Challenges and Possible Approaches for Sustainable Digital Twinning

Stijn Bellis
University Of Antwerp, Flanders Make@UAntwerpen
Antwerp, Belgium
Stijn.Bellis@uantwerpen.be

Joachim Denil
University Of Antwerp, Flanders Make@UAntwerpen
Antwerp, Belgium
Joachim.Denil@uantwerpen.be

ABSTRACT

The advance in digital twin technology is creating value for lots of companies. We look at the digital twin design and operation from a sustainability perspective. We identify some challenges related to a digital twin's sustainable design and operation. Finally, we look at some possible approaches, grounded in multi-paradigm modelling to help us create and deploy more sustainable twins.

CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability.

KEYWORDS

datasets, neural networks, gaze detection, text tagging

ACM Reference Format:
Stijn Bellis and Joachim Denil. 2022. Challenges and Possible Approaches for Sustainable Digital Twinning. In *ACM/IEEE 25th International Conference on Model Driven Engineering Languages and Systems (MODELS '22 Companion), October 23–28, 2022, Montreal, QC, Canada.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3550356.3561581

1 INTRODUCTION

Sustainable development focuses on development with three dimensions in mind: (a) Economic, (b) Social and (c) Environmental [22]. Sustainable development is defined as the development that fulfils today's needs without compromising future generations' needs. However, human emission of greenhouse gasses and aerosols create an unbalance in the Earth's energy system [1]. Consequently, we see an increase in the earth's temperature, melting polar caps and permafrost regions. Computing as an industry is currently responsible for 2% to 6% of the emissions of greenhouse gasses globally, with a predicted share of 6% - 22% in 2030 [21].

One of the key transformational technologies for the industry is the digital twin. Several definitions of a digital twin exist in the literature; as such, we will only state one: "A set of virtual information constructs that mimics the structure, context and behaviour of an individual / unique physical asset, or a group of physical assets, is dynamically updated with data from its physical twin throughout its life cycle and informs decisions that realise value." [9]. Digital twins create value for companies by integrating the physical and digital world to address complexities and high demands from the market [18]. Digital twins have multiple functionalities. Applications include but are not limited to real-time monitoring, system optimisation, quality control and waste management [2, 16]. Digital twins are applied in all industrial sectors like aeronautics, medical, smart city, and manufacturing.

While digital twins help in sustainable development's social and economic dimensions by optimising system usage, the design of twins largely ignores its impact on electricity consumption. Digital twin implementations heavily rely on computing and networking infrastructure to monitor, predict, and optimise their analogue counterparts. In this paper, we discuss the challenges and possible approaches for the sustainable development of digital twins.

To make the challenges more tangible, we use a motivating example of a simple system with a digital twin. Our system is a heat-sink used to cool a critical component of another system. It has fans to pull in air cool air which is guided over several fins to convectively cool the critical part. The original design of the system is done using a distributed parameter model. The twin makes predictions based on the predicted system critical component loads, air temperature and current conditions if the system will be able to keep cool enough.

To give an overview of the paper, in section 2 we give some background about digital twins and energy and power consumption. In section 3 we discuss the sustainable design and deployment of digital twins and 4 challenges connected with the design and deployment of digital twins. Lastly, in section 4 we describe some possible approaches to these challenges.

2 BACKGROUND

2.1 Digital Twins

Digital twins are virtual counterparts of real systems. Digital twins are created using digital models that are augmented with data from the real-world. Figure 1 shows various forms of the use of digital models. During the design process, digital models are often used to make choices on various decisions. The digital models created serve as the basis for the digital twin. The digital model can also manually be used in the operations phase of the system to identify faults, etc. To detect faults or monitor the system more efficiently, engineers have created tracking simulators that augment and keep the model in synchrony with the real-life system; these are commonly referred to as digital shadows. Digital twins go one step further and close the loop between the virtual model and the real-world system to make decisions. The term filter is defined as components that estimate the state of a system from data, e.g. Kalman filter or particle filter.

