

The Degree of Entanglement: Cyber-Physical Awareness in Digital Twin Applications

Marco Picone*, Stefano Mariani*, Roberto Cavicchioli*, Paolo Burgio*, Arslane Hamza Cherif†

* *University of Modena and Reggio Emilia, Italy*; † *Vedecom, France*

{name.surname}@unimore.it, arslane.hamza-cherif@vedecom.fr

Abstract—A defining feature of a Digital Twin (DT) is its level of “entanglement”: the degree of strength to which the twin is interconnected with its physical counterpart. Despite its importance, this characteristic has not been yet fully investigated, and its impact on applications’ design is underestimated. In this paper, we define the concept of “Degree of Entanglement” (DoE), which provides an operational model for assessing the strength of the entanglement between a DT and its physical counterpart. We also propose an interoperable representation of DoE within the Web of Things (WoT) framework, which enables DT-driven applications to dynamically adapt to changes in the physical environment. We evaluate our proposal using two realistic use cases, demonstrating the practical utility of DoE in supporting, for instance, context-awareness decisions and adaptiveness.

Index Terms—Digital Twin, Entanglement, Degree of Entanglement, Cyber-Physical Systems

I. INTRODUCTION

Heterogeneity and fragmentation of the physical world in terms of devices, networks, and protocols, make the engineering of intelligent Cyber-Physical Systems (CPSs) difficult and incentivize the creation of isolated applications operating within their local scope with limited long-term evolution capabilities. The recent growth of CPSs has been accelerated by the concept of Digital Twin (DT) [1] as an effective and powerful paradigm to build digitalized and augmented versions of Physical Assets (PAs) across multiple application domains (as schematically depicted in Figure 1).

However, there is still confusion around one of DTs’ foundational aspects: *entanglement*. It has been recently introduced and informally defined in [2] as the characterization of the communication relationship between DT and PA and recognized as a fundamental property that strongly defines the DT. While entanglement and its measurement play a crucial role in cyber-physical intelligent systems [3]–[5], its comprehensive characterization, along with considerations regarding interoperability and structured definition, is worth further investigation and evaluation.

Also, DTs are still mostly used as *passive* data repositories to shield applications from PAs heterogeneity and provide some level of fault tolerance. They are rarely considered as *active* entities able to *adapt* to the contingencies and opportunities offered by the CPS where the PAs are located [6], [7]. In this context, the possibility to monitor the level of entanglement between DT and PA over time is crucial to fully support applications and services’ requirements.

In this paper, we first introduce the *Degree of Entanglement* as a measure of the “entanglement strength”, then we exploit such notion to enable *adaptiveness* of the DT, showcasing its practical utility, and finally, evaluate our proposal in the domain of Digital Factories.

II. DIGITAL TWINS & ENTANGLEMENT

The DT concept represents the linkage between a physical object (or, asset) and its digital counterpart, meaning that all the information that fully describes the object in the softwarized space must be synchronized in real (or very close to) time according to the requirements of the target application scenario. This relationship has been recently defined in the literature as *entanglement* [2], [8] and still represents a new and under-explored topic and research opportunity ranging from simpler application scenarios with lower binding constraints to real-time deployments with challenging and mandatory requirements. Such a core property of a DT is detailed taking into account three characteristics (schematically illustrated in Figure 2), upon which we define the Degree of Entanglement (DoE) (in Section III):

a) Connectivity: There should be a way for the DT and its associated PA to exchange information, in either direction (see “Association” below). The PA should be able to communicate its status changes to the DT, or the DT should be able to perceive them somehow (e.g. a CCTV camera monitoring a robot in the production line). The same is needed in the opposite direction: changes in the DT should be communicated/perceived by the PA (e.g. commands for action issued by the application layer). Such a connection can be direct, when the DT and the PA are “physically” connected (there included wireless connections), indirect, when they are connected via a mediator (such as a monitoring camera), or absent when they are not connected at all. For direct and indirect connections, the “strength” of such connectivity may be defined based on arbitrary network-related parameters, such as speed, bandwidth, etc.

b) Promptness: The DT and the PA must synchronize their respective states so that users and applications cannot notice any misalignment, except for intended ones (e.g. in the case of a DT-based simulation of PA behavior, that should not be reflected in the PA). In other words, the time needed to update any of the two should be negligible w.r.t. the typical access time of users and applications. A real-time connection between the DT and its associated PA is not necessary, in




Fig. 1. The digital abstraction provided by DTs to tackle (by decoupling) heterogeneity between deployed assets and external applications.

general: it depends on the requirements put forth by the application at hand.

c) Association: The flow of information between a DT and its PA may be uni-directional, that is from PA to DT (the most common case), such as for monitoring applications, or from DT to PA (quite uncommon yet), such as for teleoperation of equipment; or bi-directional, where the DT and the PA may exchange data (PA \rightarrow DT and DT \rightarrow PA) and commands (usually, DT \rightarrow PA only).

