

Universität Regensburg  
Fakultät für Wirtschaftswissenschaften  
Lehrstuhl für Wirtschaftsinformatik I - Informationssysteme

**A two-fold Perspective on Enterprise Security in the Digital Twin  
Context**



Dissertation

zur Erlangung des Grades eines Doktors der Wirtschaftswissenschaft  
eingereicht an der Fakultät für Wirtschaftswissenschaften der Universität Regensburg

vorgelegt von

Marietheres Dietz, M.Sc. with Honors

Berichterstatter:

Prof. Dr. Günther Pernul

Prof. Dr. Stefan Schönig

Tag der Disputation: 07.07.2022

*To my husband, Dr. Julian Berwanger  
and to my son, Leon Berwanger.*

# Abstract

Digital twins represent and can manage an enterprise asset virtually along its lifecycle. The vital technologies the twin relies upon (e.g., Internet of Things) have only recently matured. Since then, literature has taken up on digital twins. The digital twin therefore constitutes a very young concept, where security is currently neglected. This dissertation aims at closing this research gap, and further contributes to the body of knowledge concerning digital twin security. To study digital twin security, a two-fold approach is necessary. On the one hand, digital twins are at risk for being attacked (*security for digital twins*). However, on the other hand, they can also be leveraged to gain novel security opportunities (*digital twins for security*). This dissertation lays the general foundations of the digital twin concept in enterprises and studies these two security perspectives hereinafter. It shows that the digital twin's security can be fostered utilizing the blockchain technology. Furthermore, it proposes digital twins to be of use in corporate security: It is shown that digital twins can collaborate with traditional security tools like Security Information and Event Management (SIEM) systems and organizational structures like the Security Operations Center (SOC). In this regard, the use of digital twins is further proven to be beneficial for digital forensics as well as Cyber Threat Intelligence (CTI).

---

# Acknowledgement

During my years researching and “cumulating” the papers for this dissertation, various persons have given me the support to persevere with my PhD studies. The research papers originating from this dissertation are the effort of multiple authors. Therefore, I would like to express my gratitude to all my co-authors.

A special thanks goes out to my main co-author and supervisor Prof. Dr. Günther Pernul who guided me steadily through the rocky parts of my studies. Even after my change of subject one year after starting my PhD studies, he has expressed his belief in my success. By letting me choose my own subject of interest, but advising me with which research audience and journals to target first, I was able to gain a foothold in the academic world. Also thanks to my second supervisor Prof. Dr. Stefan Schönig for the support and helpful suggestions concerning my research.

Furthermore, I would like to thank my office partner Fabi for the enjoyable atmosphere and his critical remarks, which have enhanced my research. A big thanks to Bene for boldly telling his opinions, introducing me to blockchain/DApp programming and the interesting coffee break conversations (and also for joining me in after-work yoga). I would like to thank Fred and Lucke, with whom I could always have a good laugh next to serious (scientific) discussions. This has greatly relieved the pressure that naturally comes with pursuing one’s PhD and has made writing papers with you a great pleasure. Many thanks to Dani for his eye for details concerning our research work. I would also like to express my gratitude to all the students who have contributed to my research (some of whom are also co-authors and some who have been loyally writing all their works under my supervision). In addition, I am grateful to Petra and Werner, who have given me support through all organizational and IT issues. During my research, I have also enjoyed the input from practitioners. I would like to thank Dr. Thomas Nowey and Andreas Reisser for their helpful input from their industrial experience.

Another steady rock during my PhD studies has been my family. With my husband Julian, who has been in the same boat for years, I could share all my concerns. He has always been understanding and has helped me see things from other perspectives, for which I am very grateful. Our son Leon has given us joy since the day he was born. Finally, I would like to thank my sister, brother, mother and father for their outstanding support and encouragement.

# Contents

<b>Abstract</b>	<b>i</b>
<b>Acknowledgement</b>	<b>ii</b>
<b>List of Tables</b>	<b>v</b>
<b>List of Figures</b>	<b>vii</b>
<b>I Overview of the Thesis</b>	<b>viii</b>
1 Introduction . . . . .	1
2 Research Questions . . . . .	5
3 Research Method . . . . .	8
4 Results . . . . .	11
4.1 Digital twin foundations . . . . .	13
4.2 Digital twins for security . . . . .	15
4.3 Security for digital twins . . . . .	24
4.4 Further publications . . . . .	29
4.5 Roots of this thesis and related works . . . . .	30
5 Conclusion and Future Work . . . . .	34
<b>II Research Papers</b>	<b>37</b>
1 Digital Twin: Empowering Enterprises Towards a System-of-Systems Approach . . . . .	37
2 Unleashing the Digital Twin’s Potential for ICS Security . . . . .	44
3 Integrating Digital Twin Security Simulations in the Security Operations Center . . . . .	54
4 Harnessing Digital Twin Security Simulations for systematic Cyber Threat Intelligence . . . . .	64
5 Enhancing Industrial Control System Forensics Using Replication-based Digital Twins . . . . .	74
6 A Distributed Ledger Approach to Digital Twin Secure Data Sharing . . . . .	93
7 EtherTwin: Blockchain-based Secure Digital Twin Information Management . . . . .	114
<b>A Literature</b>	<b>138</b>

---

<b>B Curriculum vitae</b>	<b>139</b>
<b>References</b>	<b>142</b>

# List of Tables

1	Overview: The seven research papers of this dissertation . . . . .	12
2	Overview of complementary research papers . . . . .	30
3	Publications on digital twin foundations . . . . .	31
4	Publications on employing the digital twin for security until 2021 . . . . .	32
5	Publications on providing security for digital twins until 2021 . . . . .	33
6	Results on digital twin research* . . . . .	138
7	Results on digital twin security research** . . . . .	138