MODELS '22 Companion, October 23–28, 2022, Montreal, QC, Canada

ACM ISBN 978-1-4503-9467-3/22/10…$15.00
https://doi.org/10.1145/3550356.3561581

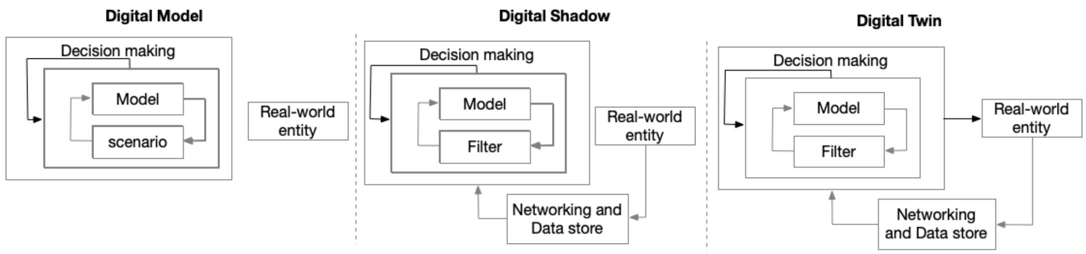

Figure 1: Digital Model, Digital Shadow and Digital Twin

A digital twin is a complex system that needs to be engineered. All the design principles thus also apply to the design of a digital twin.

2.2 Energy and Power Consumption

Energy (E) is the amount of work a system performs over a certain amount of time. The energy consumption of a digital twin is measured in Joules (J) or Watt-hours (W.h). It is the energy required to operate all aspects of the digital twin. Power (P) is the rate at which a system performs work. Power is expressed in Watt. Power and energy are related through $E = \int_a^b P dt$.

Reasoning over energy and power consumption and their associated models can include several levels of impact [4]:

- First order impacts: Impact on the direct production and operation of the digital twin. We focus on this impact in this article.
- Second order impacts: Secondary impact related to the effect of digital twins on, e.g. production and product usage. For example, the decrease of energy consumption of a device because of the optimisation possible by the twin.
- Third order impacts: indirect effects caused by twins, e.g. impacting the structure of an industry or the lifestyle of persons.

3 SUSTAINABLE DESIGN OF TWINS

To understand the problem of sustainable digital twinning, we need to understand where energy is consumed during the life-cycle of the system and its twin.

We break down energy consumption into an additive model where the consumption of the energy occurs: $E_{total} = E_{design} + E_{local} + E_{networking} + E_{cloud} + E_{update}$. With

- E_{design} is the energy consumed for creating the twin. Building a simulation model of a twin might not have a large impact on this factor. However, this term might have a significant impact when using data-driven methods.
- E_{local} is the energy consumption at the analogue side of the system (e.g., by storing the data, pre-processing the data, and executing a part of the twin model locally).
- $E_{networking}$ is the system's energy consumption by sending and receiving messages on the network.
- E_{cloud} is the energy consumption by executing the twin in the cloud environment.
- E_{update} is the energy necessary to redesign and update the model during the system's life-cycle.

This model allows us to reason on the impact developers have during the design and deployment of digital twins.

In the next subsections, we will look at the challenges in developing and deploying a digital twin. The first subsection will look at choosing a digital twin formalism. Then the second subsection looks at how this choice affects the value of the digital twin. The third subsection will look at what happens if the system evolves and lastly the fourth subsection looks at the deployment of a digital twin.

3.1 Twin Formalisms and Boundary Conditions

Different formalisms and workflows exist to create a twin model to use in the digital twin architecture. Most of the techniques shown here transform the model into a less approximate model or a model that is only valid within a certain context. For the approximation, an operating point or operating region needs to supplied. These are commonly referred to as boundary conditions. In this paper we call all these different models surrogate models of the model. We show a formalism transformation graph in Figure 2 to visualise the choice of formalism for the motivating example. A Formalism Transformation Graph shows how formalisms can be transformed into each other [8].

On one side of the spectrum, there are the fully data-driven digital twins. Engineers create this twin by combining big data together with machine learning algorithms. An example is the use of classification and clustering for predictive maintenance [7]. Historical data on performance and failure is combined to train machine learning algorithms to flag when a certain component, machine part or system is likely to fail. The current sensor information is used together with the twin to predict the health status. Engineers prioritise their maintenance schedules based on these outcomes. Note that systems do experience drift and that it is likely that retraining or continual learning techniques might be necessary [15].

