) to the total predicted positive observations (true positive plus false positive), where high precision relates to a low-false positive rate. In this case, such a ratio is critical because it expresses the times in which the system failed to detect crossing events of the geofence borders. More formally the index is calculated as

$$\text{precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}$$

where TP (true positive) represents the true positive values, i.e., when the DTWT correctly predicts that the drone will remain inside the geofence area after t_{la} seconds. The value of FP (false positive) indicates the number of cases when the DTWT predicts that the drone will stay inside the area while the real position is outside.

In Figs. 10(a) and 10(b), we plot the *precision* index when varying the look-ahead time t_{la} and the geofence area d_{thr} , respectively. Fig. 10(a) shows the index for two different training time, $t_{\text{spawn}} = 5$ s and $t_{\text{spawn}} = 15$ s, where $d_{\text{thr}} = 5$ m. It is easy to notice that, like for the previous experiments, the training length impacts significantly the estimation of future values. Moreover, the higher the t_{la} look-ahead time, the lower the prediction index. This is due to the accumulated errors as the DTWT moves to future time intervals. However, when the DT is employed with $t_{\text{spawn}} = 15$ s, the number of harmful events — in our case geofence exit — can be considerably reduced compared to a case when the DT is not employed at all. The same trends can be observed in Fig. 10(b) where $t_{\text{la}} = 5$ m. The radius of the geofence area impacts the precision value in a way that false positive values are less frequent when the target area is larger.

The last experiment shows the possibility for the RDT framework to continuously improve the accuracy of the DTWT model over time. As described in the previous Sections, the behavioral model for the DT is based on a set of free variables α_*^D that must be properly tuned by the RDT framework. In the previous experiments, each training stage was initialized with the P^0 that are read from the WT TD. On the




Fig. 8. The predicted values of x , y , and z from the spawned DTWTs are shown in Figs. 8(a), 8(b), and 8(c), respectively. In the Figures, the red dots show the collected data from the real drone, the blue lines are the predicted values from the DTWT spawned after 5 s (dashed vertical line), and the green lines are the predicted values from the DTWT spawned after 15 s (solid vertical line).

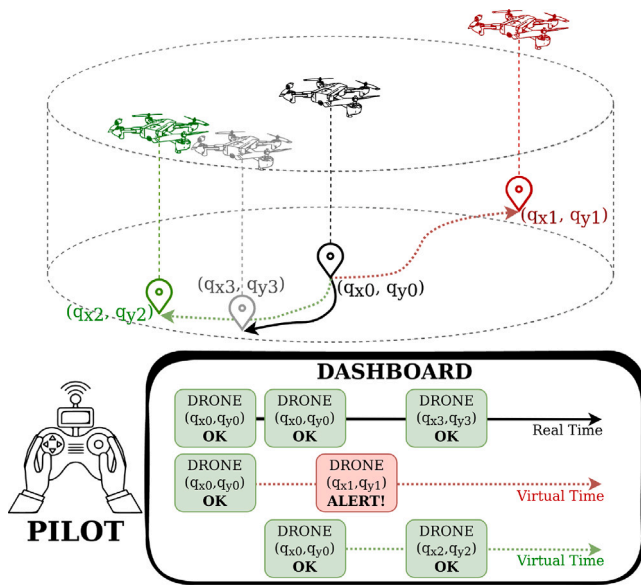


Fig. 9. A use-case scenario for the RDT framework. The pilot checks the control commands on the DTWT before issuing them to the drone.

contrary, in this experiment, we used the same estimated parameters set during consecutive training phases. More formally, let $\hat{P}^1, \hat{P}^2, \dots$ be the trained parameters set after 1, 2, ... training rounds that used P_1^0, P_2^0, \dots as initial parameters guess, respectively. To exploit previous knowledge, in each training round we used as initial guess the trained parameters of the previous round, i.e., $P_i^0 = \hat{P}^{i-1}, \forall i > 1$ with $P_1^0 = P^0$. The rationale of this experiment is to evaluate the capability of the RDT framework to continuously train its DT model during its lifetime. In Fig. 10(c) we show the distance error between the real drone and the predicted position ($d_{r,e}$) after consecutive training rounds. Here, the training time is set to 10 s. We can notice that the error decreases significantly over the training rounds, hence justifying the effectiveness of continuous learning mechanisms. We plan to further elaborate on this feature of the RDT framework as future works.

