Digital Twin Strategies for SoS

4 Challenges and 4 Architecture Setups for Digital Twins of SoS

Michael Borth
ESI (TNO)
Eindhoven, The Netherlands
Michael.Borth@tno.nl

Jacques Verriet
ESI (TNO)
Eindhoven, The Netherlands
Jacques.Verriet@tno.nl

Gerrit Muller
University of South-Eastern Norway
Kongsberg, Norway
Gerrit.Muller@gmail.com

***Abstract*—Cyber physical systems of systems operate on and with machine-generated data that form the foundation of many of their information-centric functions and processes. Especially within infrastructure systems of systems and the applications of the internet of things, data became an asset for additional tasks regarding their efficient and effective operations, e.g., predictive maintenance which lowers the total costs of ownership. Many of these tasks are well supported by digital twins but building such digital replicas of systems and processes within a system of system is complicated by that context, e.g., due to the managerial independence of the contributing systems that often restricts available information and data sharing. We discuss these challenges together with strategies and architectures that address them and illustrate this digital twin perspective on cyber physical systems and the internet of things with examples from the smart grid and smart building domain.**

Keywords—Digital Twin; SoS Engineering; CPS; IoT

I. Digital Twins for SoS

The progress of cyber physical systems (CPS) led to a rapid raise in the generation and use of machine-generated data [1]. Many such system of systems (SoS) are in fact information-centric, i.e., transport and processing of data and information are central to their performance and dependable operations, while they rely on internet of things (IoT) technology [2] to extend their networked connectivity towards a wide device range.

Manufacturing systems, e.g., execute their tasks by coupling their operations to data from factory automation systems that execute planning or resource allocation as well as to internal and relayed sensor measurements stemming from sensor devices or other systems [3]; other examples include smart buildings [4], in which light settings are no longer controlled by the people in the building, but derived from observed activities and presence data, plus from global concerns like energy saving requirements [5].

The systems are using the resulting machine-generated data for their functional purposes, but it is also valuable for various responsibilities of their owners and operators, e.g., optimization, diagnosis, or predictive maintenance. Digital twins are a technology of choice for such purposes [6] – but, just like system of systems engineering differs from system engineering, building them for cyber physical SoS and IoT comes with its own set of challenges. We discuss those and appropriate strategies and architectures in this paper based on case studies from our work on sustainable power usage.

A. Digital Twins, Shadows, and Threads

Digital twins, connected and synchronized digital replica of physical assets, represent both the elements and the dynamics of how systems and devices operate within their environment and live throughout their lifecycle [7]. They are consistently named as key technology in recent years, e.g., by Gartner [6], especially now that IoT technology becomes widely available for a cost-efficient realization of their connection to the data generated by the sensors and processes of the physical assets. This connection allows their concurrent synchronization with their physical twin that forms the basis for the various analyses they enable, e.g., for predictive maintenance, where actual data is compared to expectations given assumptions of the states of components, allowing the computational assessment of these non-observable states and thus the planning of cost-efficient maintenance schedules that prevent breakdowns [8].

As Grieves explains in [9] [10], the data that connects the physical and the digital world and the bidirectional dynamic interaction of the physical objects and virtual models are key elements of digital twins. For some authors like Kritzinger et al., the level of data integration distinguishes different categories of digital twins [11]. This is visible, e.g., in the term *digital shadow* for setups in which only an automated one-way data flow between the state of an existing physical object and the digital object exists, reserving the term digital twin for cases in which the data flows are fully integrated in both directions and the digital twin might thus even act as controlling instance of the physical object.

In our work, we also see the use of the term digital shadow directed at setups that are restricted to single analysis purposes, e.g., diagnostics, reserving the term digital twin for a complete virtualization of all aspects, but the definitions are compatible: as Fig. 1 below shows, specialists can conduct various tasks with digital twins and automated process control is one of the possible tasks, albeit one requiring the back-flow of control information towards the systems.

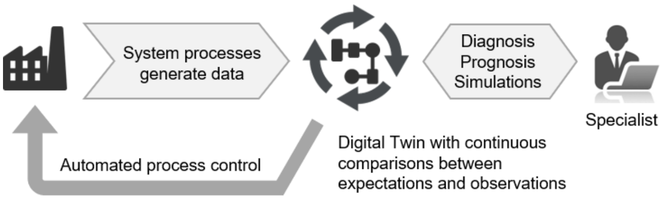

Fig. 1. Digital Twins.

