



Digital Twin Evolution for Sustainable Smart Ecosystems

Judith Michael
RWTH Aachen University
Aachen, Germany
michael@se-rwth.de

Istvan David
McMaster University
Hamilton, Canada
istvan.david@mcmaster.ca

Dominik Bork
TU Wien
Vienna, Austria
dominik.bork@tuwien.ac.at

ABSTRACT

Smart ecosystems are the drivers of modern society. They control infrastructures of socio-techno-economic importance, ensuring their stable and sustainable operation. Smart ecosystems are governed by digital twins—real-time virtual representations of physical infrastructure. To support the open-ended and reactive traits of smart ecosystems, digital twins need to be able to evolve in reaction to changing conditions. However, digital twin evolution is challenged by the intertwined nature of physical and software components, and their individual evolution. As a consequence, software practitioners find a substantial body of knowledge on software evolution hard to apply in digital twin evolution scenarios and a lack of knowledge on the digital twin evolution itself. The aim of this paper, consequently, is to provide software practitioners with tangible leads toward understanding and managing the evolutionary concerns of digital twins. We use four distinct digital twin evolution scenarios, contextualized in a citizen energy community case to illustrate the usage of the 7R taxonomy of digital twin evolution. By that, we aim to bridge a significant gap in leveraging software engineering practices to develop robust smart ecosystems.

CCS CONCEPTS

• **Software and its engineering** → *Software evolution*; • **General and reference** → *Reference works*; • **Hardware** → *Smart grid*.

KEYWORDS

cyber-physical systems, digital twins, evolution, sustainability

ACM Reference Format:

Judith Michael, Istvan David, and Dominik Bork. 2024. Digital Twin Evolution for Sustainable Smart Ecosystems. In *ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems (MODELS Companion '24)*, September 22–27, 2024, Linz, Austria. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3652620.3688343>

ACKNOWLEDGMENTS

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2023 Internet of Production – 390621612. Website: <https://www.iop.rwth-aachen.de>.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MODELS Companion '24, September 22–27, 2024, Linz, Austria

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0622-6/24/09

<https://doi.org/10.1145/3652620.3688343>

1 INTRODUCTION

Our modern world runs by smart ecosystems—large-scale, decentralized systems, capable of self-organization and self-optimization [13]. Examples of smart ecosystems include smart cities [8], smart energy communities [9], and smart grids with renewable components [10].

Much like natural ecosystems, smart ecosystems are open-ended and need to allow for continuous changes in their structure and behavior. These evolutionary dynamics, in turn, challenge the technical sustainability [18] of smart ecosystems, i.e., their ability to maintain the quality of service over a prolonged period of time [12]. To improve the sustainability of smart ecosystems, proper evolution mechanisms are required to be put in place. While evolution has a substantial body of knowledge in model-driven software engineering [6, 11], hybrid cyber-physical components of smart ecosystems, such as digital twins [14], give rise to challenges traditional software engineering techniques fall short of addressing.

Digital twins are real-time, virtual representations of physical system components [14]. They govern smart ecosystems and provide essential mechanisms and services to assess, simulate, and control the physical infrastructure of smart ecosystems for optimal behavior [17]. Thus, to ensure the technical sustainability of smart ecosystems, first, the technical sustainability of digital twins must be managed. Changes in digital twins boil down to a heterogeneous set of components, including software, hardware, middleware, and IoT devices. The interdependency of concerns severely hinders the applicability of software engineering techniques and even challenges the very understanding of evolutionary needs.

To help software engineers apply their expertise in digital twin evolution scenarios, we provide a *case-based demonstration of the 7R taxonomy* in this paper. The 7R taxonomy of digital twin evolution [4] defines seven elementary activities to support the technical sustainability of digital twins.

This paper is structured as follows. In Sec. 2, we elaborate on a case of an evolving smart ecosystem, driven by digital twin evolution. In Sec. 3, we recommend action points to apply the 7R taxonomy. In Sec. 4, we draw the conclusions. We provide background information about key concepts in sidebars.

2 CASE: CITIZEN ENERGY COMMUNITY

To illustrate the usage of the 7R taxonomy (see sidebar), we rely on a practical case of an evolving smart ecosystem, called the citizen energy community.