# List of Figures

1	Publications on digital twin (security) over the last years . . . . .	2
2	Iterative, generic design-oriented process [50] complemented by design-science guidelines [27] . . . . .	9
3	Overview of research papers, corresponding research pillars and the research questions answered . . . . .	11
4	The digital twin paradigm [Paper DT-F1] . . . . .	14
5	A system-of-systems approach based on [52], realized by networking systems through their digital twins [Paper DT-F1] . . . . .	15
6	Digital twin security operation modes (dashed lines) and overall model [Paper DT-Sec1] . . . . .	16
7	Process-based security management framework integrating digital twin security simulations [Paper DT-Sec2] . . . . .	17
8	Implementation of digital twin use case and SIEM infrastructure [Paper DT-Sec2] . . . . .	19
9	BPMN-based process for creating CTI from digital twin security simulations [Paper DT-Sec3] . . . . .	20
10	Visualized CTI report of a simulated DOS attack on a digital twin conveyor belt [Paper DT-Sec3] . . . . .	21
11	Replication-based digital twin concept for digital forensics [Paper DT-Sec4]	23
12	Control flows for a single digital twin [Paper Sec-DT 1] . . . . .	25
13	Blockchain-based architecture for secure digital twin data sharing [Paper Sec-DT 1] . . . . .	26
14	Entity Relationship Model of the DApp [Paper Sec-DT 2] . . . . .	27
15	User interface of the EtherTwin DApp [Paper Sec-DT 2] . . . . .	28
16	Overview: publications on digital twin (security) over the last years . . .	138



## **Part I**

# **Overview of the Thesis**

## 1 Introduction

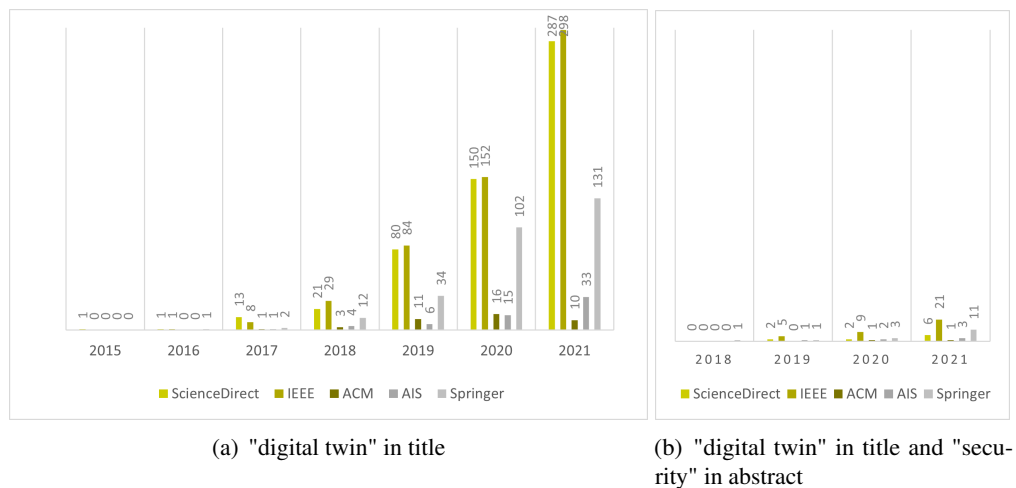
Digitization has altered the way we communicate and deal with information disruptively. From an enterprise perspective, it firstly changed the way to process data allowing more precise decision making: Digital databases and finally enterprise information systems emerged from simple physical documents. Nowadays, the installation of sensors is tackled to acquire additional data in order to further optimize business operations. This phenomenon is known as the *Internet of Things (IoT)*, which describes the connection of sensors and various other devices to Internet-based networks in order to receive information about physical conditions etc. [54]. Up to a few years, corporate information technology (IT) was the main focus of digitization – including concepts like Business Intelligence and Big Data. However, recently, the industrial domain is catching up in the form of the *Industry 4.0* movement. Traditional industrial environments consist of operational technology (OT) – also known as industrial control systems (ICS), which control physical processes [46]. In Industry 4.0, sensors and OT are connected to corporate networks and the respective information systems – resulting in the IT/OT convergence [54]. This enables, for instance, the direct communication between IT and OT systems to determine material reorders, which can then be triggered automatically. Also, sensor data is analyzed with the help of machine learning and Big Data techniques, e.g. to prevent machine outtakes and initiate maintenance tasks.

However, there are always two sides of the same coin: Although digitization presents several benefits, it also entails issues like security and privacy concerns. Ever since the emergence of digitization, *cyber security* incidents occur. Especially by introducing novel concepts (e.g., IoT) without consideration of potential vulnerabilities, the attack surface enlarges considerably. For example, the Mirai botnet abused insufficiently secured IoT devices to execute distributed denial-of-service (DDoS) attacks [34]. Regardless how advanced the digitization concepts are becoming, cyber security attacks are keeping pace: Over the recent years, attacks have become more sophisticated [65]. The so-called advanced persistent threats (APTs) demonstrate deep knowledge about system behavior as well as a targeted exploitation of this knowledge. Thereby, an APT tries to stay undetected as long as possible, mostly to gather information about the targeted victim [65]. One of the first known APTs, Stuxnet [37], targeted the industrial domain. Since then, a myriad of attacks aiming at corporate, industrial or other critical infrastructures have emerged (e.g., attack on the Ukrainian power grid [38] and the Triton APT [47]). Most notably, through the Internet-based connection of various systems (i.e., the IT/OT convergence), the malicious entrance to corporate networks by exploiting vulnerabilities of the connected systems, in order to reach the targeted system, is facilitated.