Preprint of https://doi.org/10.1109/SYSOSE.2019.8753860 IEEE SOSE 2019 in Anchorage, Alaska, USA ©2019 IEEE

The point of using the term digital shadow to illustrate the focused aspect of a virtualization is that a full digital twin that mirrors all aspects of its physical counterpart, and thus intents a complete coverage of relevant entities, processes, incl., e.g., their causality and typical behavior is desirable, but difficult to achieve. A digital shadow is, by comparison, restricted, figuratively speaking 1-dimensional, but simpler to realize. The difference can, e.g., be seen on the diagnosis task. A digital twin achieves diagnosis, as it links failure modes to observations in a cause-effect model. A shadow may be restricted to a case-based link between past observations and the respective outcome of diagnosis efforts by human experts, thus omitting the search for the root causes and eliminating the need to model causality.

The imprecision in the terminology shows even more on the concept of the *digital thread*. It can describe the data flow between the digital and the physical twins, refer to the communication framework that allows a connected data flow [12], or to the traceability of the digital twin back to the requirements, parts and control systems that make up the physical asset [13]. Fig. 2 illustrates the latter interpretation, which we find most relevant:

System engineering weaves the digital thread during the matching development of the physical and digital twins, which *(i)* allows to use the digital twin early, especially for quality assurance like virtual verification as Tao et al. describe in [14], *(ii)* prevents any reverse engineering during a cumbersome hindsight setup of digital twins, and *(iii)* provides the link between both digital and physical twin and the engineering knowledge that led to their creation, which is an indispensable asset in the support and management of them over their full lifetime **Error! Reference source not found.**.

B. *The Need for Dedicated Twin Strategies for SoS*

Given the communicated advantages of pairing systems with a digital twin, it is not surprising that both engineering and data companies made their services available to serve this growing market. We see this especially within the manufacturing domain where Industry 4.0 – and thus the promise of intelligent production networks that realize high customization efficiently – led to many activities that extended existing production automation systems towards digital twin concepts (see e.g. [16]).

Our analysis of such activities does indicate a missing point though: these early achievements typically mirror cyber physical assets that are realized as distributed systems; even though they are composed of individual systems, they do not exhibit the core of Maier's criteria [17] for SoS, i.e., operational independence and a goal-oriented dynamic composition. As we illustrate in this paper, these differences matter and require the development of dedicated twinning strategies for cyber physical SoS.

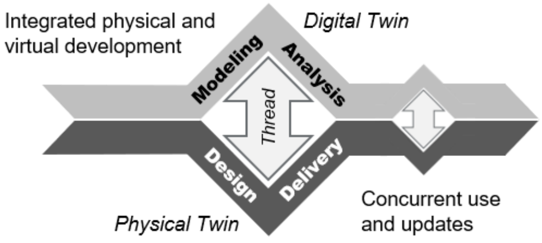

Fig. 2. Digital Twin and Thread. After the MBE Diamond from [18].

II. Digital Twin Perspective on SoS Scenarios

We encountered the challenges to address unique properties of SoS – operational independence, goal-orientation, and the dynamic nature of SoS – in our work on sustainable power regimes, i.e., both the generation and distribution of power and minimization of power consumption in buildings, especially regarding smart lighting. Here, the successful realization of sustainable digital twins or shadows together with their digital thread faces challenges that do not exist in more controlled environments, like factories, and we will use the following sections to present our underlying use cases and the resulting SoS scenarios, incl. the breakdown on which challenges they pose for digital twinning. It is our opinion, though, that our analysis and strategies will apply to other domains as well.