8. Conclusions and future works

In this paper we presented the Relativistic Digital Twin (RDT) framework, which is able to replicate a physical asset, being it a single system or a System of Systems (SoS) into an abstract Digital Twin Web Thing (DTWT) that emulates its behavior and its reaction to events. The resulting framework has the capability of (i) generating heterogeneous and generic Digital Twins that are not necessarily tied to the use case on which they operate, and (ii) interoperating with well-known IoT

frameworks that use established standards. To this aim, we targeted an initial modeling of a real-world system based on the Web of Things (WoT) standard promoted by the W3C. As a research contribution, we proposed to extend the Thing Description (TD) component of the WoT standard in order to take into account the system behavior as a set of algebraic expressions or differential equations. We presented the new TD vocabulary and all the components of the RDT framework allowing to generate a trained behavioral model by fitting the model parameters to the observed data. The RDT framework was evaluated in two different scenarios, namely the DT of a simulated smart home environment with heterogeneous sensors/actuators and, and the DT of a real-world quadcopter predicting its kinematics over time. In both cases, we assessed the effectiveness of the parameter training process and the capability of the RDT framework to support predictive and what-if-analysis. The work presented in this paper is pioneering and we believe it can trigger the interest of the scientific community in pursuing further extensions of the RDT framework. One of them could be making the generation of these Digital Twins federated so that a SoS of Digital Twins can be composed by building a mashup of pre-generated Digital Twins of appliances that have not yet worked together in the real world. In addition, we plan to include support for additional ML techniques (also model-free) and extend the evaluation to other case studies.

CRedit authorship contribution statement

Luca Sciuillo: Conceptualization, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Alberto De Marchi:** Conceptualization, Methodology, Software, Writing – original draft, Writing – review & editing. **Angelo Trotta:** Formal analysis, Investigation, Validation, Visualization. **Federico Montori:** Methodology, Writing – original draft, Writing – review & editing. **Luciano Bononi:** Supervision. **Marco Di Felice:** Funding acquisition, Project administration, Supervision, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be shared along the source code in an open source repository.

Acknowledgment

This research is funded by the “Progetto Casa delle Tecnologie Emergenti” - Comune di Bologna - PSC MISE 2014-2020.

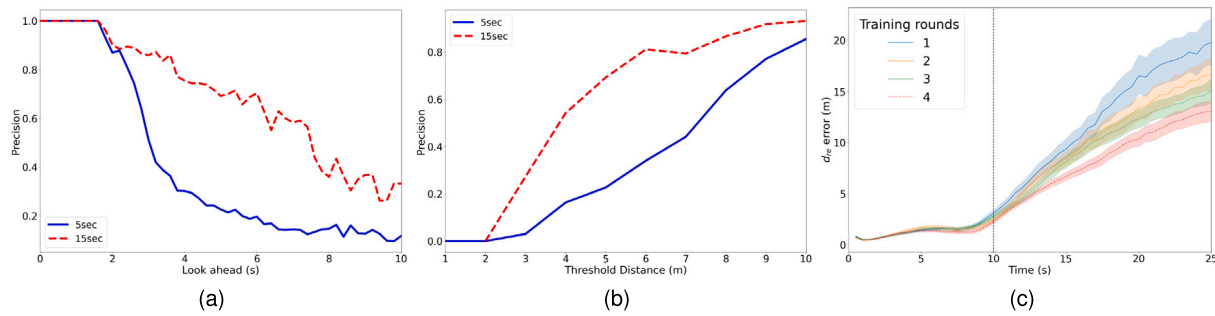


Fig. 10. The precision index is depicted in Figs. 10(a) and 10(b) by varying the look-ahead time (t_{la}) and the distance threshold (d_{thr}), respectively. Fig. 10(c) shows the distance error d_{e} over time for consecutive training rounds.