Energy communities enable collective, citizen-driven energy actions to support a clean energy transition [2]. In citizen energy communities (Fig. 1), citizens and small commercial entities are equipped with *energy generation* and *storage capacity*, promoting them to first-class generators of energy. As opposed to traditional regulatory models, a citizen energy community gives rise to a *smart*


The 7R taxonomy of digital twin evolution

Taxonomies are a form of classification, aiming to systematically organize knowledge of a specific research field or problem. Classification of objects helps to understand the specific field and systematically treat a particular problem. The 7R taxonomy of digital twin evolution [4] identifies seven areas of action to react to the evolutionary needs of digital twins.

Re-calibration of a model parameter is required when the model is not a faithful representation of the physical twin anymore and simulations become incorrect, leading to imprecise assessment, analysis, and control of the physical twin.

Re-modeling the physical twin might be required in more elaborate cases, e.g., when the model does not reflect the real phenomenon properly. Specific software engineering tasks, such as re-architecting re-packaging a software component might be considered as refinements of this R-imperative.

Reconciliation of data, i.e., updating the data schema and migrating data might be needed when data discrepancies occur, and data might become inconsistent.



Re-collecting data is needed when events are missed due to transient errors. It might necessitate reconciliation, re-modeling, and re-calibration.

Re-deploying the evolved digital twin is needed after at least one of the previous steps has been taken.

Re-configuration of the physical twin is required after the digital twin has evolved. Re-configuration entails a wide range of potential actions, from changing the settings of a physical component to the installation of new ones.

Reuse of the large amounts of data, knowledge, and know-how that have been amassed during the operation of the digital twin is paramount in ensuring cost-efficient digital twin projects.

ecosystem, in which participation is voluntary and egalitarian; and cyber-physical components compose the infrastructure.

A digital twin is developed to govern the smart ecosystem [9] from the very beginning. The digital twin provides stakeholders with tools to monitor and optimize energy trading processes, simulate energy provision and usage scenarios, analyze what-if scenarios, and predict maintenance requirements.

Throughout the lifespan of the system, new features are developed, new components are added, and core elements—often as critical as a power plant—are retired. In the following, we discuss four evolutionary scenarios in an escalating order of impact. By discussing the scenarios through the 7R framework of digital twin evolution for technical sustainability, we demonstrate how to organize the chain of thought about digital twin evolution into a structured set of arguments to support engineering tasks.

2.1 Scenario 1: From a monitoring digital twin to a predictive digital twin

The local government decided to provide financial incentives to residents, who provide the excess energy of their photovoltaic systems within the citizen energy community network. In the new setting, client end-points do not only consume but also produce electricity. However, this setup necessitates accurate forecasting of electricity fluctuations, especially excess electricity to prevent damage, e.g., due to overheating components.

Re-model. Forecasting excess electricity requires a suitable model of the electrical grid. Engineering models that leverage laws of physics are a typical choice. Thus, the grid operator decides to improve the models of the digital twin and re-model the grid by adding models of thermodynamics and external factors, such as atmospheric pressure and relative humidity.

Re-calibrate. With the new models added, the digital twin needs to be re-calibrated. Without calibration, the models would not match

the real system, resulting in inaccurate forecasts. Re-calibration is achieved by manual tuning based on high-quality operational data collected by the digital twin.

2.2 Scenario 2: AI-driven predictions

After realizing the benefits of a predictive digital twin—e.g., improved resource efficiency and safety—the grid operator decides to further improve the predictive capabilities of the digital twin. One problem with the engineering model-based techniques in place is the computing power they require for detailed simulations. As an alternative, AI-based predictive methods are proposed and realized.

Re-collect. The development of the new AI model requires large volumes of data, including data that has not been considered before. Typically, data points that were excluded from the manually-built engineering models due to increased complexity are now becoming of particular interest, such as environmental data (e.g., cloud cover). Therefore, the data collection strategy needs to be revised, and the digital twin should start harvesting the required data points.

Reconcile. Collected data needs to pass through various data processing pipelines aiming to clean and consolidate data and eventually store it in a database. The data management infrastructure needs to be reconciled with the newly collected data. This includes technical aspects (e.g., updating data schemas and processing scripts); and in some cases, addressing the organizational or legal framework (e.g., when working with personal or sensitive data).

Re-model. After reconciliation, re-modeling is required to generate AI-based prediction models that are trained on data from the new data pipelines. The re-modeling, here, concerns the addition of new data quantities and qualities to establish an adequate model for predicting the behavior of the energy community using AI.

Re-calibrate. The evolution of the data and the model require a re-calibration of the model to adjust it to the evolved (i.e., extended) scope, end eventually, again faithfully reflect the physical twin.