A. *Use Case: Smart Grid*

Looking at the electricity grid, we see that energy production has become highly distributed with the rise of renewable energy. End users can be both consumers and producers of energy, so-called *prosumers*. Consequently, the grid forms a SoS that requires advanced strategies for its control, giving rise to the term smart grid. To manage the available network capacity, which always needs to be in balance, several initiatives aim to use prosumer flexibility to guarantee grid stability and to avoid or postpone costly and time-consuming network improvements. An example is the Universal Smart Energy Framework (USEF), a partnership of smart energy industry players [19]. Using energy flexibility, a prosumer can indicate requirements and preferences with respect to energy consumption and production. This can be used to achieve the SoS-level goal grid stability using auctioning instruments like PowerMatcher [20] that match consumption and production profiles in a virtual market. Such SoS mechanisms benefit greatly from monitoring and control with digital twins.

B. *Use Case: Smart Lighting*

Smart lighting systems are distributed cyber physical systems containing thousands of sensors and actuators connected via a network together with a mix of local and global control that often includes low-level AI techniques for decision making [21]. Typical for modern lighting systems is their cooperation with other systems and thus their realization as SoS. This cooperation includes other building management systems, such as HVAC, blinds and shutters, and security systems, but also data services, e.g., on weather, and external systems.

A prime example of an external system with which a smart lighting system has to cooperation is the smart grid. When the demand for energy exceeds the grid's capacity, smart lighting systems can be demanded to lower their power consumption using Demand-Response, which is defined as *"Changes in electric usage by end-use customers from their normal consumption patterns in response to changes in the price of electricity over time, or to incentive payments designed to induce lower electricity use at times of high wholesale market prices or when system reliability is jeopardized"* [22]. Smart lighting systems may respond to a reduction demand in several ways, e.g., by switching off lights or lowering the light level in non-essential locations, which requires an awareness of the situation, e.g., building occupation, usage, and power consumption. Both monitoring and response are facilitated well with digital twins.

C. Long Lifetime of Infrastructure SoS

As an infrastructure SoS, the power grid has a lifetime of decades which leads to all kind of challenges for engineering, e.g., as upgrades are hampered by legacy systems that do not offer modern interfaces or the lack of knowledgeable engineers who could address this. Pileggi et al., e.g., explain in [23] how this complicates the creation of a digital twin for a smart grid SoS. One identified challenge involves the *unavailability of adequate data*. Typically, the smart grid has not been designed with a digital twin in mind and the data being communicated in the system is not sufficient for the purposes we discuss here. Taken together with the fact that digitalization or data techniques update more quickly than the infrastructures, we see that such SoS typically lag behind the state-of-the-art in IoT and digital twins. While this is often accepted to a certain degree, the grid's infrastructure and constituent systems will be upgraded from time to time, resulting in an *evolution process* that renews the SoS piece by piece. Any such step will invalidate an existing digital twin and lead to the necessity to upgrade it in parallel – especially if the digital twin itself is part of the SoS control structure, as it is the case for power regiments.

A similar complication arises from a different origin, i.e., the degradation of physical properties over long lifetimes. Over time, the performance of smart lighting systems' LED luminaires and the smart grid's batteries and PV panels degrades. This results in the need to change the corresponding models in the digital twin accordingly, e.g., by continuous model calibration.

The ability to realize such updates of both physical systems and digital twins is bound by the above-mentioned lack of available knowledgeable engineers, which directly results from SoS lifespans of several decades. Knowledge gained during design and upgrades of both systems and twins should therefore not depend on people but must be kept in the organizations [24], as it is essential for system evolution.

D. Goals and Conflicts in Coalitions of Systems

Control systems operate well on data readings from sensors, but cooperation in SoS is organized via shared goals and the needed contributions to these goals. It makes sense that a twin for SoS needs to monitor if that contribution is fulfilled. However, that may be neither straightforward, nor unequivocal, as our use cases show. Here, grid stability is the main purpose of the smart grid, as it is of interest to all stakeholders. On the other hand, there are individual goals: energy producers want the highest possible gain whereas consumers want the lowest possible cost. Grid stability at an optimal price requires a balance of production and consumption: in case of a large surplus of energy, energy producers are required to reduce or even stop their production. For this, they will often receive a compensation for the missed income. As energy surpluses can be predicted, producers may try to unjustly benefit from this compensation and disadvantage others by indicating a foreseen production that is higher than their actual production capacity.