Appendix A. Code snippets — Rooms scenario

In this Appendix we present three code snippets that exemplify the temperature-based use case⁹ presented in Section 3. The code below shows the three newly introduced JSON fields for the `rProp` that represents the temperature in one of the rooms (T_A). As the property is connected with a sensor value, we set “valueFrom”: “readProperty”. The `model` field is specified on top of three local parameters and global parameters 2 and 3. The model also depends on three `modelInputs`: `heaterPower`, `coolerPower` and `temp1`. `temp1` is the temperature of the second room, which is imported into the model through `input(temp1)`. Since the model of `temp1` only specifies “self”, then we can assume that the optimizer will import the model of the property `temperature1` as it is. Furthermore, both `coolerPower` and `heaterPower` are summed up and imported together through `sum(inputType(@heatPower))`. This happens because they are responsible (one opposed to the other) for the same phenomenon, i.e., a power causing a temperature variation over time. Both the models of “heater” and “cooler” are not imported as they are: the first one is multiplied by a parameter (`params[0]`, which is different from the `params[0]` in the model of the temperature) and the second one is multiplied by -1 , because its power contributes negatively to the temperature rise.

```
...
"valueFrom": "readProperty",
"modelInput": [
{
  "title": "heaterPower",
  "propertyName": "heater",
  "type": "number",
  "model": "params[0] * self | params[0] >= 0.0
    | params[0] = 1.0",
  "modelType": "@heatPower"
},
{
  "title": "coolerPower",
  "propertyName": "cooler",
  "type": "number",
  "model": "-self",
  "modelType": "@heatPower"
},
{
  "title": "temp1",
  "propertyName": "temperature1",
  "type": "number",
  "model": "self"
}
],
```

⁹ The full TD can be found at <https://github.com/UniBO-PRISMLab/relativistic-wot/blob/main/optimizer/thing.jsonld>

```
"model": "dot(self) = params[0] * (params[1] *
(global[2] - self) + sum(inputType(
@heatPower))) + global[3] * (input(temp1) -
self) | params[0] >= 0.0, params[1] >= 0.0
, global[3] >= 0.0 | params[0] = 0.001,
params[1] = 0.1, global[2] = 15.0, global[3]
] = 0.1",
...
```

Listing 2: DTWT TD, temperature

The reason behind how we are treating differently `heaterPower` and `coolerPower` lies on top of how they are defined. Below we can see how `heaterPower` is defined: its model only outputs the actual value of the `rProp`. This identifies a set of cases where the manufacturer only provides the behavior of the component independently from the environment in which it acts. In fact, in this case, the impact of the `heaterPower` has to be tuned via a parameter within the `modelInputs` in Listing 2.

```
...
"valueFrom": "readProperty",
"model": "self = value()",
...
```

Listing 3: DTWT TD, heaterPower

Differently, the `coolerPower`, as it is defined in Listing 4, indicates a `rProp` that encloses pretty much all the information about the cooler. In fact, the `value()` of the cooler is never used, instead the `rProp` always outputs the value calculated by a model that defines how the cooler would impact an environment on top of two model parameters. Furthermore, the model imports the model of the `rProp` called “`coolerSetPoint`” as a `modelInput`. The latter is another `rProp` that only outputs the power lever of the cooler knob as a number constrained between 0 and 9.

```
...
"valueFrom": "model",
"modelInput": [
{
  "title": "coolerSetpoint",
  "type": "number",
  "model": "max(0, min(round(value()), 9))"
}
],
"model": "dot(self) = params[0] * (params[1] *
input(coolerSetpoint) - self) | params[0]
>= 0.0, params[1] >= 0.0 | params[1] = 0.1,
params[0] = 0.1",
...
```

Listing 4: DTWT TD, coolerPower

Appendix B. Quadcopter model

In this Appendix, we describe in detail the quadcopter model and its relative translation in the TD.

Following the drone's model description of Section 7.2 with states and inputs, we define the system dynamics as follows:

$$\dot{q}_x = q_{vx} \quad (\text{B.1})$$

$$\dot{q}_y = q_{vy} \quad (\text{B.2})$$

$$\dot{q}_z = q_{vz} \quad (\text{B.3})$$

$$\dot{q}_\psi = q_{v\psi} \quad (\text{B.4})$$

$$\dot{q}_{vx}^B = \alpha_1^D (\alpha_2^D El - q_{vx}^B) \quad (\text{B.5})$$

$$\dot{q}_{vy}^B = \alpha_3^D (\alpha_4^D Ai - q_{vy}^B) \quad (\text{B.6})$$

$$\dot{q}_{vz}^B = \alpha_5^D (\alpha_6^D Th - q_{vz}^B) \quad (\text{B.7})$$

$$\dot{q}_{v\psi}^B = \alpha_7^D (\alpha_8^D Ru - q_{v\psi}^B) \quad (\text{B.8})$$

complemented by the kinematic relationships:

$$\begin{bmatrix} q_{vx}^B \\ q_{vy}^B \end{bmatrix} = \text{RT}(q_\psi)^\top \begin{bmatrix} q_{vx} \\ q_{vy} \end{bmatrix} \quad (\text{B.9})$$

$$\begin{bmatrix} \dot{q}_{vx} \\ \dot{q}_{vy} \end{bmatrix} = \text{RT}(q_\psi) \begin{bmatrix} \dot{q}_{vx}^B \\ \dot{q}_{vy}^B \end{bmatrix} + \text{RT}'(q_\psi) q_{v\psi} \begin{bmatrix} q_{vx}^B \\ q_{vy}^B \end{bmatrix}. \quad (\text{B.10})$$

Notice that the dynamics are expressed in the drone's body reference frame and that the first-order models for the velocities stem from Newton's basic laws of classical mechanics and the inertial properties of the copter, namely from its mass distribution. Furthermore, although uncoupled to the body reference frame, the dynamics equations involve different axes of the inertial reference frame. The model above contains 8 free variables, $\alpha_1^D, \dots, \alpha_8^D$, which represent the set of parameters P specific to the quadcopter in use, and which must be properly tuned. Given such a proposed behavioral model, we defined the TD of the WT associated with the quadcopter; the latter will be used by the RDT framework to generate the DT of the appliance. Due to the length of the resulting TD, we report here only the most significant fragments¹⁰.

```
"positionX": {
  ...
  "valueFrom": "readProperty",
  "modelInput": [
    {
      "title": "vX",
      "propertyName": "velocityX",
      "type": "number",
      "model": "self"
    }
  ],
  "model": "dot(self) = input(vX)"
},
```

Listing 5: TD of the quadcopter, x-position property

The snippet in Listing 5 shows the x -position property of the drone, *i.e.*, q_x . Similar to q_y and q_z , the position of the drone is read directly from the onboard sensors and hence the `valueFrom` field is set to `readProperty`. Following (B.1) describing the drone dynamics for the x -position, we included the `model` field that reflects the position's dynamic using the `dot(self)` notation. We can also notice from Listing 5 the usage of the `modelInput` field that lists the properties to be used in the model.

¹⁰ The full TD can be found at <https://github.com/UniBO-PRISMLab/relativistic-wot/blob/main/optimizer/thingDrone.jsonld>

```
"velocityX": {
  ...
  "valueFrom": "model",
  "modelInput": [
    {
      "title": "yaw",
      "propertyName": "yaw",
      "type": "number",
      "model": "self"
    },
    {
      "title": "yawrate",
      "propertyName": "yawrate",
      "type": "number",
      "model": "self"
    },
    {
      "title": "vbX",
      "propertyName": "velocitybodyX",
      "type": "number",
      "model": "self"
    },
    {
      "title": "vbY",
      "propertyName": "velocitybodyY",
      "type": "number",
      "model": "self"
    },
    {
      "title": "abX",
      "propertyName": "accelerationbodyX",
      "type": "number",
      "model": "self"
    },
    {
      "title": "abY",
      "propertyName": "accelerationbodyY",
      "type": "number",
      "model": "self"
    }
  ],
  "model": "dot(self) = math.cos(input(yaw)) *
    input(abX) - math.sin(input(yaw)) *
    input(abY) - input(yawrate) * (math.sin(
    input(yaw)) * input(vbX) + math.cos(
    input(yaw)) * input(vbY))"
},
```

Listing 6: DTWT TD, x-velocity property

Another fragment of the TD is reported in Listing 6, focusing on the modeling of the quadcopter's velocity. More specifically, we show the `velocityX` property that corresponds to the q_{vx} state of the quadcopter and whose dynamics are described by (B.10). In this case, the `valueFrom` field is set to `model` which means that the property has no real sensor in the drone and hence the value is derived from the model. The `model` field encodes (B.10) using the syntax described in Section 4. Finally, the `modelInput` field lists the other properties used inside the model.

References

- [1] AIAA. Digital Engineering Integration Committee, Digital Twin: Definition & Value. An AIAA and AIA Position Paper, Tech. Rep., American Institute of Aeronautics and Astronautics (AIAA) and Aerospace Industries Association (AIA), 2020, URL <https://www.aiaa.org/advocacy/Policy-Papers/Institute-Position-Papers>.
- [2] A. Rasheed, O. San, T. Kvamsdal, Digital twin: Values, challenges and enablers from a modeling perspective, IEEE Access 8 (2020) 21980–22012, <http://dx.doi.org/10.1109/ACCESS.2020.2970143>.