On the other end of the spectrum, we have physics-based or simulation-based digital twins. These digital twins run simulations of the system to reason on the system behaviour. The simulation model can be modelled using several formalisms and at several levels of abstraction and approximation (or a combination of). Based on the complexity of the model, more (respectively less) accuracy is obtained at a higher energetic and run-time cost (respectively lower).

In our running example, we start from a distributed parameter model. For this type of formalism, several techniques exist to create surrogate models. Projection-based methods lower the state

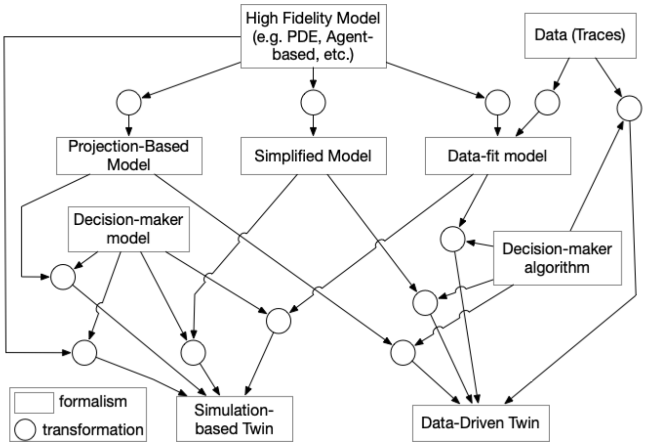

Figure 2: Formalism Transformation Graph showing different means to create the twin model

space dimensions of the problem. An example is proper orthogonal decomposition where the singular value decomposition is used to create a lower-dimensional orthonormal basis. The nonlinear terms still need to be evaluated in an efficient way [6].

Another means to trade run-time and energetic performance with accuracy is to model a simplified model by hand. The modeller needs to have sufficient domain knowledge to do this. The manual modelling requires much insight into the underlying physics of the problem. Yet another means of approximation is at the numerical algorithm level. We can use coarser approximations of the spatial and temporal discretisation employed by methods such as finite differences. Finally, linearising and neglecting non-linear terms is also often employed around operating points.

A third technique are data-fit models that are trained using the input and outputs from simulating the simulation model. Different basis functions are available to map onto. For example, neural networks are shown to universally approximate functions [14].

Besides using a single model and formalism for the twin, some hybrid techniques and architectures are described in the literature that combine several of these techniques. Most notable are ensemble methods that combine several of these techniques together in a single twin.

Challenge 1: Digital Twin Model Energy Consumption: The challenge for sustainable twinning is to create a model that explains and predicts the energy consumption of all possible digital twin architectures created during the system's design. Furthermore, the model should also predict the ramifications of this architectural choice on the energy consumption in the operational phase of the system. The plethora of methods and techniques for creating surrogate models (and thus with the number of available paths toward a digital twin model) is especially challenging. Furthermore, the formalism transformation graph is highly domain-dependent and starts from the initial formalism. If there is no such base model with a very large validity context, the developers are already forced into using combinations of models, where for some of these models, surrogate models are needed.

For the running example, it is possible to create a neural network that clones the physics behaviour of the system within a certain operational area (boundary conditions). This will require a lot of energy to produce such a model as we need a lot of data for training a neural network model (E_{design} is increased significantly). However, once trained, the running of the model in the digital twin will not consume much energy (E_{local}). In comparison, if we use the distributed parameter model with coarser temporal and spatial approximations, the design phase (E_{design}) will not require any additional energy. However, calculating the twin requires a lot of energy every time we need to make a prediction.

3.2 Twin Purpose and Its Value Proposition

In the previous challenge, we neglect several important concepts. The foremost concept is that of the purpose of the twin. Each twin is made for a specific purpose, e.g., the monitoring of the system's health, and the online optimisation of the system. Based on the purpose, a company gains a certain amount of value from the twin. This has to be balanced. When the cost of developing and operating the twin is higher than the value it brings, there is no need for a twin.