Such kind of deceit calls for fraud detection on SoS level [25], but we face more general challenges as well: sensor readings, financial compensation, and grid stability do not mesh and any notion of goal contribution is rather abstract and ill-defined, eluding the coupling of digital and physical twins.

E. No Sharing due to Organizational Independence

Lack of data and information sharing hampers or down-right prevents the realization and use of digital twins. This is even an issue for single-organization / single-system setups, as Grieves and Vickers explain in [26], e.g., if different departments of a company fail to cooperate. For SoS, our use cases show a set of motives to prevent access to operational data and the sharing of information needed to build the infrastructure that facilitates it.

Privacy concerns and regulations can, as they should, restrict or prevent the sharing of data. We see that they already hamper the introduction of smart meters in family homes, as those can be misused to deduce private information from behavioral patterns that might be computed from electricity usage data. Similar to privacy is the protection of business interests. Even the room occupancy data that smart lighting uses to realize energy savings is considered confidential information by some, e.g., if it allows to deduce the state of a company in fluctuating markets. Furthermore, information hiding is a valid strategy to protect against malicious intent, e.g., obscuring a prolonged absence that generates an opportunity for burglaries.

Both the sharing of data and of relevant information about a system's capabilities and methods to generate that data can moreover be restricted to protect the intellectual property of the system's builder or operator. Direct visibility or possible deduction of a lighting system's message passing can, e.g., help in the reverse engineering of its control system, and will thus often be hidden from a building's management system, the higher-order SoS, even though sharing might ease maintenance.

F. Dynamic Nature of Cyber Physical SoS

Cyber physical SoS show dynamics in three different ways: *(i)* evolution through updates and upgrades, which we addressed in Section C, *(ii)* adaptation to changes in their environment, which is foremost acted upon by constituting systems, and *(iii)* changes to the SoS configuration, aka the join-and-leave scenario. For our use cases, join and leave manifests, e.g., in the completion of new solar installations or the decommissioning of a coal plant. Adaptation on the SoS level is necessary, e.g., to adjust expectations on the generated energy to seasons on a longer time scale, while expectations that digital twins use to monitor the individual systems' performance is both long and short term, e.g., taking the local weather into account.

Because of these dynamics, digital twins for SoS must also have the ability to be changed or reconfigured with little effort. Furthermore, their setup but also their goals must account for unforeseen emerging effects, which have a higher likelihood to occur in these circumstances.

The challenge posed by the dynamic nature of SoS is aggravated by the fact that the isolation of individual systems of a SoS for a divide-and-conquer approach is not always possible; in fact, they might become interwoven by the interdependencies of their actions or goals. In the energy market SoS studied by Pileggi et al. in [23], e.g., this involved the input bids and output contracts of auctioning agents. As these are dependent, it is a challenge to create a digital twin of a prosumer without having to create a digital twin of the auctioneer. This is especially true when one would like to do 'what-if' analyses, as the necessary modelling would include even the mechanisms of game theory.

III. ARCHITECTURE CONCEPTS

Next to suitable processes and SoS operations, which we cover in Section IV, we resolve the special SoS challenges for the generation and use of digital with four architectural concepts: *(A)* the focus on upper echelons with regard to what the digital twin mirrors, *(B)* a modular approach based on causal thinking with regard to how we structure the inner models of digital twins, *(C)* the integration of reflection instruments that provide the twin with something akin to self-awareness regarding its performance, and *(D)* the use of points of loose coupling within the SoS to realize the data connection between the digital and physical twin.

We detail these four concepts in the upcoming sections and illustrate how they help to realize suitable digital twins.

A. *Focus on Upper Echelons*

We can build digital twins at various levels of fidelity. Which level of fidelity we use is a dilemma between the effort for building and maintaining the twin versus the required fidelity. High fidelity requires a detailed twin. The challenge of a detailed twin is that the effort for creation is large, and, even worse, that the effort to maintain it is large – which is ill advised given the long lifetime and dynamic nature of SoS, two of the challenges we identified in Section II.

When we take the DIKW hierarchy [27] as starting point for this consideration, we see that we can build digital twins at data, information, or knowledge levels. The most detailed level is the use of (raw) data, which probably is closest to the ideas of digital twin foundations. For detailed diagnostics, data, e.g. individual sensor readouts, may be a necessity. However, when we use the digital twin for analysis of capabilities, then information (e.g. processed, aggregated, and analyzed data) may be suitable. We may even move higher up to knowledge, i.e. associate meaning to the information and relate various information elements to come to a form of situational awareness, especially regarding how well the high-level goals of an SoS are achieved.