Entanglement is then introduced in the state-of-the-art literature [2], in terms of a combination of these three characteristics, as *strong* when the PA is constantly linked to the DT with a bidirectional link; *simple* when the communication is unidirectional, or not real-time, or the linkage may be interrupted for a certain time; and *weak* when information about the PA is derived from data gathered by third parties.

These “levels of entanglement” mix the nature of the relationship between DT and PA with its direction: The strong notion does not mention communication or synchronization quality (or real-time constraints) but does mention the bi-directional nature of the entanglement, whereas the simple option mentions both. In fact, real-time communications and data freshness (promptness) are orthogonal to the direction of entanglement (association), and there are potentially many more factors contributing to determining what we call the *Degree of Entanglement* (DoE) between a DT and its PA. In this challenging context, the convergence and interplay of entanglement with real-time physical assets or even real-time DTs pose a novel vantage point that has remained unexplored within the scientific discourse. This uncharted territory emerges as a potential progression subsequent to this study and the introduction of the DoE.

Accordingly, the next section describes our operational notion of DoE, as well as how it could be practically used to enable the adaptation of DT behavior and to extend the awareness of digital applications on the cyber-physical relationship between the twin and its physical counterpart.

III. THE DEGREE OF ENTANGLEMENT

Only recently the concept of DT entanglement has been emphasized as a fundamental property to characterize the communication relationship between DT and PA. However, there is no consensus yet on a specific (formal, operational) definition of what it means and how it can be computed and used.

We postulate that the entanglement between a DT and its associated PA is not static, but can *dynamically* vary during the DT (and PA) lifespan. At each point in time, hence, the DT has a certain *Degree of Entanglement* (DoE) with its PA. Such a DoE should *measure how tight* – whatever this means – is the coupling between the DT and its PA. For instance, in terms of how promptly the DT reacts to PA changes, or how reliable is the DT representation of the PA state, or similar nuances. In addition, the DT itself should be able to compute its own DoE, so as to be *aware* of it at any time. Only in this way, in fact, the DT would be able to *autonomously adapt* to changes in the *contextual conditions* within which it operates (along with its associated PA).

In the following, we define the DoE in terms of connectivity, promptness, and association following the state of the art literature and already discussed in Section II.

A. Generic formulation

Let $C, P, A \in [0, 1] \subset \mathbb{R}$ be three real numbers representing, respectively, the properties of *connectivity*, *promptness*, and *association* of a DT with respect to its PA. How to compute such numbers and based on what data is out of the scope of this paper, as it is hard or even impossible to determine in general, being likely to depend on many factors at multiple levels: PA nature (e.g. connected object or passive resource), infrastructural constraints (e.g. available network speed and bandwidth), and even domain-specific requirements (e.g. real-time or data freshness requirements). However, the next section provides two possible instantiations of the conceptual framework here formalized, and Section V further specifies the formulae in a specific use case in the domain of digital factories, as a way to exemplify how to use our framework in practice.




Fig. 2. The main aspects characterizing the concept of DT's entanglement.

In the following instead, we simply reasonably assume them to have some practical, operational, measurable meaning; for instance, connectivity may range from absent (0) to direct (1), with possibly many shades of connection quality in between, promptness may range from non-real-time (0) to hard real-time (1), with all the nuances of real-timeliness in between [9], and association may range from uni-directional (0) to bi-directional (1), possibly also considering the direction of the association (PA to DT, or DT to PA) and its nature (for monitoring the PA or controlling it).

In its most general terms, the *Degree of Entanglement* (DoE) between a DT and its PA is defined in terms of C, P, A by a function f mapping those three values to another real number between 0 and 1:

$$f : C \times P \times A \mapsto \theta \in [0, 1] \subset \mathbb{R}$$

Such a definition captures the already discussed intuition of the term entanglement – and, consequently, of the term DoE – in the sense of how tight is the coupling between the DT and its associated PA. In fact, it is intuitively reasonable to consider that a DT connected to its associated PA via a direct, real-time or high-quality communication channel, and bi-directional connection (e.g. the DT of a remotely controllable vehicle in a warehouse) has a higher DoE with respect to a DT connected to its PA via an indirect, non-real-time, and uni-directional connection (e.g. the DT of an urban intersection, whose state is captured by CCTV at a remote monitoring room). As already discussed for C, P, A , defining such a function, in general, may be impossible, and is out of the scope of this paper.

The specific criteria upon which to trigger f can be customized according to the application domain, however, it is desirable that such a function is evaluated *periodically* independently of the reception of events from the PA, as the absence of any communication with the PA is likely to indicate that the DoE is degrading. As regards the duration of such period, heuristically, it must be set so as to enable the DT to become aware of a change in its DoE faster than (i) the maximum frequency at which the PA changes state (h_{PA}), and (ii) the maximum frequency at which applications interact with the DT (h_{DT}).

$$p_f < \frac{1}{h_{PA}} \wedge p_f < \frac{1}{h_{DT}}$$

This way, as soon as an interaction occurs, the DT has the most recently updated DoE available for usage. In practice, one could set up a process triggering f when either a PA event arrives to the DT, or the configured period p_f expires.