From the perspective of simulation engineering, the purpose has a huge influence on the engineering of the simulation model. The same thus applies to engineering the model of the twin. Figure 3 shows a concept framework, from [5], based on [3], extended with a decision maker. When running a simulation, the semantics $[[.]]$ are given in terms of traces. However, we might be interested in certain features of the trace (e.g. the maximum, minimum, integral, etc.). A function $f()$ translates to this value. Based on the system's requirements, we interpret this feature in the ontological domain (e.g. the system is cooling sufficiently). A decision maker (human or automatic) uses this ontological interpretation of the feature and a decision-making model to make a certain decision. The same applies to the approximated model. When the ontological interpretation of the feature is the same, the decision is the same, and the models are both valid (and substitutable) in the context. However, depending on the context, this is not always true. As such, approximations introduce uncertainty on the simulation outcome. Based on this conceptual framework, we see that uncertainty and tolerance play a critical role in the outcome of a decision maker.

The value proposition and the allowed uncertainty introduced by the model are related. If the twin decision maker is tolerant to more uncertainty, then an approximate model with more uncertainty results in the same value as a twin with a detailed model. This allowable uncertainty can be exploited by selecting a model that has the lowest energetic cost that still is valid (and thus substitutable) in the context where the system operates in. Furthermore, in certain problems, it is not necessarily that the decision maker is always perfect for gaining sufficient value from the digital twin. As such, the requirements on the uncertainty can be further reduced.

In our example, the decision maker does not allow for much tolerance within the operational region of the system. This is because the changes in the load have a very immediate response on the temperature of the cooling system. Furthermore, failing to warn the system to be cooled to reduce the thermal load, results in a catastrophic restart.

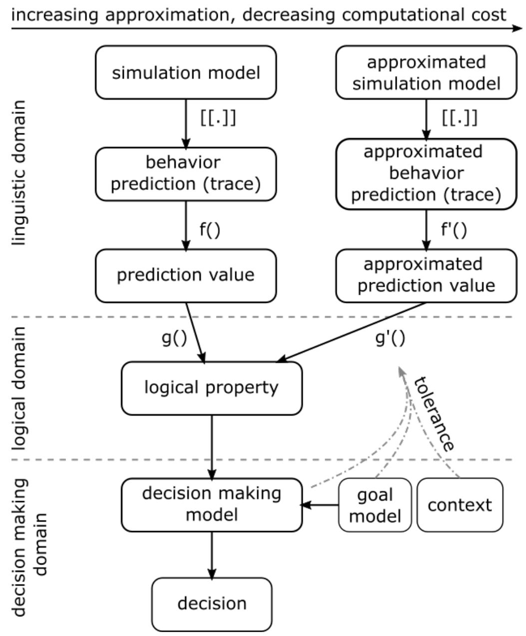

Figure 3: Conceptual framework for the validity of surrogate models, from [5]

Challenge 2: Including Uncertainty into the Formalism Selection Process: Estimating how much uncertainty the decision maker can tolerate gaining enough value from the digital twin is a difficult problem. Furthermore, once we know the allowable tolerance, we still need to include this into the formalism selection process of challenge 1.

3.3 System Evolution

To allow for the long and continuous operation of the digital twin, we need insight into the range of validity of the model in combination with insights into the system's evolution.

Different forms of evolution exist in systems and their digital twin:

- Change in boundary conditions: The evolution entails a change in the system's boundary conditions. In our case study, we could see that the thermal load on the heat sink is beyond the specified load during the design of the digital twin.
- Change in a component of the system: The evolution entails that a sub-component of the system is changed. As such, the digital model should be updated accordingly to the specified change. For example, the fan is changed within the heat sink system, creating a different cold air stream over the sink and thus changing the thermal cooling properties of the system.
- Change in the purpose of the system (and its twin): On the longer term, the system's evolution can entail a different function and purpose. For example, the heat-sink is used to cool an entirely other system than it was initially designed for.

All of the previous evolutions of the system require that the validity of the simulation twin model is checked for the new context of the system or for the changes in the system operation.

Challenge 3: The Impact of System Evolution: Including the dimensions of evolution within the formalism selection process is needed. Having a good estimate of these evolutions in frequency and severity helps determine the needed boundary conditions and validity of the model. If not taken into account, a new model needs to be used, possibly a model that consumes more energy. In data-driven models, it could be that a new model needs to be trained, adding significant overhead to the total energetic cost of the setup.

In our motivating example, evolutions in the system to cool has an impact on the thermal load it supplies to our cooling system. The boundary conditions thus change. The model that underlies the twin should be checked for this new operational context: is the model still valid in this new context? If not, a new model must be selected, trained or developed.