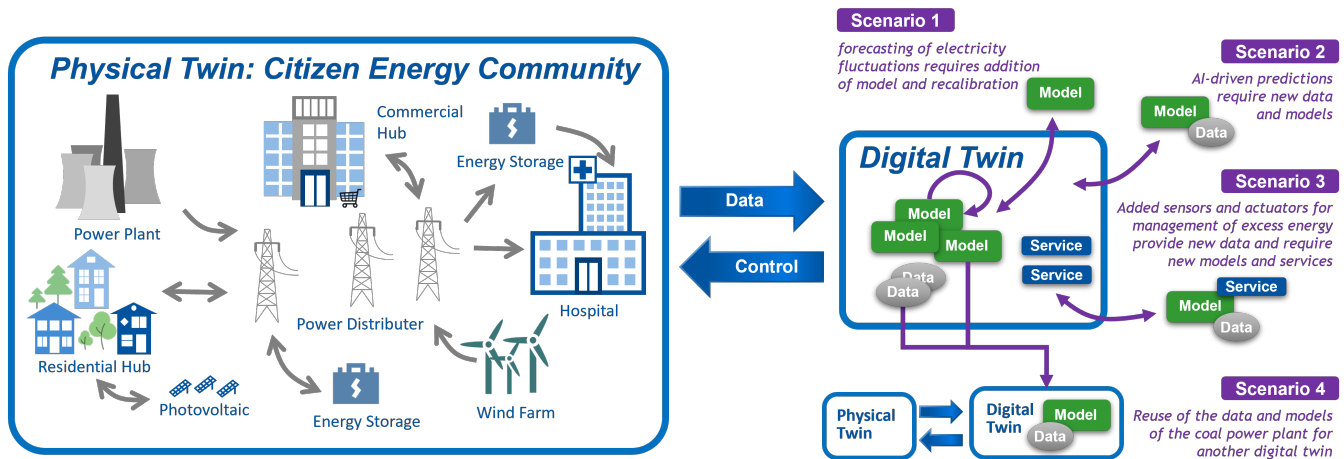


Figure 1: Digital Twin of an Energy Citizen Community evolving over time

Citizen energy communities

A citizen energy community [2] is a localized entity, established with the purpose of generating, distributing, supplying, and storing energy. It enables local energy trading and facilitates the purchasing and selling of energy and energy services to optimize local consumption [9]. Such a citizen energy community consists of citizens, their buildings, small commercial or public entities consuming energy, and different sources producing energy including the citizens and small commercial or public entities. Energy communities are crucial in driving the clean energy transition.

Digital twins of citizen energy communities

A digital twin of a citizen energy community provides a faithful virtual replica of the overall socio-techno-economic system. By that, the digital twin enables the assessment of key indicators, e.g., of sustainability and overall system health, and supports the continuous improvement and evolution of the ecosystem. A digital twin also helps monitor and optimize energy trading processes [22], simulate energy provision and usage scenarios, detect incorrect sensor information, and predict maintenance tasks of power lines, energy storages, or other physical components.

2.3 Scenario 3: Management of excess energy

Too much energy can lead to voltage frequency disturbances in the system. As a result, transformers might trip off to protect themselves from being damaged. This can cause localized blackouts. To further improve the safety of the grid and optimize its efficiency, the operator decides to equip the grid with the latest generation of safety components—sensors that detect potentially hazardous patterns, and actuators that can act upon hazardous situations. As usual, the digital twin operates these components.

Re-configure. First, the physical infrastructure of the grid needs to be re-configured. This re-configuration concerns putting new sensors and actuators in place. The new equipment enables the grid operator to localize causes for inefficient use of the grid and, consequently, to also actuate on identified grid components (e.g., temporal removal of consumers/producers from the grid, or establishment and enforcement of bandwidth limits).

Re-collect. As new sensors are in place that are producing data not considered before, the digital twin has to collect these new data points about hazardous situations such as voltage frequency disturbance or energy overload in specific areas of the grid.

Re-model. For the optimization of the smart grid efficiency, the operators decide to use the existing sensor and actuator components and integrate them to realize an agent who is in continuous interaction with the physical components by an actuation and sensing

relationship. In this respect, a new model is created that supports a reinforcement learning approach [21].