When digital twins operate at higher levels of the DIKW hierarchy, then we make the models less sensitive for dynamics of individual instances of constituent systems, thus lessening our needs to maintain them. This goes together with the advantage that the upper echelons of the DIKW hierarchy may also be less privacy and confidentiality sensitive, addressing the concerns that challenge data sharing. While the latter requires careful design of the data abstraction to achieve such desensitizing, its effects are easily illustrated within our use cases: the use of a prosumer's raw data within a digital twin that runs at grid level allows the energy network provider to run detailed profiling, generating, e.g., predictive models that help to stabilize the grid, but also deduce life situation, habits, and otherwise violate the prosumer's privacy. If, on the other hand, only high-level information is exchanged, e.g., the confirmation that the prosumer will adhere to demand-response requests to save energy, this violation becomes impossible while the goal of grid stability can still be achieved and monitored.

Given the need for more detailed data access for other purposes, e.g., diagnosis as mentioned above, we work on SoS twins that can mimic the DIKW hierarchy within their own modular structure, as laid out in the next section.

B. *Ensure Modularity with Causality*

It is our vested opinion that building digital twins for SoS must follow a modular approach that is founded in causality. Modularity is necessary to address several of the challenges we identified above: SoS evolution with partial and local updates or upgrades, join-and-leave scenarios, organizational boundaries that protect sensible information or relate to goal conflicts all require that the parts or compartments of the digital twin that mirror the respective constituents of the SoS can be changed without a re-design of the whole twin. Any other realization implies far too great efforts, for the realization of the twin as well its verification and validation prior to operational use. Moreover, modularization naturally provides comparable units, addressing, e.g., the fraud detection challenge, as laid out below in Section C.

Composing a digital twin from modules while focusing on the upper echelons of the information hierarchy is, however, not trivial and must establish which data or information is used to *(i)* link systems and to *(ii)* traverse the SoS hierarchy. Our guiding architectural concept for this is causality. Causality ensures conceptional and computational composability, as Pearl explains in [28]. Consequently, it allows us to compute expected effects of changes, investigate 'what-if' scenarios for interventions, but also to have a digital twin make steps between the individual systems and the SoS level. The latter is realized in various ways: If need be, this follows live physical flows – power generation and consumption incl. their fluctuations that impact grid stability as well as building conditions [21]. On a higher level, however, it embodies the effects of behavior and therefore knowledge about contributions.

C. *Safeguard Digital Twin and SoS with Reflection*

Reflection, in the sense of a system-level inner awareness, is the ability of a system – or a digital twin in this case – to reason about its own state and performance. For SoS, we combine these aspects in a notion of system health that is based on a system's ability to contribute to the SoS goals. Reflection for a digital twin offers insight into the twin's information health and its capability to perform its function, e.g., to enable control actions that keep the power grid stable. Information health is based on the availability of needed data and information plus their correctness and, for twins that take uncertainty into account, expected margins of error on observations and distributions over non-observables. The capability to perform its function is based on the fulfillment of right-time, right-place requirements within the digital twin's calculations or reasoning; the calculation of an anticipated spike in energy consumption, e.g., must be done in a real-time window that still allows for corrective actions.

Regarding the challenges that we see for digital twins of SoS, reflection mechanisms are key to detect misbehavior like fraud or other actions that arise from goal conflicts. They can, e.g., compare a system's behavior and contribution to expectations based on engineering knowledge or on both past observations and those about other, similar systems. Furthermore, reflection allows us to detect the need to update the digital twin if it no longer reflects the environment, the SoS configuration, or the SoS goal settings. It is thus an asset for addressing the dynamics we identified in Section II.F.