3.4 Twin Deployment Architecture

As previously explained, a digital twin is a complex system on its own. A twin deployment architecture thus needs to consider all different parts of the system. Below, we give a non-exhaustive list of several architectural choices that have an influence on the energy consumption of the twin:

- A first choice that might have a significant impact is where to run part of the twin architecture. Parts of the model can be run in the cloud, edge or locally. This choice could have a significant impact on the energy consumption of the system. One of the effects is that data streams from a local system to the cloud might increase or decrease because of these choices.
- A second choice, influenced by the purpose of the twin, is how the digital twin interacts with the real system. There can be several frequencies at which the system and twin work. If there is a frequency difference between the system and the twin, the difference can be exploited to batch process signals and compress and pack the signals before sending.
- A third choice is on the networking technologies for the data's telemetry. Different technologies are available to provide networking and telemetry, and all impact the twin system's final energy consumption. Some networks are very much low-power, e.g. LoraWAN [10], or configurable, which might result in less power consumption [23]. Similarly, the decision-making and state estimation technologies also have an influence on the consumption of energy. Some algorithms provide very good results but have a heavy computational burden (e.g., particle filters in comparison with Kalman filters).
- Finally, all energy consumption related to storing and retrieving the data from and into data lakes and time-series databases has also a serious impact.

Challenge 4: Deployment of the Twin Architecture : The final challenge is to reason on the deployment choices related to the deployment architecture used for the digital twin. Most choices are impacted by the requirements of the system, which in turn are depending on the value proposition of the twin. Integrating these value and sustainability-related questions into the requirements and architecture phase is needed but challenging as requirements and constraints can change.

4 POSSIBLE APPROACHES IN THE SUSTAINABLE DESIGN OF TWINS

The proposed approaches are very much grounded in the philosophy of Multi-Paradigm Modelling (MPM) [20]. MPM advocates that all relevant system parts are explicitly modelled using the most appropriate formalism(s) at the most appropriate abstraction level(s). The additive model shown in section 3 is already the beginning to examine what models are required to solve the proposed challenges. We first start with a requirements-level view, then a system level view and than look at some individual components.

In this section, we describe some approaches to solve the challenges from the previous section. In the first subsection, we describe the approaches to the challenge of value evaluation. In the second section, we describe approaches for selecting a digital twin formalism. In sections 3 and 4 we describe approaches to predict the energy during the design and operation of the digital twin.

4.1 Value and Allowable Uncertainty

Taking the value of the choices into account, we need to fix on how much uncertainty brings an optimised value for the designer of the digital twin. To quantify the allowable uncertainty, models of the decision maker need to be co-simulated with the model of the simulation model (or traces thereof that are already existing), augmented with uncertainty. The results of these evaluations should be combined with value models from economics. Value is a formal encoding of preference. The outcome is a (response) surface that encodes the relationship between uncertainty and provided benefits. It allows to a trade-off between the incurred energetic costs with the value provided by an uncertain model. However, the energetic cost cannot be the only criterion in the trade study, as we would end with always similar architectures. The design and trade-offs are discussed in the next subsection.

4.2 Selection of the Correct Twin Model Architecture

The selection of a twin model architecture is perhaps the most difficult challenge presented here. The different dimensions of the problem (including evolution and value) make the problem very difficult as most of the information is or might not be available to the twin designers. As the design space of the problem is quite big, design-space exploration techniques have to employed to solve the problem. As we are in a multi-paradigm modelling setting, design-space exploration should not just be done at a single level of abstraction, but at multiple levels of detail, with explicit knowledge, similarly to [11, 17, 25]. As such, we need models at different levels of abstraction for all of the additive parts of the energy consumption model (we discuss the development of such models in the next two subsections). As stated before, we cannot only look at the energetic cost of a system. Other criteria should be considered (e.g., monetary cost of operation, design time, etc.).

Sensitivity information is of particular interest between the different choices and the model's energy consumption. We want to fix the choices with the highest impact first. To evaluate decisions, we will use competitive analysis [12]. Competitive analysis focuses on minimizing the regret incurred by making decisions under uncertainty. Regret allows us to rank different alternative designs and make decisions online. This is needed, as backtracking a choice has a serious consequence on the total energy consumed over the model's life-cycle, e.g., we first learned a neural network with specific boundary conditions to broaden the boundary conditions later.