Re-calibrate. The new model in support of reinforcement learning needs to be calibrated. This ensures that the model is a faithful representation of the grid. Calibration is achieved step-wise, by ingesting pieces of data as they arrive on the data stream.

Re-deploy. The data from the added sensors and actuators as well as the results of the developed reinforcement learning approach should be visualized to the users of the digital twin. This requires that the digital twin as a software system has to be re-deployed.

2.4 Scenario 4: Retiring the coal power plant

Eventually, the distributed citizen energy community reaches the level of self-sustainability, efficiency, and safety, where the central coal power plant component is not needed anymore; and political trends drive the obsolescence of coal-fired power generation. As a consequence, the coal power plant is retired. The digital twin, however, is a source of important information thanks to the data collected throughout the lifespan of the coal power plant. Additionally, legal constraints require the grid operator to keep this data for several years for documentation purposes.

Reuse. The grid operator is now able to reuse important design documents, design rationale (engineering decisions), experimental simulation traces, and operative information collected by the digital

Table 1: Primary roles of Software Engineering in Digital Twin Evolution

R-imperative	Involvement of software engineers Primary role	Extent
Re-calibrate	Update models. In major cases: support model engineers and scientists.	🟡
Re-model	Support model engineers and scientists, and refactor models for scalability.	🟢
Re-collect	Integration with sensor APIs and middleware (e.g., messaging).	🟠
Reconcile	Maintenance of data management pipelines, ETL processes, data schemas.	🟡
Re-deploy	Infrastructure-as-Code, DevOps, CI/CD.	🔴
Re-configure	Middleware development, embedded software development.	🟡
Reuse	Software componentization for reuse. Transfer learning from AI components.	🟢

twin during the lifespan of the original power plant. However, effective reuse might require further actions, e.g., re-calibrating models or re-collecting additional data.

Here, we maintain a focus on software aspects. In a system-wide focus, resource value retention options would become additionally important [1, 5], e.g., reusing particular components of a power plant, repairing or replacing parts in the smart grid, or re-purposing buildings leading to changed energy needs.

3 ACTION POINTS FOR APPLICATION

We aim to ease the application of the 7R taxonomy for digital twin evolution. Generally, applying the taxonomy requires answering two questions related to the affected R-imperatives on the one hand and the existing evolutionary processes on the other.

3.1 Which of the R-imperatives does an evolutionary scenario touch upon?

Answering this question helps in understanding the primary roles of software engineering in support of digital twin evolution, and the extent to which software engineering is involved in these phases. Tab. 1 provides typical examples of such roles to every R-imperative.

Re-calibration. This imperative often does not require the involvement of model engineers and scientists; software engineers who are familiar with the model might take care of re-calibration in their own scope. Calibration and re-calibration of models is a moderately software-intensive R-imperative (🟡).

Re-modeling. This imperative, on the other hand, is primarily the concern of model engineers and scientists. The role of software engineers is to take such models and refactor them for scalability. This is typical, e.g., with machine learning models, in which algorithms are fine-tuned by scientists, enabling software engineers to integrate the model into the software architecture. Re-modeling is one of the least software-intensive R-imperatives (🟢).

Re-collecting. Re-collecting data typically requires working with device APIs or interacting with a messaging middleware. It is a fairly software-intensive imperative (🟠) that touches upon distributed components and often runs into testing challenges.

Reconciliation. The software engineering effort focuses on maintaining data management pipelines as the underlying data collection infrastructure changes. This is a fairly critical and software-intensive imperative (🟡), as it touches upon data, a key value driver for companies [15].

Re-deployment. This imperative is typically the most software engineering-intensive one (🔴). As computing is typically located in the cloud nowadays, software engineers need to define the overall infrastructure-as-a-code [20] for deployment, as well as enact the end-to-end DevOps or, in rare cases, CI/CD processes.

Re-configuration. Re-configuration of the physical infrastructure mostly requires interacting with middleware as physical components are mostly hidden behind messaging and procedural layers. Occasionally, developing and maintaining embedded software for physical devices might be required, which is typical in specialized cases, e.g., where custom measurement equipment is used. Still, this imperative is only moderately software-intensive (🟡).

Reuse. This imperative can be supported by software engineering [16] by proper componentization of software, preparing it to be used in other digital twinning projects. AI-heavy companies might want to retain value from their previously trained AI components by transfer learning [7]. As reuse in digital twin settings is a more pressing challenge on the physical side of things, this R-imperative is one of the least software-intensive tasks (🟢).