- [3] W. Kritzing, M. Karner, G. Traar, J. Henjes, W. Sihn, Digital twin in manufacturing: A categorical literature review and classification, *IFAC-PapersOnLine* 51 (11) (2018) 1016–1022, <http://dx.doi.org/10.1016/j.ifacol.2018.08.474>, 16th IFAC Symposium on Information Control Problems in Manufacturing INCOM 2018.
- [4] S.A. Niederer, M.S. Sacks, M. Girolami, K.E. Willcox, Scaling digital twins from the artisanal to the industrial, *Nat. Comput. Sci.* 1 (5) (2021) 313–320, <http://dx.doi.org/10.1038/s43588-021-00072-5>.
- [5] A. Gyrard, S.K. Datta, C. Bonnet, A survey and analysis of ontology-based software tools for semantic interoperability in IoT and WoT landscapes, in: 2018 IEEE 4th World Forum on Internet of Things, WF-IoT, 2018, pp. 86–91, <http://dx.doi.org/10.1109/WF-IoT.2018.8355091>.
- [6] W3C Working Group, WoT reference architecture (proposed recommendation 9 April 2020), 2020, URL <http://www.w3.org/TR/wot-architecture/>.
- [7] T. Jäckle, W3C WoT (Web of Things) Integration, 2022, URL <https://www.eclipse.org/ditto/2022-03-03-wot-integration.html>.
- [8] P. Varga, F. Blomstedt, L.L. Ferreira, J. Eliasson, M. Johansson, J. Delsing, I. Martínez de Soria, Making system of systems interoperable – the core components of the arrowhead framework, *J. Netw. Comput. Appl.* 81 (2017) 85–95, <http://dx.doi.org/10.1016/j.jnca.2016.08.028>.
- [9] I. Zyrianoff, L. Gigli, F. Montori, C. Kamienski, M. Di Felice, Two-way integration of service-oriented systems-of-systems with the web of things, in: IECON 2021–47th Annual Conference of the IEEE Industrial Electronics Society, 2021, pp. 1–6, <http://dx.doi.org/10.1109/IECON48115.2021.9589619>.
- [10] P. Pesic, Einstein and the twin paradox, *Eur. J. Phys.* 24 (6) (2003) 585, <http://dx.doi.org/10.1088/0143-0807/24/6/004>.
- [11] Z. Huang, Y. Shen, J. Li, M. Fey, C. Brecher, A survey on AI-driven digital twins in Industry 4.0: Smart manufacturing and advanced robotics, *Sensors* 21 (19) (2021) 6340, <http://dx.doi.org/10.3390/s21196340>.
- [12] T. Wang, J. Cheng, Y. Yang, C. Esposito, H. Snoussi, F. Tao, Adaptive optimization method in digital twin conveyor systems via range-inspection control, *IEEE Trans. Autom. Sci. Eng.* 19 (2) (2022) 1296–1304, <http://dx.doi.org/10.1109/TASE.2020.3043393>.
- [13] F. Jaensch, A. Csiszar, A. Kienzlen, A. Verl, Reinforcement learning of material flow control logic using hardware-in-the-loop simulation, in: 2018 First International Conference on Artificial Intelligence for Industries, AI4I, 2018, pp. 77–80, <http://dx.doi.org/10.1109/AI4I.2018.8665712>.
- [14] R. Minerva, G.M. Lee, N. Crespi, Digital twin in the IoT context: a survey on technical features, scenarios, and architectural models, *Proc. IEEE* 108 (10) (2020) 1785–1824, <http://dx.doi.org/10.1109/JPROC.2020.2998530>.
- [15] L. Sciuillo, A. Trotta, F. Montori, L. Bononi, M. Di Felice, WoTwins: Automatic digital twin generator for the web of things, in: 2022 IEEE 23rd International Symposium on a World of Wireless, Mobile and Multimedia Networks, WoW-MoM, 2022, pp. 607–612, <http://dx.doi.org/10.1109/WoWMM54355.2022.00095>.
- [16] M.G. Kapteyn, J.V.R. Pretorius, K.E. Willcox, A probabilistic graphical model foundation for enabling predictive digital twins at scale, *Nat. Comput. Sci.* 1 (5) (2021) 337–347, <http://dx.doi.org/10.1038/s43588-021-00069-0>.
- [17] C. Agrell, K. Rognlien Dahl, A. Hafver, Optimal sequential decision making with probabilistic digital twins, *SN Appl. Sci.* 5 (4) (2023) 114, <http://dx.doi.org/10.1007/s42452-023-05316-9>.