Thus, novel digitization concepts are to be considered with these two sides in mind. The *digital twin* presents an emerging digitization paradigm, which is especially pursued in the Industry 4.0 [66]. Despite the digital twin was firstly mentioned a decade ago in an aerospace context [62], the maturity and economic profitability of its main technological enablers (e.g., the IoT) has been reached just recently. Until then, the digital twin would

have been too costly for common enterprise deployment. The digital twin can be referred to as a virtual representation of a real-world counterpart such as a system, product, process or any other enterprise asset over its lifecycle [6]. At its core, the digital twin unifies all available asset-specific data and uses semantic technologies, like the industrial standards AutomationML<sup>1</sup> or SysML<sup>2</sup>, to provide context. On this basis, virtual models can be built and analyses (e.g., predictive maintenance) are performed. Furthermore, the virtual models can serve as foundation for conducting simulations [49]. In contrast to similar concepts like digital models or digital shadows, the digital twin distinguishes itself by providing a bidirectional communication to its real-world counterpart [35]. All these unique characteristics entail various opportunities. Consider the digital twin of a windmill collecting system logs and several sensor data from its real-world counterpart. With data analytics and simulations, the digital twin can predict machine fatigue and compare different scenarios for optimization. For instance, the simulation could be used to investigate the degree of the rotor blades in order to reduce velocity and prevent further impairment. Beyond these analytical operations, the digital twin might even send the commands for the best option to the windmill and to stop further damage at once. To this end, the digital twin could also be deployed for security simulations. Nevertheless, the digital twin is also prone to become the target of attacks. Hence, it is necessary to study digital twins in regard to security.

Regarding the articles published on the digital twin concept, it becomes clear that the digital twin's connection to enterprise security has been scarcely researched to date. Figure 1 compares the publications concerning digital twins (see Subfigure 1(a)) in contrast to those concerning digital twin security (see Subfigure 1(b)). Thereto, the



**Figure 1:** Publications on digital twin (security) over the last years

most relevant business informatics databases (IEEE<sup>3</sup>, ACM<sup>4</sup> and AIS<sup>5</sup>) and two more

<sup>1</sup><https://www.automationml.org/o.red.c/home.html>

<sup>2</sup><https://sysml.org/>

<sup>3</sup><https://ieeexplore.ieee.org/Xplore/home.jsp>

<sup>4</sup><https://dl.acm.org/>

<sup>5</sup><https://aisel.aisnet.org/>

broadly focused databases (ScienceDirect<sup>6</sup> and Springer<sup>7</sup>) have been queried. The results<sup>8</sup> show that, even by constraining the publications from the year 2000 onward, the very first publication on digital twins appeared in 2015. This indicates that the digital twin represents a very young field of research. Moreover, works on digital twin security firstly emerged in 2018 and currently present a small amount of those papers on digital twins in general: While in total 1,507 works on digital twin are published until 2021, only 69 (4.6%) of them focus on security. Overall, Figure 1 illustrates the importance of the topic by the strong upwards trend of the number of publications starting in 2015 (resp. 2018). It further indicates that the more technical IEEE database is the one providing the most output on digital twin security (see Subfigure 1(b)). Here, 6.1% (35 out of 572) tackle the security perspective, which amounts almost one and a half as much as regarding all other databases combined. Nevertheless, it has to be noted that the low percentage exists also owing to the fact that security has only come into the focus of digital twin research with a 3-year-delay. Consequently, the compared time spans<sup>9</sup> the topics have been researched on differ (7 years of digital twin research, 4 years of digital twin security research).

Next to this scientific view on digital twins (and security), a look at practice shows that the digital twin has become a vital tool, which has been targeted fairly early. General Electric (GE) is among the pioneers in creating digital twins<sup>10</sup>: Already in late 2017, they monitor about 551,000 DTs referring to products, part of products, processes and systems [57]. Also, other firms, e.g. Tesla in the automotive industry [57] and Siemens [58] work with the digital twin technology. For instance, Siemens Mindsphere<sup>11</sup> is a platform used in many industrial domains, which fully supports harvesting data from connected IoT. Interestingly, the Mindsphere software not only supports digital twin creation when the real-world counterpart is in existence but also when there is no physical counterpart yet available. Furthermore, today's industrial digital twins are integrated in traditional enterprise resource planning (ERP) software (e.g., from SAP) [58]. Glancing into the future, the IEEE Computer Society has marked digital twins in manufacturing as one of the key technology areas in their "Technology Predictions 2022" report [4]: Although digital twins are now considered a relatively mature technology in their current state, they are expected to become even more autonomous. Similar to research, digital twin security is just starting in practice. Once again, GE is one of the first companies that propose a digital twin concept: Their research department recently introduced "Digital ghosts"<sup>12</sup>, a prototypical digital twin used for cybersecurity incident detection and defense.

To conclude, the digital twin presents a novel digitization concept and thus, a young research field relevant to practice. As a consequence, its application is accompanied by various challenges and problems – especially concerning security. This dissertation aims

---

<sup>6</sup><https://www.sciencedirect.com/>

<sup>7</sup><https://link.springer.com/>

<sup>8</sup>Detailed results are given in Appendix A.

<sup>9</sup>Please note that the current results on the year 2022 are not representative as some works are yet to be published until this year's end. Thus, the year 2022 is not considered in the course of this thesis.