D. *Digital Twin Accesses Loose Coupling of SoS*

We realize digital twins for IoT and cyber physical SoS that follow the three architectural concepts we described above by exploiting SoS mechanisms of 'loose coupling'. Such mechanisms, which, e.g., Borth discussed in [29] for situation awareness systems, lift the data exchange between constituent systems to the information level, e.g., by processing steps that attach semantic meaning to observations. As such steps typically have functional and non-functional demands beyond simple access to the actual data points, they are often realized with buffers or short-term memory structures within data distribution services, thus partially decoupling the constituent systems in technical terms while allowing the described content lift for the SoS.

We illustrate such a mechanism for a modular SoS *and* digital twin structure from the domain of our use cases in Fig. 3, where we instantiate the digital twin as part of the information processing that the SoS requires anyhow. As the figure shows, this avoids double structures (thus the reference to exploitation above), but it also allows us to insert reflection mechanisms at sensible points while protecting sensible information:

> *(i)* Raw data remains within the buildings where it is available for local digital twin supported tasks like diagnosis. *(ii)* Reflection mechanisms run consistency checks that include comparisons, but only over aggregated behavior data. *(iii)* Overarching SoS goals are maintained via communication of high-level concepts. *(iv)* Modularity provides the means to have a new prosumer join the grid. *(v)* Causal reasoning allows the digital twin to forecast the effects of the new situation.

Overall, we see our architectural concepts in place to address the challenges that our – and other – application area posed to us.

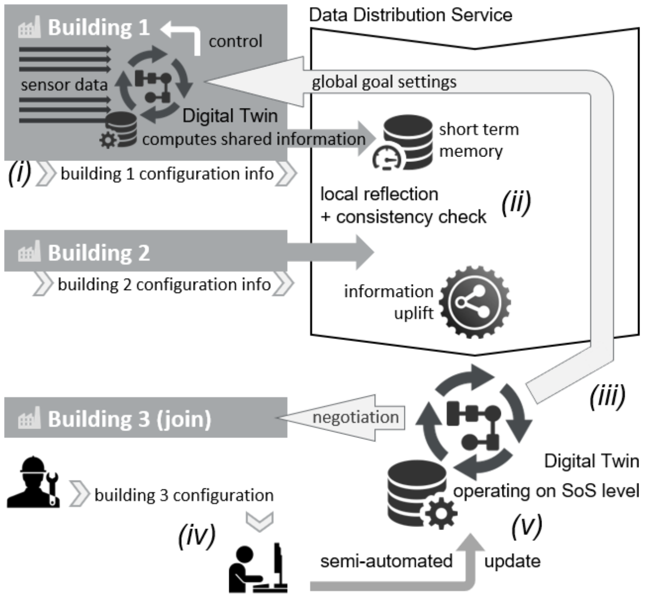

Fig. 3. Conceptual digital twin architecture with loose coupling.

IV. PROCESSES AND OPERATIONS

In the previous sections, we illustrated that setting up digital twins for SoS and IoT systems holds its own challenges and how architecture help to address those. However, we caution that any technical structure is only part of a solution: dedicated processes and operations are needed as well. We precise our main points for those in this section but must refer to other work for details.

A. *Digital Thread for Lifecycle Management*

Similar to established product lifecycle management (PLM) approaches, which integrate data, processes, and business information systems to provide product information management from ideation to end of product life [30], the digital thread is set to connect information about the mirrored system with the constituent parts of the digital twin (Section I.A). As we show in Fig. 2, this is not a one-time activity, but needs to be a continuous workflow in which the digital thread is updated to ensure concurrent operations of the digital twin and the systems it mirrors, which both adapt to changes and evolve over time.

Given fitting knowledge and model management processes, this addresses challenges we identified with regard to the long lifetime of especially infrastructure SoS. For this, we advise to establish, along with many other elements, twin maintenance procedures that work efficiently on modularized computational models (see above), as Borth and van Gerwen illustrate in [31].

B. *Update Automation without Twinning the Twins*

Given that we see the need for continuous maintenance of digital twins, we wish to automate the respective processes – which is also helpful to cope with the organizational independence of SoS, as manual efforts get cumbersome across organizations. At first glance, this seems to warrant a twinning of the twins, i.e., model-based mechanisms that detect a deviation between physical and digital twin and that trigger automated adjustments, just like a digital twin detects a deviation between its model-based expectations and real-world observations to trigger control actions. This is not sustainable if taken to the extremes, as the analogy of the infinity effect that two mirrors facing each other generate shows.