- [18] L. Sciuillo, L. Gigli, F. Montori, A. Trotta, M. Di Felice, A survey on the web of things, *IEEE Access* 10 (2022) 47570–47596, <http://dx.doi.org/10.1109/ACCESS.2022.3171575>.
- [19] M. Kovatsch, R. Matsukura, M. Lagally, T. Kawaguchi, K. Toumura, K. Kajimoto, Web of Things (WoT) Architecture, W3C Recommendation, W3C, 2020, <https://www.w3.org/TR/wot-architecture/>.
- [20] C. Aguzzi, L. Gigli, L. Sciuillo, A. Trotta, F. Zonzini, L. De Marchi, M. Di Felice, A. Marzani, T.S. Cinotti, MODRON: A scalable and interoperable web of things platform for structural health monitoring, in: 2021 IEEE 18th Annual Consumer Communications & Networking Conference, CCNC, 2021, pp. 1–7, <http://dx.doi.org/10.1109/CCNC49032.2021.9369492>.
- [21] G. Privat, T. Coupaye, S. Bolle, P. Raipin-Parvedy, WoT graph as multiscale digital-twin for cyber-physical systems-of-systems, in: Proc., 2nd W3C Web of Things Workshop. Grenoble, France: Gilles Private Orange Labs Services, 2019, pp. 1–5, URL <https://www.w3.org/WoT/ws-2019/>.
- [22] S. Muralidharan, B. Yoo, H. Ko, Designing a semantic digital twin model for IoT, in: 2020 IEEE International Conference on Consumer Electronics, ICCE, 2020, pp. 1–2, <http://dx.doi.org/10.1109/ICCE46568.2020.9043088>.
- [23] H. Park, A. Easwaran, S. Andalām, TiLA: Twin-in-the-loop architecture for cyber-physical production systems, in: 2019 IEEE 37th International Conference on Computer Design, ICCD, 2019, pp. 82–90, <http://dx.doi.org/10.1109/ICCD46524.2019.00019>.
- [24] A. Agrawal, V. Singh, M. Fischer, A new perspective on digital twins: Imparting intelligence and agency to entities, *IEEE J. Radio Freq. Identif.* 6 (2022) 871–875, <http://dx.doi.org/10.1109/JRFID.2022.3225741>.
- [25] G.E.P. Box, Science and statistics, *J. Amer. Statist. Assoc.* 71 (356) (1976) 791–799, <http://dx.doi.org/10.1080/01621459.1976.10480949>.
- [26] J.R. Dormand, P.J. Prince, A family of embedded Runge-Kutta formulae, *J. Comput. Appl. Math.* 6 (1) (1980) 19–26, [http://dx.doi.org/10.1016/0771-050X\(80\)90013-3](http://dx.doi.org/10.1016/0771-050X(80)90013-3).
- [27] M. Gerdt, Optimal Control of ODEs and DAES, De Gruyter, 2011, <http://dx.doi.org/10.1515/9783110249996>.
- [28] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S.J. van der Walt, M. Brett, J. Wilson, K.J. Millman, N. Mayorov, A.R.J. Nelson, E. Jones, R. Kern, E. Larson, C.J. Carey, Í. Polat, Y. Feng, E.W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E.A. Quintero, C.R. Harris, A.M. Archibald, A.H. Ribeiro, F. Pedregosa, P. van Mulbregt, SciPy 1.0 Contributors, SciPy 1.0: Fundamental algorithms for scientific computing in python, *Nature Methods* 17 (2020) 261–272, <http://dx.doi.org/10.1038/s41592-019-0686-2>.
- [29] C. Rackauckas, Q. Nie, DifferentialEquations.jl – a performant and feature-rich ecosystem for solving differential equations in Julia, *J. Open Res. Softw.* (2017) <http://dx.doi.org/10.5334/jors.151>.
- [30] J. Nocedal, S.J. Wright, *Numerical Optimization*, second ed., Springer, New York, NY, USA, 2006.
- [31] C. Cartis, N.I.M. Gould, P.L. Toint, Evaluation Complexity of Algorithms for Nonconvex Optimization: Theory, Computation and Perspectives, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2022, <http://dx.doi.org/10.1137/1.9781611976991>.
- [32] X. Zhang, X. Li, K. Wang, Y. Lu, A survey of modelling and identification of quadrotor robot, *Abstr. Appl. Anal.* 2014 (2014) 320526, <http://dx.doi.org/10.1155/2014/320526>.