The DoE, in general terms, can be utilized for at least three goals: to drive the DT lifecycle; to condition the *shadowing function* responsible for mirroring the PA with the DT according to its abstraction model [7]; and to alert applications about its change, so that they can then react appropriately (e.g. if θ is below a given threshold, or decreases at a given pace, they could consider the DT no longer faithfully representative of the PA state and thus ignore DT events until θ returns acceptable).

As regards the shadowing function, let $S_{DT} = \langle \mathcal{P}, \mathcal{R}, \mathcal{E}, \tau \rangle$ be the state of a DT at any given time τ (the logical notion of time local to the DT itself), where: \mathcal{P} is the set of PA properties that the DT mirrors, \mathcal{R} is the set of relationships with other DTs that the DT has, and \mathcal{E} is the past events thread that the DT stores, composed of the sequence of events e_{PA} that the DT has already received from the associated PA, as well as the sequence of events e_{DT} that the DT has already received from the applications it serves.

Let also $e_{PA} = \langle \mathcal{P}, \mathcal{M}, t \rangle$ where: \mathcal{M} is a set of meta-data the DT can measure while interacting with the PA, such as estimated network latency, processing times, etc., and t is the notion of time locally available to the PA (that may be, in general, totally unrelated with τ).

Given the above, we can define the shadowing function of the DT as the one responsible to compute the next S_{DT} based on the current S_{DT} , the most recent event e_{PA} received, and the currently available DoE:

$$Shad_{PA \rightarrow DT} : S_{DT} \times e_{PA} \times \theta \mapsto S'_{DT}$$

As for function f , defining precisely how such a function operates with its parameters may be hard or even impossible in general; however, Section V instantiates the above-described general framework in a specific use case of digital factories. There, precise characterisations of P, C, τ, t , and \mathcal{M} are proposed, and a specific structure is given to both f and $Shad_{PA \rightarrow DT}$.

With this formalization of the shadowing function, the DoE can influence the next DT state: for instance, a property $p \in \mathcal{P}$ may be no longer exposed when θ decreases below a certain




Fig. 3. The relationship over time between the DT-PA shadowing process and the evaluation of the DoE.

threshold, so that

$$S_{DT} = \langle \mathcal{P} \cup p, \mathcal{R}, \mathcal{E}, \tau \rangle \xrightarrow{Shad_{PA \rightarrow DT}} S'_{DT} = \langle \mathcal{P}', \mathcal{R}', \mathcal{E}', \tau' \rangle$$

In summary, having an operational notion of DoE, connected to the DT shadowing function, opens up the way for *adaptation* of DT behaviour depending on contextual conditions, which is a necessary form of autonomy to enable more advanced ones [6].

B. Specific formulations

The generic formulation just given serves as a conceptual framework to think about the DoE on top of the definitions and models already proposed in the state of the art literature, such as connectivity, promptness, and association, as well as the shadowing function and PA events. However, no operational formulation can be given with no assumptions at all about the application domain or the deployment infrastructure where DTs will operate within. As a consequence, the generic formulation is of little practical use.

This section remedy by proposing two specific formulations serving complementary goals: the former is a discretization meant to provide the most simple baseline for a practically useful (operational) definition of the DoE; the second is the one adopted in Section V in the context of the described case study focusing on the impact of network-related aspects on the DoE.

1) *Baseline discrete DoE*: Let's assume P , C , and A to be the simplest discrete values possible:

- $P \in [r, n]$ whether the promptness of PA-DT entanglement is *real-time* (r) or not (n)
- $C \in [d, i, a]$ whether the connectivity between PA and DT is *direct* (d), *indirect* (i), or even *absent* (a)
- $A \in [b, u]$ whether the PA-DT association is *bi-directional* (b) or *uni-directional* (u)

Such values can be easily dynamically computed, based on appropriate (and widely known) mechanisms such as heartbeat messages and deadlines. For instance, a DT may change P from r to n as soon as a soft/hard real-time deadline is missed (or the opposite at the earliest deadline got right), C from d to a as soon as a heartbeat is missed and no “backup connection” is available (such as a camera inspecting a road junction, replacing direct connections with roadside

units sensing vehicles), and so on. Then, the DoE can be given by function $f(P, C, A) \in [0, 8] \subset \mathbb{N}$ as:

$$\theta = \begin{cases} 8 & \text{if } P = r, C = d, A = b \text{ (maximal)} \\ 7 & \text{if } P = r, C = d, A = u \text{ (deep)} \\ 6 & \text{if } P = r, C = i, A = b \text{ (strong)} \\ 5 & \text{if } P = r, C = i, A = u \text{ (synchronous)} \\ 4 & \text{if } P = n, C = d, A = b \text{ (standard)} \\ 3 & \text{if } P = n, C = d, A = u \text{ (weak)} \\ 2 & \text{if } P = n, C = i, A = b \text{ (shallow)} \\ 1 & \text{if } P = n, C = i, A = u \text{ (minimal)} \\ 0 & \text{if } P = \cdot, C = a, A = \cdot \text{ (null)} \end{cases}$$

where symbol \cdot denotes any admissible value (i.e. complete absence of connectivity makes P and A irrelevant). There is one subtle aspect of such discretization that DT designers must keep in mind: it implicitly establishes priority in evaluating P , C , and A (in particular, $C > P > A$). This criterion is arbitrary, however reasonable, as (i) absence of connectivity makes the DoE immediately null; (ii) a bi-directional association does not *per se* indicates a stronger DoE, as a uni-directional association may simply be an application requirement.

2) *Network-aware DoE*: Let us now assume that all that matters for computing DT-PA DoE are network-related aspects. For instance, taking inspiration and reference from the analysis presented in [3] latency of network communications and packet loss. We can then define P as a *weighted normalised message rate* amongst an arbitrary number $N \in \mathbb{N}$ of signals:

$$P = \frac{\sum_i^N \omega_i * \frac{M_i^r}{M_i^e}}{N} \in [0, 1] \subset \mathbb{R}$$

where M_i^r is the *actual* message rate of PA signal i (not necessarily PA events) received, M_i^e is the *expected* message rate of PA signal i to be received in a reference time window – that could naturally correspond to the periodicity of f calculation, that is, p_f , for instance –, and ω_i are weights enabling to emphasise individual signals contribution to the (weighted) average differently. This definition of P provides a simple yet effective way (amongst the many possibly available) to let network latency have a measurable impact on the DoE. It is a way to operationally reify promptness (P) with network latency. We do not restrict M to be PA events e_{PA} since the DoE could be computed more frequently than the maximal or average frequency at which e_{PA} are received. For instance, M could represent a heartbeat specifically meant to enable the DT to compute its DoE.

Similarly, we could define C as a *weighted normalised packet loss* amongst the N signals:

$$C = \frac{\sum_i^N \phi_i * \frac{P_i^r}{P_i^e}}{N} \in [0, 1] \subset \mathbb{R}$$

where P_i^r is the number of packets *received* for a specific PA signal i , e.g. a specific PA property, P_i^e is the *expected* number of packets for signal i (again, in a reference time




Fig. 4. Example of structured modeling and description of the DoE associated with a DT with the WoT Thing Description.

window), and ϕ_i are weights enabling to emphasize individual signals' contribution to the (weighted) average differently. This definition of C provides a way to let packet loss have a measurable impact on the DoE, as a practical way to measure (the quality of) connectivity C .

If we now assume association A to be fixed at 1, as we don't care about the association in this specific example, the DoE could be then computed by function $f(P, C, A = 1) \in [0, 1] \subset \mathbb{R}$ as:

$$\theta = \frac{w_1 * P + w_2 * C}{2}$$

where w_1, w_2 are adjustable weights of the weighted average between P and C .

This formulation is similar to the one adopted in Section V, where a connection with the discretized formulation is given.

IV. INTEROPERABLE DOE REPRESENTATION

DTs have been significantly shaped by a multitude of heterogeneous definitions, models, protocols, and communication patterns. The same applies to the notion of entanglement. This leads to fragmentation of implementation that hinders interoperability. To tame it, a way to characterize the DoE with a well-structured and interoperable representation is crucial. The *Web of Things (WoT)* standard is already well established in the industry, hence is a natural choice for doing so.

The idea of combining WoT and DTs started appearing in the literature [7], [10] and even in some open platforms such as [11]. In this paper, we foster this vision where a DT can exploit the descriptive characteristics provided by the WoT

Thing Description (WoT-TD) in order to represent itself to the external world as the digitalization of a specific physical asset, overcoming the challenges associated with the fragmented IoT ecosystem, where each DT platform and designer employs custom and closed functionalities to structure DT's properties, generated events, and actionable capabilities.

We start with the core WoT-TD definitions and features and extend them to provide a structured knowledge representation for DTs and their DoE. Figure 4 illustrates three WoT-TD's extensions related to:

- *Context & DoE Schema:* We use the *context* field to provide information about the DoE data model used in the thing description, and in particular to define the meaning of the properties, events, and actions with respect to entanglement. This can be used by applications to correctly interpret and use the properties, events, and actions of the DT, even if they use different terminologies or data models;
- *DoE Definition:* We introduced a new optional field denoted as *doeDefinition* to extend the original WoT-TD including custom DoE metadata. Entanglement is thus characterized in terms of *type*, to identify if it is computed through a discrete representation or using network-based metrics (as illustrated in Section III), *version* to identify the used variant (for example associated to the formula to compute the DoE), a human-readable *description*, and *support* to provide a link to additional information about how the DoE is computed or can be accessed, controlled, or interacted with;




Fig. 5. Forklift baseline scenario with two different variations associated to: (a) a critical braking situation; (b) a significant network degradation.