Given an architecture that realizes reflection (see above), we advise combining automated low-level adaptations for local updates with expert-driven revisions on higher levels. The idea here, which is depicted in Fig. 4, is to reserve procedures with high efforts for responses to rare or unpredictable dynamics and changes that involve human decision makers, like changes to the goals or ways of cooperation within a SoS. The reflection mechanisms that we introduced, on the other hand, note local drifts, like the effects of seasons for a smart building, and can trigger local adjustments that do not disturb higher-level goals.

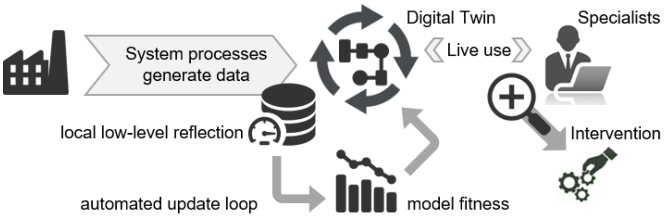

Fig. 4. Digital twin update automation.

V. DIGITAL TWINS TO THE RESCUE

While we listed special challenges for building digital twins for cyber physical systems of systems and IoT installations, we see them as valuable assets to overcome many challenges that are unique to SoS engineering and operations. Using appropriate architectures and strategies, for which we introduced our main approaches here, digital twins help us to

1. overcome managerial shortfalls and pitfalls in the operation of SoS that stem from the operational and managerial independence of the constituting systems, e.g., by providing for fraud detection,
2. address knowledge and version management challenges introduced by the evolutionary development of SoS via the digital thread,
3. safeguard against unintended emerging effects by providing efficient health monitoring mechanisms, and
4. counter extra costs in updates and upgrades caused by the SoS' geographical distribution, as they allow for rigorous virtual experiments and tests during evolutionary progress.

These points – and similar ones that others may see in their own work – show that digital twins address consequences of all of Maier's criteria [17], which make SoS engineering unique, but also challenging. It is in this sense, that we claim that digital twins come to the rescue, as stated in this final section's title.

We require advanced tools and mechanisms to ensure that cyber physical SoS and IoT are worthy of the trust their users (have to) invest in them. Digital twins are among those tools.

REFERENCES

[1] Hewlett Packard Enterprise, "The Intelligent Use of Big Data on an Industrial Scale", InsideBigData, 2017, at https://insidebigdata.com/white-paper/guide-big-data-industrial-scale/.

[2] F. Mattern, Friedemann, C. Floerkemeier, "From the Internet of Computer to the Internet of Things", *Informatik-Spektrum*, vol. 33 no. 2, pp. 107-121, 2010.

[3] J. Schlick, "Cyber-physical systems in factory automation - Towards the 4th industrial revolution," *9th IEEE International Workshop on Factory Communication Systems*, Lemgo, pp. 55-55, 2012.

[4] Building Efficiency Initiative, "What is a Smart Building", 2011, at https://buildingefficiencyinitiative.org/articles/what-smart-building

[5] M. H. Albadi and E. F. El-Saadany, "A summary of demand response in electricity markets", *Electr. Power Syst. Res.*, vol. 78, no. 11, pp. 1989-1996, 2008.

[6] Gartner: Chrysty Pettey, "Prepare for the Impact of Digital Twins", Gartner, 2017, at https://www.gartner.com/smarterwithgartner/prepare-for-the-impact-of-digital-twins/.

[7] Wikipedia, "Digital Twin", 2019, at: en.wikipedia.org/wiki/Digital_twin

[8] R. K. Mobley, "An introduction to predictive maintenance," 2nd edition, Butterworth-Heinemann, 2002.

[9] M. Grieves, "Digital Twin: Manufacturing Excellence through Virtual Factory Replication", 2014, at http://innovate.fit.edu/plm/documents/doc_mgr/912/1411.0_Digital_Twin_White_Paper_Dr_Grieves.pdf.

[10] M. Grieves, "Origins of the Digital Twin Concept," 2016, DOI: 10.13140/RG.2.2.26367.61609.