<sup>10</sup><https://www.ge.com/digital/applications/digital-twin>

<sup>11</sup><https://siemens.mindsphere.io/en>

<sup>12</sup><https://www.ge.com/research/offering/digital-ghost-real-time-active-cyber-defense>

at contributing to research on digital twin security in order to strengthen the concept and enrich it for enterprise deployment. Therefore, the dissertation takes an *enterprise-centric view on digital twin security*<sup>13</sup> resulting in the following three perspectives:

- **Foundations:** Research carried out first by regarding the digital twin concept in an corporate context and by studying its foundations. On this basis, the digital twin can be researched with respect to security.
- **Security benefits:** Despite the digital twin presents an important digitization concept that builds upon IoT and the IT/OT convergence, it is also discussed to benefit security [21, 55]. Since the emergence of advanced cyber security attacks (e.g. APTs), the detection requires novel techniques. Thereto, the digital twin with its potential of monitoring complex, interconnected systems and their security might become vital for the future enterprise: Digital twins continuously gather information about their real-world counterparts' state, and provide enhanced analysis and simulation capabilities to improve attack detection. With digital twin-based simulations, corporate security teams can further explore the virtually represented system in a potential attack scenario. Such a simulation might even be used as a basis for a cyber range to train threat detection and response.
- **Security problems:** In contrast to the beneficial side of the digital twin concept in terms of security, the second perspective on digital twin security has to be considered as well: First and foremost, the enormous amount of knowledge about an asset gathered in a digital twin might attract attackers. Additionally, the bidirectional connection to its corresponding counterpart enlarges the attack surface and presents a new attack entry point to the real-world system [24]. Moreover, the multiple lifecycle parties involved in the data sharing process of a digital twin requires the implementation of access control strategies.

In conclusion, this dissertation addresses *digital twin security in a two-fold perspective*: While studying digital twin's deployment for enterprise security, it also investigates concepts to enhance the security of digital twins. This provides a novel and comprehensive view at the application of digital twins in enterprises.

The remainder of this Part I provides an overview on this dissertation and is structured as follows: After this introduction (Section 1), the problems in digital twin security and the corresponding research questions are stated in Section 2. Afterwards, the method of this dissertation is described (Section 3). Section 4 presents the results of this dissertation, which comprises seven research papers that provide answers to the research questions. Finally, Section 5 concludes the work and points out future research. Details on the research articles as well as the articles themselves are presented in Part II.

---

<sup>13</sup>Nevertheless, digital twin security might also be of importance in societal and private contexts.

## 2 Research Questions

As indicated in Figure 1, the digital twin presents a novel, and promising research domain. Since digital twin research started less than a decade ago, many published concepts still present drafts and propositions, which often lack a scientifically sound basis, profound evaluation or implementation. Especially security is hardly addressed in digital twin research to date – despite being one of the major issues of digitization concepts. To overcome these issues, it is necessary to investigate the main characteristics of digital twins from an enterprise-centric view and to study digital twins in respect to security on this basis. It is further vital to research digital twin security from the following two different perspectives: On the one hand, digital twins must be sufficiently secured to prevent cyber security attacks. On the other hand, digital twins offer manifold potentials for security. To gather these aspects, this dissertation is structured into the following three main research pillars:

- **DT-F**: Introduction of scientific foundations on the digital twin paradigm
- **DT-Sec**: Investigation of the digital twin’s usage for security
- **Sec-DT**: Provision of security measures for the digital twin

Each of these three research pillars contains one to two central questions with the goal to contribute comprehensive research on these topics. The answers to these research questions (**RQ1-RQ4**) provide this dissertation’s scientific basis.

### **Digital twin foundations: DT-F**

The first research pillar and focused area of this work tackles the foundations on digital twins from an enterprise-centric viewpoint.

Regarding the works published on the digital twin, the majority still investigates its foundations. Most of the works shown in Table 3 (see Section 4.5) give an overview of the state-of-the art in digital twin research by conducting a literature review. However, the majority of them tend to not provide very rigorous methods. For instance, instead of following a published and common method for the literature review, some authors only provide their search term, the queried databases and how they screened the literature. Interestingly, the papers with the most citations are published in manufacturing journals, where they provide less rigorous methods (or no method at all) and lack to propose their own compelling new ideas. In contrast, the general view on digital twins has not gained much interest in terms of citations, despite being published in well-known and respectable journals or proceedings (e.g., [22, 48, 68]). Although, the works presented in Table 3 represent just an extraction, these issues were constantly perceived in the course of this thesis. To sum up, the main focus of these works is set on the manufacturing industry and neglect a view on enterprises in general. In correspondence with the industrial-centered content of most digital twin works, the audience of these papers represent readers from industry. Therefore, these papers mainly appear in industrial conferences or the like,

which often neglect scientific rigor in favor of practice. As a result, the foundations commonly lack a holistic perspective on the topic, and sometimes suffer from a missing methodological structure and arguments.

To close this research gap, the scientific focus of this dissertation's first research pillar (**DT-F**) lays the foundations on digital twins from an enterprise perspective instead of merely concentrating on industry, cyber physical systems (CPS) or IoT. Thereby, current enterprise strategies are investigated to be combined with the digital twin paradigm in order to provide novel approaches to empower modern enterprises. Research question **RQ1** summarizes the scientific intention of this research pillar.

**RQ1:** *How can the digital twin paradigm empower current enterprises in general?*

### **Digital twins for security: DT-Sec**

After studying the foundations of digital twins from an enterprise viewpoint, the next step is to focus on security. The second research pillar of this dissertation tackles the employment of digital twins for enterprise security.

As digital twins are able to monitor and sometimes even control their real-world counterpart's daily operations, literature suggests that they might also provide novel insights and opportunities to security [21, 54, 55]. However, at the beginning of this dissertation no work has tackled this issue so far (see Table 4 of Section 4.5).

As the application of digital twins for security is still scarcely researched to date, a multitude of open issues and aspects are yet to study. Existing publications often present individual solutions to specific areas (like smart grids [2, 9, 56]), with no general overview or investigation on how different digital twin characteristics can be used for security. Some works provide a rather strong focus on the technical implementation – often with restrictions. For instance, the publications of Eckhart and Ekelhart [19, 20] are limited by requiring the data format AutomationML to build the virtual environment. Other works merely present concepts for using digital twins for security, without evaluation or implementation (e.g., [2, 7]). Moreover, the digital twin is commonly regarded as a stand-alone solution when used for security. However, they might also be incorporated with already existing security tools or structures, which could potentially increase the acceptance of digital twins. From a practical perspective, replacing old structures by novel and young concepts like the digital twin will, especially in a field as vital as security, will run into opposition. Therefore, it is necessary to propose ways to combine digital twins and existing security procedures.