- **DoE Properties:** This field provides information about the DoE, such as whether it is direct or indirect, real-time or not, and bidirectional or unidirectional, as well as the metrics for its computation (e.g., promptness, connectivity, and association). It can also provide additional metadata about the entanglement, such as the type of sensor or actuator used to capture or control the PA, the frequency and accuracy of the measurements, any relevant interaction pattern (e.g., MQTT with a dedicated topic), and the type of input and output data. Through this, applications can query and reason about the DoE to make informed decisions about how to interact with the PA based on their specific business goals and application requirements.
- **DoE Events:** The events metadata field in a WoT-TD can be used to model information about changes in the DoE between a DT and its PA. For example, if the DoE is reduced due to observability issues with the PA or degradation in network connectivity, an external application can be notified through the declared event *doe_status* and the associated interaction pattern (e.g., MQTT with a dedicated topic). This enables the application to take appropriate actions, such as triggering a maintenance request or adjusting the data processing and decision logic accordingly.

This is a first step toward a structured and interoperable representation of the entanglement between a DT and the physical world through the proposed DoE definition. Alternatives such as *ETSI OneM2M* [12] or *Semantic Web* [13] can also be used to support the proposed vision of extending the awareness of digital applications when interacting with DTs. In Section V we present an implementation and experimental evaluation of the proposed WoT-TD-based DoE with the aim to show how it can effectively support the design and development of intelligent cyber-physical applications.

V. EXPERIMENTAL EVALUATION

The aim of this Section is to present a first implementation of the proposed DoE approach together with an experimental evaluation of its capability to decouple the estimation of the entanglement level from digital applications and their business logic by providing a uniform and standardized characterization and description of cyber-physical bound between the physical

asset and the DT. Through the adoption of the DoE, an external application interested in interacting with the twin will operate with a uniform interface exposing a structured representation of the entanglement level without the need to embed in its internal logic the complexity of reading, analyzing, and understanding low-level telemetry messages that are strongly bound to the nature of the thing, its behavior, and implementation. Furthermore, one of the additional objectives of the experimental evaluation is to show how through different configurations and scenarios, the DoE can quickly react to underlying physical variations allowing an extended awareness of the quality and the freshness of the DT digitalization process and the evolution of its status over time.

We designed an experimental scenario mapping a digital factory and executed a group of dedicated tests in a representative safety-critical use case to exemplify the practical implementation of the DoE concept and to highlight its advantages in accurately depicting and communicating the state and behavior of a PA where upholding a high DoE is fundamental for ensuring that the system's essential functions remain uncompromised.

A. Use Case Description

In the target envisioned scenario, a forklift is deployed onto a warehouse, and undergoes several “mission-critical” tasks, such as raising and lowering boxes onto/from the shelves, and




Fig. 6. Evolution of the relationship between DT and PA over time and the consequent adaptation of the application according to the computed DoE.




Fig. 7. (a) Scenario 1: Baseline configuration where everything is operating normally without any variation or degradation associated with the PA, the DT, or the network.; (b) Scenario 2: PA variation to promptly react to change in the amount, quality, and frequency of transmitted data; (c) Scenario 3: Network degradation

moving them around. These tasks are sent by a central control unit *via* wireless connectivity to the onboard computers and then displayed to a human operator. When a task, or a batch of tasks, is completed, the operator notifies the control unit, which in turn sends another (batch of) tasks.

The forklift is equipped with an intelligent camera that can instantaneously detect obstacles, whether human or object, in its path. Should danger arise, the forklift is equipped with emergency braking mechanisms and warning signals for the operator. These functionalities are executed by an independent onboard subsystem housing a dedicated camera and an object/people detector, functioning through a Convolutional Neural Network (CNN) software module. This onboard subsystem ensures that computational resources remain consistently available, enabling swift responses to potential events. During emergencies, however, it might be challenging to uphold a high DoE, as resource allocation becomes vital for crisis resolution.

In this context, direct connectivity is the preferred mode, while synchronization promptness primarily matters when the system control is remote. The association between PA and DT could encompass both state updates (PA \rightarrow DT) and command updates (DT \rightarrow PA). As schematically depicted in Figure 6 the idea is to investigate how the dynamic computation of the DoE can support both the DT and external applications to extend their awareness of the cyber-physical context and consequently adapt to changing environmental conditions. The reported schema merges the modeling introduced in Section III and the target use case characteristics to show two things. On one hand, that the DT is responsible for receiving events and data from the associated PA (the forklift) in order to compute its internal state and evaluate the current DoE. On the other hand, that the forklift management application (e.g., running in the digital factory control room) observes both the DT's state and the DoE in order to understand the currently exposed properties and functionalities provided by the DT (that may change over time according to the context), as well as check the status of the shop floor and decide how to interact with

the PA through its DT (e.g., sending information for the next mission or forcing the device to stop).