[11] W. Kritzinger, et al., "Digital Twin in manufacturing: A categorical literature review and classification", *IFAC PapersOnLine*, vol. 51, no. 11, pp. 1016-1022, 2018.

[12] Conrad Leiva, et al., "Enabling the Digital Thread", eBook, iBASEt, 2016, at https://info.ibaset.com/enable-the-digital-thread.

[13] L. S. Gould, "What Are Digital Twins and Digital Threads?", Automotive Design & Production, 2018, at https://www.adandp.media/articles/what-are-digital-twins-and-digital-threads.

[14] F. Tao, J. Cheng, Q. Qi et al. "Digital twin-driven product design, manufacturing and service with big data," *International Journal of Advanced Manufacturing Technology*, vol. 94, pp. 3563–3576, 2018. DOI: 10.1007/s00170-017-0233-1.

[15] G. Cline, "Product Development in the Era of IoT: Tying the Digital Thread," Aberdeen Group, 2017.

[16] Siemens, "One step closer to Industry 4.0," in The Magazine, 2016, at https://www.siemens.com/customer-magazine/en/home/industry/one-step-closer-to-industrie-4-0.html.

[17] M. W. Maier, "Architecting Principles for System of Systems," *Systems Engineering*, vol. 1, no. 4, pp. 267-284, 1998. doi:10.1002/(SICI)1520-6858(1998)1:4<267::AID-SYS3>3.0.CO;2-D

[18] S. J. Hatakeyama, D. W. Seal, Don Farr, and S. C. Haase, "An Alternate View of the Systems Engineering "V" in a Model-Based Engineering Environment," in *Systems Engineering Newsletter* PPI SyEN 73, Project Performance International, 2019.

[19] USEF Energy, "Universal Smart Energy Framework", 2019 at https://www.usef.energy/.

[20] J. K. Kok, C. J. Warmer, and I. G. Kamphuis, "PowerMatcher: Multiagent control in the electricity infrastructure," *Fourth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2005).* ACM, pp. 75–82, 2005.

[21] M. Borth and M. Hendriks, "Engineering the smarts: An illustration of the disconnect between control engineering and AI," *11th System of Systems Engineering Conference (SoSE 2016).* IEEE, 2016.

[22] V. S. K. Murthy Balijepalli, V. Pradhan, S. A. Khaparde and R. M. Shereef, "Review of Demand Response under Smart Grid Paradigm" in *ISGT2011-India.* Kollam, Kerala, pp. 236-243, 2011.

[23] P. Pileggi, J. Verriet, J. Broekhuijsen, C. van Leeuwen, W. Wijbrandi, M. Konsman, "A Digital Twin for Cyber-Physical Energy Systems", *7th Workshop on Modeling and Simulation of Cyber-Physical Energy Systems*, 2019.

[24] M. Sandberg, "Knowledge Based Engineering - In Product Design," Tech. rep. 2003-05, Luleå University of Technology, Department of Applied Physics and Mechanical Engineering, Division of Computer Aided Design, Luleå, 2003.

[25] V. Ford, A. Siraj and W. Eberle, "Smart grid energy fraud detection using artificial neural networks", in *2014 IEEE Symposium on Computational Intelligence Applications in Smart Grid (CIASG)*, pp. 1-6, 2014. doi: 10.1109/CIASG.2014.7011557

[26] M. Grieves, J. Vickers, "Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems", in *Kahlen F. J., Flumerfelt S., Alves A. (eds) Transdisciplinary Perspectives on Complex Systems*, Springer, 2017.

[27] R. L. Ackoff, "From data to wisdom," *Journal of Applied Systems Analysis*, vol. 16, pp. 3-9, 1989.

[28] J. Pearl, "Causality", Cambridge University Press, 2009.

[29] M. Borth, "On the architecture of systems for situation awareness," in *Situation Awareness with Systems of Systems*, Springer-Verlag New York, pp. 39-53, 2013.

[30] PLM Technology Guide, "What is PLM?", 2019, at http://plmtechnologyguide.com/site/?page_id=435.

[31] M. Borth and E. van Gerwen, "Tracking Dynamics in Concurrent Digital Twins," *International Conference on Complex Systems Design & Management*, Springer, pp. 67-78, 2018.