To provide a general perspective on digital twin security, different aspects of digital twins are to be examined for their potential use in security operations. Furthermore, it should be suggested how these digital twin properties might support current enterprise security areas (e.g., system testing, digital forensics). These requirements result in research question **RQ2**.

**RQ2:** *What are the foundations to enable security operations with a digital twin?*

While existing literature often presents the digital twin as an individual security solution, the integration of digital twin security operations into corporate security strategies is currently neglected. To bridge this research gap, the contribution of digital twins to modern holistic enterprise security structures like the Security Operations Center (SOC) or Cyber Threat Intelligence (CTI) are to be investigated. Research question **RQ3** addresses this issue.

**RQ3:** *How can security simulations with the digital twin be integrated in the corporate security strategy?*

### **Security for digital twins: Sec-DT**

The third research pillar of this dissertation tackles the second viewpoint on digital twin security. While the second pillar studies digital twins as an instrument to enhance corporate security, this pillar considers the enlarged attack surface digital twins entail and researches security measures for digital twins.

At the beginning of this dissertation, no other works have provided profound concepts to secure digital twins. But literature agrees on the importance of sufficiently securing digital twins [33, 55, 67]. Despite stressing the security issue when referring to digital twins, these works do not propose concepts or implementations to strengthen the security of digital twins. This may also be due to digital twins being a very young research field (see Figure 1). Following from that, early works suggest concepts and approaches. Only a minority of the publications have already tackled digital twin implementation. Without implementation and technical details, it is hard to determine which mechanisms might help to secure the digital twin. However, literature expresses that some aspects, like the secure management of digital twin data storage and exchange [45], are of special importance regarding digital twins.

Commonly, digital twins contain highly confidential and crucial information. They further communicate with their counterpart bidirectionally. Given these special digital twin characteristics, novel approaches are needed to secure digital twins. Especially, since digital twins manage knowledge about an enterprise asset, they will be at the center of attacks. In correspondence with the second perspective and the above mentioned issues on digital twin security, new scientific technologies for data management and exchange (e.g., the blockchain) might prove valuable to secure digital twin data management. Thus, research question **RQ4** summarizes the focus of the dissertation's third research pillar.

**RQ4:** *How can digital twin data storage and exchange be carried out in a secure way?*



### 3 Research Method

This dissertation originates in *Wirtschaftsinformatik (WI)*, a research field in the German-speaking area, which is considered the equivalent to the international information system (IS) research [50]. WI emerged along the digitization movement in enterprises, where new ways of applying IT in future were sought. Thereby, research was mostly conducted pragmatically without the support of rigorous methods. As a result, it initially proved difficult to compete at an international level.

With the growth of WI/IS, the topics diversified. Accompanied by this diversity of research topics is the heterogeneity of applied research methods. The most common methods today employ either a behavioral or a design-oriented approach [50]. As the behaviorist approach starts off with hypotheses and finishes with their subsequent empirical evaluation, it strongly promotes research rigor [27]. However, it has proven to be less useful for practical, real-world problems [50]. This caused a debate about the balance between rigor and relevance among scholars, who consequently broadened their focus with alternative methods, such as the so-called "design science" approaches. Design-based approaches are less rigorous than behavioral approaches but more practice-oriented. They orientate on designing a scientific artifact. This includes the construction of an artifact to solve a relevant problem in IS/WI as well as evaluating the performance of the constructed artifact [27].

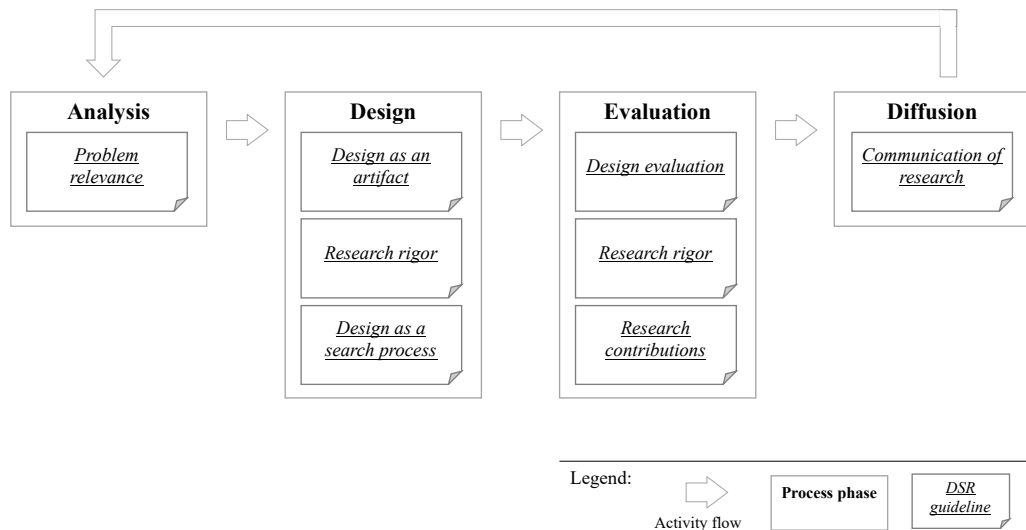
As the digital twin is especially employed in industrial environments [66], its research requires an approach qualified to improve practical problems. Hence, the *design-based research approach* is best suited to build a profound basis for answering the research questions introduced in Section 2. This dissertation follows the iterative design-oriented research process proposed by Österle et al. [50]. Seven guidelines concerning design science research (DSR) provided by Hevner et al. [27] complement this process. Figure 2 illustrates the four phases of this design-oriented process linked with the DSR guidelines.