We consider three different scenarios as depicted in Figure 5, in order to isolate the possible factors that can affect the DoE over time. At first, the PA, the network, and any other computational resources are operating correctly without any disruption or variation. This configuration represents our baseline condition where we can measure the DoE according to the involved telemetry traces and the associated metrics. In the second scenario, the communication infrastructure is flawless, but the onboard system encounters some issues, namely emergency stops due to people crossing its path, and must reduce the amount and quality of data sent to the infrastructure to devote more resources to promptly react to danger. In the third case, the onboard system is running smoothly, compelling to send data at the supposed rate, but the communication network within the digital factory has issues such as unexpected antenna failures, or insufficient coverage.

The forklift, in optimal conditions, sends to the DT every 100ms a timestamp (a keep-alive heartbeat), the image of a 4K RGB camera, and the position in GPS coordinates obtained by means of sensor fusion between the camera and an inertial measurement unit (IMU). This corresponds to a *maximal* or *deep* DoE. When obliged to reduce the quality of the data, the first option is to reduce the size of the image sent. This will still maintain a reasonably high DoE. In case further reduction of resources for sending data is required, the system can skip the image and send only the position, since it is still computed and needed for the localization and path planning of the autonomous forklift and the data amount is just some Bytes. This would lower the DoE to *standard* or *weak*. Utmost degradation might be sending only the keep-alive timestamp or stopping the sending altogether, leaving potentially the DoE to *null*.

B. Implementation & Results

The depicted scenarios together with the involved DoE computation have been implemented adopting the standard IoT protocol MQTT (Message Queue Telemetry Transport)

[14] configured with QoS (Quality of Service) zero to avoid duplicated messages and using a local WiFi network to connect both the PA and the DT. PAs have been implemented through dedicated Python applications that emulate the behaviors of real devices using pre-recorded telemetry and data traces through different metrics from prior experiments. Each application offers the flexibility to dynamically adjust message rate, latency, packet loss, and transmission windows, ensuring the proper configuration of the target scenario for experimental evaluation. DTs have been implemented using the White Label Digital Twin (WLDT) Java Library¹, a modular software stack based on a shared multi-threaded engine able to effectively implement DTs behavior and to define their shadowing procedures, data processing, and the interaction with external applications [15]. Considering the involved message rates and transmission windows, the evaluation of the DoE is performed directly by the DT in order to provide a responsive metric to monitor its entanglement with the PA.

Results in Figure 7(a) illustrate the baseline configuration where each entity operates normally with the maximum performance and without any degradation in terms of networking and computation resources. The top plot in each image reports the evaluated DoE, while the underlying plots indicate the measured delay in *ms* for each involved telemetry signal (keep-alive heartbeat, position, and image respectively). As expected, the reported DoE is significantly high for the whole experiment, only slightly affected by the normal oscillations of the network with respect to the different payload sizes (involving the transmission of 4K frames) and message rates.

Plots in Figure 7(b) illustrate the second scenario and show how the DoE can quickly capture the variation in the PA allowing the DT and external observers to detect the difference and consequently adapt the behavior. In that specific case, for example, the DT disables both the incoming communication channel (associated to mission control messages) and pauses synchronization of the missing properties. The external applications notify the anomaly to industrial control systems communicating that both telemetry and control functionality have been disabled due to a disruption in the entanglement. During the experiment timeline, we can appreciate multiple variations in the behavior of the PA and how they influence the DoE.

In the beginning, we have a degradation of the image quality but without any negative impact on the delay, message rate, or the number of telemetry signals transmitted. In that first phase, the DoE preserves a high value since the DT has all the required elements to be synchronized with its physical counterpart and the image quality adaptation does not affect the level of entanglement (it might influence an additional metric on the quality of the digital representation according to the DT's model). In the second phase, the forklift progressively stops sending image frames and then also positioning information. In that case, we have a direct impact on the DoE in terms of the number of received telemetry

signals and the consequent capability of the DT to rebuild a trustworthy replica of the PA. Later, telemetry data gradually returns to their original performance and also the DoE reports the expected high value. At the end of the timeline, a disruptive event at the physical layer (e.g., emergency stop due to sudden crossing of a person from a corner) forces the PA to stop sending telemetry data, and without receiving any signals the DT promptly re-evaluates the DoE setting it to 0.

In the last scenario, the onboard system of the forklift is working correctly and smoothly, but suddenly (in the middle of the experiment timeline) the network infrastructure experiences performance degradation and consequently introduces a significant delay (in our scenario randomly distributed between *10ms* and *100ms*) on all the telemetry transmissions and the message rates. Figure 7(c) illustrates the associated timeline, clearly depicting how for a specific time window network performance decreases directly affecting entanglement. The rapid recalculation of the DoE enables both the DT and external observers to swiftly adjust to changes in the operational context. This adjustment prevents erroneous decisions and actions that could result from using outdated information about the physical world. In the later phase of the experiment, as network issues are resolved, the measured DoE values promptly revert to the higher levels that signify seamless operation.