3.2 What are the processes in the organization?

Answering this question helps organize the R-imperatives into a coherent flow. Taxonomies only define a classification of concepts and defer the operationalization to the specific context of the organization or company. Thus, a process model or DevOps variant [3] is required to operationalize the taxonomy. These operationalizations might differ in their extent, intent, and vendor dependence.

Extent: short versus long loops. In the demonstrative case, Scenario 1 is a relatively short loop. It requires implementing a new model and re-calibrating it. In contrast, Scenario 3 is a more elaborate one, touching upon all but one R-imperative. Clearly, the shorter the loop, the easier it is to oversee and manage. Evidence from the industry also shows that shorter loops, especially on the digital side of things (i.e., touching upon re-modeling, re-calibration, and re-deployment), are more frequently situated within the traditional realm of software engineering companies. Longer loops tend to extend into other domains and require more elaborate cooperation.

Intent: data-first versus model-first. In the demonstrative case, we show one particular sequence of R-imperatives for each scenario. In practice, R-imperatives can be chained in a different order and with more cycles to achieve the evolutionary goals of digital twins.

Often, the preferred order of R-imperatives depends on company best practices and employed paradigms.




Figure 2: Operationalizations of the taxonomy in Scenario 3

Fig. 2 shows two typical operationalizations of Scenario 3. In a *data-first* approach, the physical twin is re-configured, and subsequently, data collection and reconciliation start immediately to drive model creation in a deductive fashion. The discussion of Scenario 3 in the running example followed a data-first view. Alternatively, in a *model-first* approach, the re-configuration of the physical twin is followed by re-modeling, re-calibration, and re-deployment of the digital twin. The benefit of this approach is that models can be used to re-generate data schemas and processing scripts, and thus, data collection can commence smoothly, almost without manual intervention. Software companies adopting model-driven practices [19] might venture into model-first evolutionary processes, but the data-first mindset is still prevalent in practice.

Vendor dependence. Operating smart ecosystems is seldom a one-person show. Software companies work with various vendors. Increasingly more often, equipment vendors ship devices coupled with models pre-configured with reasonable defaults. In such cases, longer loops are to be expected, and re-modeling, re-calibration, and re-configuration tasks, in particular, need to be scheduled appropriately. In contrast, internal re-modeling and re-calibration speed up the process but pose challenges in technical aspects, such as maintenance, and non-functional aspects, such as certification.

4 CONCLUSION

This paper provides a case-based introduction to the application of the 7R taxonomy of digital twin evolution. We focus on the role of software engineering in the key tasks outlined by the taxonomy (i.e., its R-imperatives).

Ultimately, the 7R taxonomy of digital twin evolution fosters better decisions in a convoluted problem space in which software engineers are key to success. There are many benefits software engineers can gain from using the taxonomy.

Comprehensive arguments. Scenarios can be explained in a structured way through the taxonomy, leading to more comprehensive arguments. The taxonomy defines the main concerns of digital twin evolution, and by explaining the scenarios through these concerns, the engineers can make sure they think of every important aspect.

Understanding the impact of change. The length of loops is indicative of the complexity of change and how demanding the evolution of the digital twin is. However, more comprehensive and structured impact analysis methods are required, given the critical nature

of smart ecosystems. In the interim, we recommend relying on established cause-effect modeling techniques.

Mostly manual evolution. Mostly due to the physical components. Scenario 1 could have been addressed by putting the right change mechanisms in place. However, Scenario 2 showed that evolution can quickly escalate into the physical realm. Thus, better automation support for digital-physical co-evolution is required.