**Analysis.** The first phase tackles the identification of the problem and its subsequent description. In this course, the research questions and goals are formulated. A research plan structures the path towards the intended research solution answering the established research questions. This phase can be aligned with the following DSR guideline:

*Problem relevance. Technology-based solutions for practical business problems should be developed.*

The field of digital twin security is of importance to literature as well as practice. The underlying problems are stated and formulated into research questions in the previous Section 2. Thereby, three sub-fields are identified, where to one or two research question belongs. The corresponding research plan is presented in Figure 3 (Section 4).

**Design.** The next phase of the DSR process focuses on the creation of a viable artifact. Thereby, well-established methods are applied and the method selection needs to be



**Figure 2:** Iterative, generic design-oriented process [50] complemented by design-science guidelines [27]

stated. Also, the artifact will be delimited from existing solutions. Associated with this phase are the following three DSR guidelines:

*Design as an artifact. A usable artifact (construct, model, method, or instantiation) is to be created.*

*Research rigor. The construction and evaluation of the artifact should apply rigorous methods.*

*Design as a search process. The design of an artifact is a search- or problem solving-process utilizing available means while complying with the laws in the problem environment.*

During this dissertation, several artifacts have been created to answer the research questions. They mostly take the form of a model, framework or a prototype. Each artifact is presented in a research paper (see Part II), where it is compared or delimited from related works. The construction of those artifacts is based on scientific literature: The models, frameworks and prototypes rely on published works, methods or standards. All prototypical artifacts have been created using publicly available technologies.

**Evaluation.** This phase studies the constructed artifact in regard to the specified research goals. Thereto, the methods selected in the research plan are carried out rigorously. The following two DSR guidelines (combined with the above declared guideline concerning research rigor) can be linked with the evaluation phase:

*Design evaluation. The utility, quality, and efficacy of the artifact must be precisely evaluated by methods and rigorously demonstrated.*

*Research contributions. Contributions in the areas of design artifact, design fundamentals, and/or design methods must be clearly demonstrated.*

To evaluate the established artifacts, the works demonstrate its usefulness by comparing it with the stated research goal and mainly apply rigorous evaluation techniques (e.g., use cases, performance measurements, expert interviews). In each work, the selection of the evaluation method is justified. The resulting publications provide the research contribution: The artifacts contribute to the present state-of-the-art and extend the knowledge, propose a novel approach in the digital twin area and/or provide a prototypical implementation.

**Diffusion.** In the final phase, the results of the previous phases are disseminated. The instruments for dissemination range from scientific publications to practitioner seminars depending on the target group. The diffusion can be coupled with DSR guideline tackling communication of research:

*Communication of research. The research targets a technology-focused as well as business-oriented audience, where it must be properly represented.*

The results of this research are summarized in this dissertation (see Section 4). Furthermore, seven publications are published or accepted for publication in scientific journals or conference proceedings. They tackle technological and/or business issues: For instance, Paper DT-F1 describes the digital twin as a system-of-systems and is more management-oriented. Whereas Paper DT-Sec4 provides a very technical view on digital twin for digital forensics. Most of the papers evolved around the Secure Industrial Semantic Sensor Cloud (SISSeC)<sup>14</sup> project, where they also influence practice. To conclude, this dissertation and its research serve a broad audience.

---

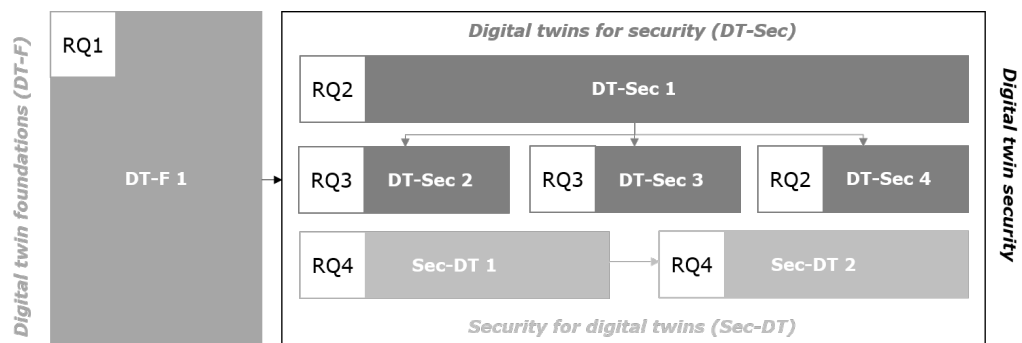
<sup>14</sup><https://www.mobilitylogistics.de/logistics/produktionslogistik/projekt-sissec>

## 4 Results

Following the research method introduced in Section 3, seven research articles answer the research questions of Section 2. Each of these papers has been published in conferences or journals renowned in the area of information systems, and especially in the IT security domain. This way, the research results of this dissertation are communicated to a suitable specialist audience.

Table 1 provides an overview of this dissertation's seven research articles. Thereby, each paper is categorized into one of this dissertation's three main research pillars, which each contain different research questions that are answered (see Section 2). The abbreviations **DT-F**, **DT-Sec** and **Sec-DT** indicate the pillar a paper is assigned to. Each paper is further assigned to a number<sup>15</sup>. Additionally, Table 1 describes each paper by its full citation, its current publication state as well as its type of article. For the latter, (C) is used for submission for publication at a conference, while (J) indicates the submission for publication in a journal. At the time of writing this dissertation all articles have been accepted and published except for paper **DT-Sec3**, which is accepted for publication and will be published shortly. Further details on the submissions, conferences and journals as well as the arrangement of authorship of the respective articles are provided in Part II of this dissertation. In Part II, the full version of the corresponding papers are attached.

The three research pillars of this dissertation form the basis for categorizing the papers and structuring this dissertation. Figure 3 highlights the three pillars in different colors and shows corresponding research papers and questions.