VI. RELATED WORKS

The DTs concept was introduced in 2003 by Michael Grieves in the aerospace field at NASA [16], [17]. DTs evolved rapidly, moving from the manufacturing industries to IoT and Cyber-Physical Systems contexts [18], extending beyond industrial domains as a bridge between physical space and cyberspace [19]–[21]. DTs are not any more mere models of physical assets; they possess core capabilities and autonomously evolve through specific behaviors and algorithms to understand the world, learn, reason, and respond to external applications [2], [8]. However, a shared set of properties and behaviors for DTs is lacking, hindering the establishment of a common background and language for working with DTs. Efforts to address this issue have led to the proposal of a DT taxonomy with dimensions like communication type, DT purpose, accuracy, and synchronization [22]. Nevertheless, there is still no standardized definition of DT properties and capabilities across the literature. The lack of a common approach has resulted in fragmented DT implementations tailored to specific sectors, limiting the potential for interoperable DTs and applications [23]. The ETSI standardization organization is actively working on defining common DT standards, architecture, requirements, and functionalities [24], [25], but it remains specific to ETSI architectures. In this context, there is a need for autonomous and adaptive DTs capable of extended awareness, independent of specific applications [6]. The vision extends to creating an ecosystem where devices, services, and users collaborate efficiently through distributed and interconnected DTs [7], [26], [27]. To achieve this vision, understanding the entanglement between DTs and their

¹White Label Digital Twin - GitHub - <https://github.com/wldt>

physical counterparts is crucial and the consequent definition of the concept of Degree of Entanglement represents a novel and strategic property to enable a complete characterization of digital twins.

VII. CONCLUSIONS

The analysis and definition of the digital twin Degree of Entanglement (DoE) provided in this paper offer new opportunities for DT autonomy and adaptation, while also posing challenges for future research in various application scenarios. The DoE provides a descriptive representation of the linkage level between a DT and its physical counterpart, enabling extended awareness and facilitating autonomous and adaptive decisions based on variations in entanglement over time. By sharing the idea and importance of defining a core set of native DT metrics, starting with the DoE, it becomes possible to describe the digital replica's status and quality without prior context knowledge or specific implementations. This approach fosters observable and interoperable DT environments, allowing entities to understand the relationship between physical and digital layers, make intelligent real-time decisions, and reduce the risk of misaligned representations. Looking ahead, the evaluation of security aspects [28] and the integration of DTs through an edge-to-cloud compute continuum can represent a challenging opportunity for example within 5G and Multi-Access Edge Computing networks [29] represents a challenging opportunity to extend the experimental evaluation of the proposed approach and support the definition, implementation, and deployment of entanglement-aware cyber-physical applications. The adoption of the DoE aligns well with a distributed ecosystem of DTs where the awareness and control of the QoS (Quality of Service) can represent a strategic pillar to dynamically monitor and manage relationships between physical assets and digital services over time.