Finally, we need detailed information on the validity of each model used. Validity frames are needed to capture such information [19, 24].

4.3 Prediction of Design Consumption

We make a distinction between data-driven models and physics-based twins. Physics-based twins do not use a huge amount of energy during the design time. Most energy consumption is involved in techniques for calibration and validation of the models, e.g., Monte-Carlo simulations for proving statistical validity. However, a lot of manual effort is involved in the design of abstractions and approximations. Estimation models for the other criteria are also thus very much needed.

For data-driven twins, the situation is different. Training of black box models does take a lot of computation and energy. Some techniques already exist to predict training time but assume that pre-trained neural networks are used and are fine-tuned for the application [26]. This part of the research still had very foundational aspects. We, therefore, want to establish the relationship between metrics of the dynamic behaviour and energy consumption of the training. Some examples of these metrics include but are not limited to:

- Maximum Lyapunov Exponents that quantify the rate of separation of infinitely close trajectories. The metric gives an insight into the predictability of the dynamic system.
- Singular Value Decomposition is a decomposition of a matrix. When used with trace information of a simulation, the SVD gives information on the orthonormal basis set needed to represent the problem. It is also used in projection methods for this reason.

If these metrics correlate, we create a data fit model to estimate the needed energetic cost (including uncertainty).

4.4 Prediction of Operational Consumption

For each technique to create the twin model, an estimate of how much energy is required to update the twin is needed. For projective techniques, we still need to estimate the non-linear terms. For the other methods, we might be able to relate metrics to the energetic cost (e.g., the number of equations needed to represent the problem). Energy consumption during operation for black box models has already some techniques available [13].

Finally, for all the deployed components, e.g. network components, etc., we need prediction models on the energy consumption at a different level of detail.

5 CONCLUSION

The advance of digital twin technology is creating value for lots of companies. In this paper, we looked at the digital twin design and operation from the perspective of sustainability. We identified some challenges related to the sustainable design and operation of a digital twin: (a) formalism selection, (b) twin value proposition, (c) system and twin evolution and (d) twin deployment. Finally, we looked at some possible approaches, grounded in multi-paradigm modelling, to help us create and deploy more sustainable twins.

REFERENCES

[1] Richard P Allan, Ed Hawkins, Nicolas Bellouin, and Bill Collins. 2021. IPCC, 2021: summary for Policymakers. (2021).

[2] Barbara Rita Barricelli, Elena Casiraghi, and Daniela Fogli. 2019. A survey on digital twin: Definitions, characteristics, applications, and design implications. *IEEE access* 7 (2019), 167653–167671.

[3] Bruno Barroca, Thomas Kühne, and Hans Vangheluwe. 2014. Integrating Language and Ontology Engineering.. In *MPM@ MoDELS*. Citeseer, 77–86.

[4] Frans Berkhout and Julia Hertin. 2001. Impacts of information and communication technologies on environmental sustainability: Speculations and evidence. *Report to the OECD, Brighton* 21 (2001).

[5] Raheleh Biglari, Joost Mertens, and Joachim Denil. 2022. Towards Real-time Adaptive Approximation. In *Embedded Real-Time Systems, ERTS*.

[6] Steven L Brunton and J Nathan Kutz. 2022. *Data-driven science and engineering: Machine learning, dynamical systems, and control*. Cambridge University Press.

[7] Thyago P Carvalho, Fabrízzio AAMN Soares, Roberto Vita, Roberto da P Francisco, João P Basto, and Symone GS Alcalá. 2019. A systematic literature review of machine learning methods applied to predictive maintenance. *Computers & Industrial Engineering* 137 (2019), 106024.

[8] Moharram Challenger, Ken Vanherpen, Joachim Denil, and Hans Vangheluwe. 2020. Ftg+ pm: Describing engineering processes in multi-paradigm modelling. In *Foundations of Multi-Paradigm Modelling for Cyber-Physical Systems*. Springer, 259–271.

[9] AIAA Digital Engineering Integration Committee et al. 2020. Digital Twin: Definition & Value—An AIAA and AIA Position Paper. *AIAA: Reston, VA, USA* (2020).