REFERENCES

- [1] Dominik Bork, Istvan David, Iris Reinhartz-Berger, Sergio España, Giancarlo Guizzardi, and Henderik Proper. 2024. The Role of Modeling in the Analysis and the Design of Sustainable Systems. *Communications of the Association for Information Systems* 54 (2024). <https://doi.org/10.17705/1CAIS.05434>
- [2] European Commission. [n. d.]. Energy communities. <https://energy.ec.europa.eu/topics/markets-and-consumers/energy-communities>.
- [3] Istvan David. 2023. *SusDevOps: Promoting Sustainability to a First Principle in Software Delivery*. Technical Report. arXiv:2312.14843 [cs.SE]
- [4] Istvan David and Dominik Bork. 2023. Towards a Taxonomy of Digital Twin Evolution for Technical Sustainability. In *ACM/IEEE International Conference on Model Driven Engineering Languages and Systems Companion*. IEEE.
- [5] Istvan David, Dominik Bork, and Gerti Kappel. 2024. Circular Systems Engineering. *Software and Systems Modeling* (2024).
- [6] Davide Di Ruscio, Ludovico Iovino, and Alfonso Pierantonio. 2011. What is needed for managing co-evolution in MDE?. In *Proc. of the 2nd Intl. Workshop on Model Comparison in Practice (IWMCP '11)*. ACM, 30–38.
- [7] A. Farahani et al. 2020. A Concise Review of Transfer Learning. In *2020 Intl. Conf. on Computational Science and Computational Intelligence (CSCI)*. IEEE, 344–351.
- [8] Valdemar Vicente Graciano Neto and Mohamad Kassab. 2023. *What every engineer should know about smart cities*. CRC Press, London, England.
- [9] Gabriele Gramelsberger, Hendrik Kausch, Judith Michael, Frank Piller, Ferdinanda Ponci, Aaron Praktikno, Bernhard Rumpe, Rega Sota, and Sandra Venghaus. 2023. Enabling Informed Sustainability Decisions: Sustainability Assessment in Iterative System Modeling. In *ACM/IEEE Intl. Conference on Model Driven Engineering Languages and Systems Companion*. IEEE, 964–968.
- [10] Mohammad Kamrul Hasan, AKM Ahasan Habib, Zarina Shukur, Fazil Ibrahim, Shayla Islam, and Md Abdur Razzaque. 2023. Review on cyber-physical and cyber-security system in smart grid: Standards, protocols, constraints, and recommendations. *Journal of Network and Computer Applications* 209 (2023), 103540.
- [11] Regina Hebig, Djamel Eddine Khelladi, and Reda Bendraou. 2017. Approaches to Co-Evolution of Metamodels and Models: A Survey. *IEEE Transactions on Software Engineering* 43, 5 (2017), 396–414. <https://doi.org/10.1109/TSE.2016.2610424>
- [12] Lorenz M. Hilty et al. 2006. The relevance of information and communication technologies for environmental sustainability – A prospective simulation study. *Env. Modelling & Software* 21, 11 (2006), 1618–1629.
- [13] Jakob Jul Jensen. 2020. Applying a “Smart Ecosystem” Mindset to Rethink Your Products. *Computer* 53, 12 (2020), 98–101.
- [14] Werner Kritzing, Matthias Karner, Georg Traar, Jan Henjes, and Wilfried Sihm. 2018. Digital Twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine* 51, 11 (2018), 1016–1022.
- [15] Douglas B. Laney. 2017. *Infonomics: How to Monetize, Manage, and Measure Information as an Asset for Competitive Advantage*. Routledge. 322 pages.
- [16] Judith Michael, Jérôme Pfeiffer, Bernhard Rumpe, and Andreas Wortmann. 2022. Integration Challenges for Digital Twin Systems-of-Systems. In *10th IEEE/ACM Int. WS on SE for Systems-of-Systems and Software Ecosystems*. IEEE.
- [17] Judith Michael, Maike Schwammberger, and Andreas Wortmann. 2024. Explaining Cyberphysical System Behavior With Digital Twins. *IEEE Software* 41, 01 (2024), 55–63.
- [18] B. Penzenstadler, L. Duboc, C. C. Venters, S. Betz, N. Seyff, K. Wnuk, R. Chitchyan, S. M. Easterbrook, and C. Becker. 2018. Software Engineering for Sustainability: Find the Leverage Points! *IEEE Software* 35, 04 (2018), 22–33.
- [19] Douglas C Schmidt. 2006. Model-driven engineering. *Computer-IEEE Computer Society* 39, 2 (2006), 25.
- [20] M. Staron, S. Abrahao, B. Penzenstadler, and L. Hochstein. 2023. Recent Research Into Infrastructure as Code. *IEEE Software* 40, 01 (2023), 86–88.
- [21] Nikita Tomin, Victor Kurbatsky, Vadim Borisov, and Sergey Musalev. 2020. Development of Digital Twin for Load Center on the Example of Distribution Network of an Urban District. In *E3S Web Conf.*, Vol. 209. 02029.
- [22] Yakubu Tsado, Olamide Jogunola, Femi O. Olatunji, and Bamidele Adebisi. 2022. A Digital Twin Integrated Cyber-physical Systems for Community Energy Trading. In *IEEE SmartGridComm*. 134–140.