ACKNOWLEDGEMENT

This work was supported by the H2020 project 5GMETA, grant agreement No.957360

REFERENCES

- [1] P. Bellavista, N. Biccocchi, M. Fogli, C. Giannelli, M. Mamei, and M. Picone, "Requirements and design patterns for adaptive, autonomous, and context-aware digital twins in industry 4.0 digital factories," *Computers in Industry*, vol. 149, p. 103918, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0166361523000684>
- [2] R. Minerva, G. M. Lee, and N. Crespi, "Digital twin in the iot context: A survey on technical features, scenarios, and architectural models," *Proceedings of the IEEE*, vol. 108, no. 10, pp. 1785–1824, 2020.
- [3] P. Bellavista, N. Biccocchi, M. Fogli, C. Giannelli, M. Mamei, and M. Picone, "Measuring digital twin entanglement in industrial internet of things," in *ICC 2023 - IEEE International Conference on Communications*, 2023, pp. 5897–5903.
- [4] T. H.-J. Uhlemann, C. Schock, C. Lehmann, S. Freiberger, and R. Steinhilper, "The digital twin: Demonstrating the potential of real time data acquisition in production systems," *Procedia Manufacturing*, vol. 9, pp. 113–120, 2017, 7th Conference on Learning Factories, CLF 2017.
- [5] M. Song, K. Zhong, J. Zhang, Y. Hu, D. Liu, W. Zhang, J. Wang, and T. Li, "In-situ ai: Towards autonomous and incremental deep learning for iot systems," in *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, 2018, pp. 92–103.
- [6] K. Hribernik, G. Cabri, F. Mandreoli, and G. Mentzas, "Autonomous, context-aware, adaptive digital twins - state of the art and roadmap," *Comput. Ind.*, vol. 133, p. 103508, 2021.
- [7] A. Ricci, A. Croatti, S. Mariani, S. Montagna, and M. Picone, "Web of digital twins," *ACM Transactions on Internet Technologies*, vol. 22, no. 4, nov 2022. [Online]. Available: <https://doi.org/10.1145/3507909>
- [8] R. Minerva and N. Crespi, "Digital twins: Properties, software frameworks, and application scenarios," *IT Professional*, vol. 23, no. 1, pp. 51–55, 2021.
- [9] D. Calvaresi, M. Marinoni, A. Sturm, M. Schumacher, and G. C. Buttazzo, "The challenge of real-time multi-agent systems for enabling iot and CPS," in *Proceedings of the International Conference on Web Intelligence, Leipzig, Germany, August 23-26, 2017*, A. P. Sheth, A. Ngonga, Y. Wang, E. Chang, D. Slezak, B. Franczyk, R. Alt, X. Tao, and R. Unland, Eds. ACM, 2017, pp. 356–364.
- [10] L. Bedogni, S. Manfredini, F. Poggi, and D. Rossi, "Wise: A semantic and interoperable web of things architecture for smart environments," in *2022 IEEE 8th World Forum on Internet of Things (WF-IoT)*, 2022, pp. 1–6.
- [11] "Eclipse Ditto." [Online]. Available: <https://www.eclipse.org/ditto/>
- [12] "ONE MACHINE-TO-MACHINE PARTNERSHIP PROJECT (ONEM2M)," accessed on 02-09-2021. [Online]. Available: <https://www.etsi.org/committee/onem2m>
- [13] I. Szilagyi and P. Wira, "Ontologies and semantic web for the internet of things - a survey," in *IECON 2016 - 42nd Annual Conference of the IEEE Industrial Electronics Society*, 2016, pp. 6949–6954.
- [14] "MQTT Version 3.1.1," September 2014. [Online]. Available: <http://docs.oasis-open.org/mqtt/mqtt/v3.1.1/mqtt-v3.1.1.html>
- [15] M. Picone, M. Mamei, and F. Zambonelli, "Wldt: A general purpose library to build iot digital twins," *SoftwareX*, vol. 13, p. 100661, 2021.
- [16] E. Glaessgen and D. Stargel, "The digital twin paradigm for future nasa and us air force vehicles," in *Proc. of the 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, 2012, p. 1818.
- [17] M. Grieves and J. Vickers, "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," in *Transdisciplinary perspectives on complex systems*. Springer, 2017, pp. 85–113.
- [18] F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital twin in industry: State-of-the-art," *IEEE Trans. on Industrial Informatics*, vol. 15, no. 4, pp. 2405–2415, 2019.
- [19] R. Saracco, "Digital twins: Bridging physical space and cyberspace," *Computer*, vol. 52, no. 12, pp. 58–64, 2019.
- [20] A. Ricci, A. Croatti, and S. Montagna, "Pervasive and connected digital twins—a vision for digital health," *IEEE Internet Computing*, 2021.
- [21] P. Bellavista, C. Giannelli, M. Mamei, M. Mendula, and M. Picone, "Application-driven network-aware digital twin management in industrial edge environments," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 11, pp. 7791–7801, 2021.
- [22] H. Van der Valk, H. Haße, F. Möller, M. Arbter, J.-L. Henning, and B. Otto, "A taxonomy of digital twins," in *AMCIS*, 2020.
- [23] F. Tao and Q. Qi, "Make more digital twins," *Nature*, vol. 573, no. 7775, pp. 490–491, 2019.
- [24] ETSI, "TR 103 844 (SmartM2M); Digital Twins and standardization opportunities in ETSI," (Last Access: 10-November-2023). [Online]. Available: <https://bit.ly/38RGbs1>
- [25] ETSI, "TS 103 846 (SmartM2M); Digital Twins: Functionalities and Reference Architecture," (Last Access: 10-November-2023). [Online]. Available: <https://bit.ly/3EnSyYx>
- [26] Members of the Digital Framework Task Group, "White paper: The gemini principles," Centre of Digital Built Britain, Tech. Rep., January 2018, available at <https://www.cdbb.cam.ac.uk/DFTG/GeminiPrinciples>. Last access: 20210401.
- [27] D. Gelernter, *Mirror Worlds or the Day Software Puts the Universe in a Shoebox: How Will It Happen and What It Will Mean*. New York, NY, USA: Oxford University Press, Inc., 1991.
- [28] M. Letafati and S. Otoum, "On the privacy and security for e-health services in the metaverse: An overview," *Ad Hoc Networks*, vol. 150, p. 103262, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1570870523001828>
- [29] M. Picone, S. Mariani, M. Mamei, F. Zambonelli, and M. Berlier, "Wip: Preliminary evaluation of digital twins on mec software architecture," in *2021 IEEE 22nd International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, 2021, pp. 256–259.