[10] Jonathan de Carvalho Silva, Joel J. P. C. Rodrigues, Antonio M. Alberti, Petar Solic, and Andre L. L. Aquino. 2017. LoRaWAN — A low power WAN protocol for Internet of Things: A review and opportunities. In *2017 2nd International Multidisciplinary Conference on Computer and Energy Science (SpliTech)*. 1–6.

[11] Joachim Denil, Antonio Cicchetti, Matthias Biehl, Paul De Meulenaere, Romina Eramo, Serge Demeyer, and Hans Vangheluwe. 2011. Automatic deployment space exploration using refinement transformations. *Electronic communications of the EASST.-Berlin, 2006, currens* 50 (2011), 1–13.

[12] Dean P Foster and Rakesh Vohra. 1999. Regret in the on-line decision problem. *Games and Economic Behavior* 29, 1-2 (1999), 7–35.

[13] Eva García-Martín, Crefeda Faviola Rodrigues, Graham Riley, and Håkan Grahn. 2019. Estimation of energy consumption in machine learning. *J. Parallel and Distrib. Comput.* 134 (2019), 75–88.

[14] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. 1989. Multilayer feedforward networks are universal approximators. *Neural networks* 2, 5 (1989), 359–366.

[15] Huiyue Huang, Lei Yang, Yuanbin Wang, Xun Xu, and Yuqian Lu. 2021. Digital Twin-driven online anomaly detection for an automation system based on edge intelligence. *Journal of Manufacturing Systems* 59 (2021), 138–150.

[16] Ziqi Huang, Yang Shen, Jiayi Li, Marcel Fey, and Christian Brecher. 2021. A survey on AI-driven digital twins in Industry 4.0: Smart manufacturing and advanced robotics. *Sensors* 21, 19 (2021), 6340.

[17] Aleksandr A Kerzhner and Christiaan JJ Paredis. 2012. A SysML-based language for modeling system-level architecture selection decisions. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, Vol. 45011. American Society of Mechanical Engineers, 1263–1276.

[18] Mengnan Liu, Shuiliang Fang, Huiyue Dong, and Cunzhi Xu. 2021. Review of digital twin about concepts, technologies, and industrial applications. *Journal of Manufacturing Systems* 58 (2021), 346–361.

[19] Simon Van Mierlo, Bentley James Oakes, Bert Van Acker, Raheleh Eslampanah, Joachim Denil, and Hans Vangheluwe. 2020. Exploring validity frames in practice. In *International Conference on Systems Modelling and Management*. Springer, 131–148.

[20] Pieter J Mosterman and Hans Vangheluwe. 2004. Computer automated multi-paradigm modeling: An introduction. *Simulation* 80, 9 (2004), 433–450.

[21] Copenhagen Centre on Energy Efficiency. 2020. Greenhouse gas emissions in the ICT sector: Trends and methodologies. (2020).

[22] Ben Purvis, Yong Mao, and Darren Robinson. 2019. Three pillars of sustainability: in search of conceptual origins. *Sustainability science* 14, 3 (2019), 681–695.

[23] Conor Sexton, Nicholas J Kaminski, Johann M Marquez-Barja, Nicola Marchetti, and Luiz A DaSilva. 2017. 5G: Adaptable networks enabled by versatile radio access technologies. *IEEE Communications Surveys & Tutorials* 19, 2 (2017), 688–720.

[24] Bert Van Acker, Paul De Meulenaere, Joachim Denil, Yuri Durodie, Alexander Van Bellinghen, and Kris Vanstechelman. 2019. Valid (re-) use of models-of-the-physics in cyber-physical systems using validity frames. In *2019 Spring Simulation Conference (SpringSim)*. IEEE, 1–12.

[25] Yon Vanommeslaeghe, Joachim Denil, Jasper De Viaene, David Ceulemans, Stijn Derammelaere, and Paul De Meulenaere. 2019. Leveraging Domain Knowledge for the Efficient Design-Space Exploration of Advanced Cyber-Physical Systems. In *2019 22nd Euromicro Conference on Digital System Design (DSD)*. 351–358.

[26] Luca Zancato, Alessandro Achille, Avinash Ravichandran, Rahul Bhotika, and Stefano Soatto. 2020. Predicting training time without training. In *NeurIPS 2020*. https://www.amazon.science/publications/predicting-training-time-without-training