**Figure 3:** Overview of research papers, corresponding research pillars and the research questions answered<sup>16</sup>

To provide a foundation on the topic of this dissertation, the general nature and characteristics of digital twins and their potential for enterprises is investigated in the paper **DT-F 1**. On this basis, the next two research pillars are strengthening the dissertation's focus on security and indicate the two perspectives the combination of digital twins and security can take. In the second pillar (**DT-Sec**), the use of digital twins for security purposes is investigated. Therefore, the article **DT-Sec 1** takes a generic viewpoint, while

<sup>15</sup>Please note that the order by pillar abbreviation or number does not reflect the chronological publication dates of the articles.

<sup>16</sup>The arrows indicate the logical order of the publications, which are not necessarily reflected by their chronological publication dates.

Pillar & No.	Publication	State	Type
DT-F 1	DIETZ, M. AND PERNUL, G. Digital Twin: Empowering Enterprises Towards a System-of-Systems Approach. <i>Business &amp; Information Systems Engineering</i> 62, 2 (2020), 179–184.	published	J
DT-Sec 1	DIETZ, M. AND PERNUL, G. Unleashing the Digital Twin’s Potential for ICS Security. <i>IEEE Security &amp; Privacy</i> 18, 4 (2020), 20–27.	published	J
DT-Sec 2	DIETZ, M., VIELBERTH, M. AND PERNUL, G. Integrating Digital Twin Security Simulations in the Security Operations Center. In <i>Proceedings of the 15th International Conference on Availability, Reliability and Security</i> . ACM, New York, NY, USA (2020), pp. 1–9.	published	C
DT-Sec 3	DIETZ, M., SCHLETTE, D. AND PERNUL, G. Harnessing Digital Twin Security Simulations for systematic Cyber Threat Intelligence. In <i>46th Annual Computers, Software, and Applications Conference</i> . IEEE Computer Society (2022).	accepted	C
DT-Sec 4	DIETZ, M., ENGBRECHT, L. AND PERNUL, G. Enhancing Industrial Control System Forensics Using Replication-based Digital Twins. In <i>Advances in Digital Forensics XVII, IFIP Advances in Information and Communication Technology</i> , vol. 612. Springer, Cham (2021), pp. 21-38.	published	C
Sec-DT 1	DIETZ, M., PUTZ, B. AND PERNUL, G. A Distributed Ledger Approach to Digital Twin Secure Data Sharing. In <i>Data and Applications Security and Privacy XXXIII</i> . Lecture Notes in Computer Science, vol. 11559. Springer, Cham (2019), pp. 281-300.	published	C
Sec-DT 2	PUTZ, B., DIETZ, M., EMPL, P. AND PERNUL, G. EtherTwin: Blockchain-based Secure Digital Twin Information Management. <i>Information Processing &amp; Management</i> 58,1 (2021), 102425.	published	J

**Table 1:** Overview: The seven research papers of this dissertation

the papers **DT-Sec 2-4** dive deeper into individual security topics. This includes the investigation of combining digital twins with security solutions like SOC and SIEM systems (**DT-Sec 1**), its potential for CTI (**DT-Sec 3**) as well as conducting digital forensic analysis with the help of digital twins (**DT-Sec 4**). The third research pillar (**Sec-DT**) tackles security from a different viewpoint: It examines how digital twins can be sufficiently secured. While there are various traditional methods for securing enterprise systems and solutions that can also be applied to the digital twin, the two papers **Sec-DT 1-2** of this dissertation propose a novel approach to provide security to digital twins by using the blockchain technology.

The following Sections 4.1, 4.2 and 4.3 each refer to a research pillar and explain the content of the corresponding publications in detail.

#### 4.1 Digital twin foundations

Due to its novelty and vast application domains, no uniform definition of a digital twin exists to date. Various works on digital twin propose different perspectives – however, mostly from an industrial setting. Nevertheless, digital twins can prove valuable for enterprises in general. Thus, it is essential to study the core characteristics of the digital twin paradigm from an enterprise-centric viewpoint.

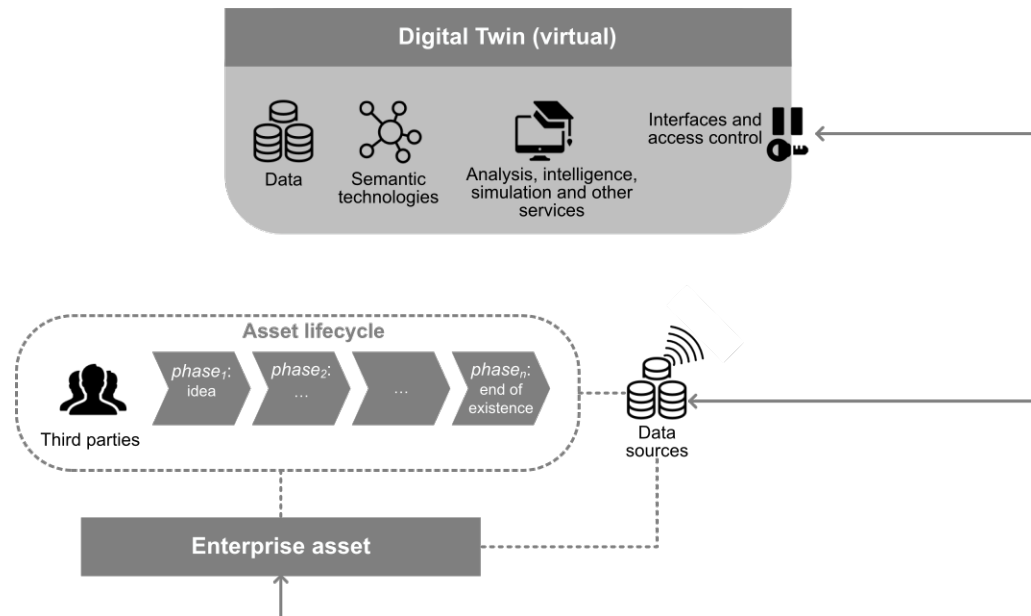
The initial work **DT-F 1** of this dissertation elaborates general parts of the digital twin paradigm, characteristics and current challenges for applying digital twins in enterprises. Moreover, it proposes how digital twins can support the enterprise towards realizing a system-of-systems approach. This work lays the foundations for digital twins from an enterprise perspective, which presents the basis of this dissertation's remaining works.