Luca Sciuillo is a Junior Assistant Professor with the University of Bologna, Italy, where he also obtained his Ph.D. in Computer Science and Engineering in the 2021 and his Master Degree (summa cum Laude) in Computer Science in 2017. He was a Visiting Researcher at the Huawei European Research Center of Munich, Germany. He is part of the IoT PRISM Lab directed by Prof. Marco Di Felice and Prof. Luciano Bononi. His research interests include wireless systems and protocols for emergency scenarios, wireless sensor networks, IoT systems and Web of Things.



Alberto De Marchi is currently a Research Associate at the Institute of Applied Mathematics and Scientific Computing, University of the Bundeswehr Munich, Germany, where he received the Dr.rer.nat. degree (summa cum laude) in 2021. He holds a Master's degree in Mechatronics Engineering from the University of Trento, Italy. He has been a visiting researcher at Curtin University, Australia, in 2022. His research revolves around computational methods for mathematical optimization, with applications ranging from control systems to data science.



Angelo Trotta is a Junior Assistant Professor with the Department of Computer Science and Engineering, University of Bologna. He received his Ph.D. degree in computer science and engineering in 2017 from University of Bologna, Bologna, Italy. He was a Visiting Researcher with the Heudiasyc Laboratory, Sorbonne Universities, UTC, Compiègne, France, and with Genesys-Laboratory, Northeastern University, Boston, MA, USA. He is co-founder of 'AI for People', an international association whose aim is to use the Artificial Intelligent technology for the social good. His current research includes nature inspired algorithms for self-organizing multirobots wireless systems, reinforcement learning, 5G slicing and IoT.



Federico Montori is an Assistant Professor with the University of Bologna, Italy, where he also obtained his Ph.D. in Computer Science and Engineering in 2019. He received the B.S. and the M.S. in Computer Science, both Summa Cum Laude, from the University of Bologna in 2012 and 2015 respectively. He has been a visiting researcher at Swinburne University of Technology, Australia, and Luleå Tekniska Universitet, Sweden. He participated in several EU projects and he is now WP leader for the H2020 Project Arrowhead Tools. His primary research interests include Mobile Crowdsensing (MCS), pervasive and mobile computing, IoT automation and data analysis for IoT scenarios.



Luciano Bononi (MSC, Summa cum laude, 1997, Ph.D, 2001), Full Professor at the Department of Computer Science and Engineering of the University of Bologna. His research fields include: Wireless Systems and Networks, Internet of Things, Smart Mobility and Mobile Applications, Simulation and Digital Twin. He is founder of the Wireless and Mobile Applications Lab in 2010, and he is actively involved in national and international research projects, and collaborations with international Universities and Industry partners. He is a member of the doctoral board for the Ph.D. program in Computer Science and Engineering of the University of Bologna. In 2012–2015 he was elected member of the academic senate of the University of Bologna. Author of more than 160 peer-reviewed publications, he received best paper awards, best editor awards, and certificates of appreciation by Elsevier, IEEE and ACM. He was Associate Editor of more than 5 international journals (10+ special issues), and he worked as chair and TPC member in more than 15 and 200

IEEE/ACM international conferences, respectively. In 2021 he is listed in the “World’s top 2% of scientists of their main subfield discipline”.



Marco Di Felice is a Full Professor of Computer Science with the University of Bologna, where he is the co-director of the IoT PRISM laboratory. He received the Laurea (summa cum laude) and Ph.D. degrees in computer science from the University of Bologna, in 2004 and 2008, respectively. He was a Visiting Researcher with the Georgia Institute of Technology, Atlanta, GA, USA and with Northeastern University, Boston, MA, USA. He authored more than 120 papers on wireless and mobile systems, achieving three best paper awards for his scientific production. He is Associate Editor of the IEEE Internet of Things journal. His research interests include self-organizing wireless networks, unmanned aerial systems, IoT, WoT, and context-aware computing.