##### **DT-F 1: Digital Twin: Empowering Enterprises Towards a System-of-Systems Approach**

Paper **DT-F 1** lays the foundations of digital twins in enterprises and proposes how they can further enable a system-of-systems approach. This work answers research question **RQ1** (see Section 2) and defines the notion on digital twins for the subsequent papers on digital twin security.

From an enterprise perspective, it is necessary to efficiently manage all corporate assets (e.g., systems, processes). Increasing complexity of infrastructures and the growth in the amount of data however, hamper the management of enterprise assets. The digital twin paradigm, although mainly discussed in industrial settings, can support this management by virtually representing a corporate asset along its lifecycle. To take advantage of the unique features of digital twins in an enterprise context, it is vital to define the key characteristics, and show how digital twins can create value in enterprises. Moreover, as the concept is still in its infancy, current and future challenges are considered in this paper.

To define the key characteristics of a digital twin, an unstructured literature review is followed. The knowledge gathered is then unified, synthesized and transferred to the enterprise-centric view. Thereby, general parts of the paradigm are identified, i.e. the environment required for the digital twin to establish. On this basis, the characteristics of the digital twin are regarded, i.e. the inner workings of a digital twin. Figure 4 presents the identified general parts of the paradigm: the *digital twin*, the *enterprise asset*, the *asset lifecycle* and the *data sources* that store data about the asset. Furthermore, the characteristics of the digital twins are highlighted. Common features provided by the digital twin are the *data* gathered about its counterpart (the enterprise asset) as well as *semantic technologies* enhancing this data. Moreover, on top of this semantically enhanced data, *analysis, intelligence, simulation and other services* can be built. To



**Figure 4:** The digital twin paradigm [Paper DT-F1]

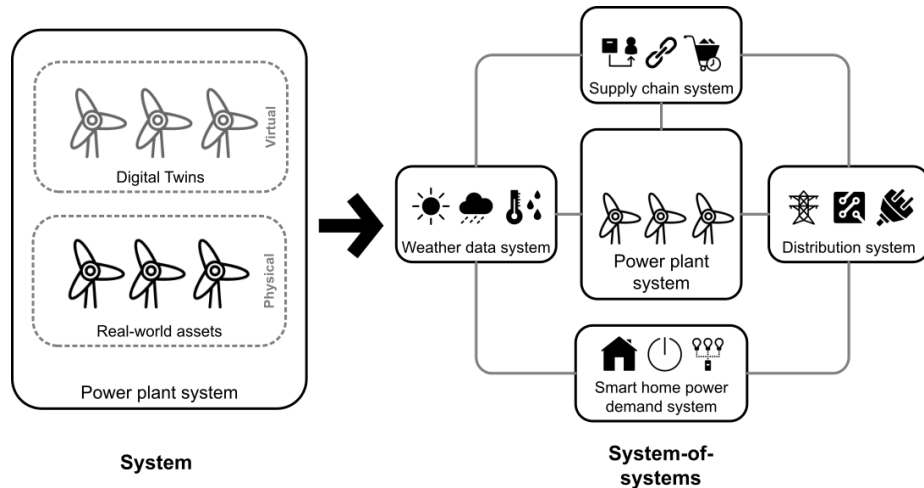
ensure security as well as usability, *interfaces* providing digital twin data are to be secured by *access control* mechanisms.

Based on the established notion of the digital twin paradigm in an enterprise context, potential manifestations of digital twins are proposed. This includes stating the different types, respectively dimensions, of digital twins as well as the proposition of creating a system-of-systems approach with digital twins. The latter is shown in Figure 5.

This paper's idea of using digital twins to create a system-of-systems relies on the research on system-of-systems in general [52]. Beyond simply connecting physical systems, digital twins should be connected as virtual system management tools in order to realize the benefits of the system-of-systems phenomenon and furthermore, to enhance its management. In general, a system is established by connecting related objects [72]. In terms of digital twins, one digital twin might present a system or various digital twins might be connected to establish a "digital twin system". By combining various, previously separated systems, a system-of-systems approach can be established [52]. Owing to the novel combination, synergies can be created with the power to overthrow competitive dynamics [52]. Figure 5 exemplary indicates, how a system-of-systems approach can be created using the digital twin paradigm: It considers a power plant (system), virtually presented by digital twins of the respective windmills. This "digital twin power plant system" can then be connected to other digital twin systems – like the distribution system managing the supply of energy provided by the power plant. Connecting these systems by their digital twins, finally enables a system-of-systems approach.

At last, current and future challenges regarding digital twins are pointed out. These are categorized into *technical efforts* and *corporate challenges*, from which the proposition of *future research* are derived.

Regarding the IEEE technology predictions identified for 2022 [59], digital twins



**Figure 5:** A system-of-systems approach based on [52], realized by networking systems through their digital twins [Paper DT-F1]

will create and interact in a so-called metaverse<sup>17</sup>, which is similar to the proposed system-of-systems approach in this paper. It can thus be concluded that this paper's main idea, which was already published in 2020, is still highly relevant for the future.

#### **Contribution of DT-F:**

This work lays the foundations for digital twins in an enterprise context. It provides the notion of the digital twin paradigm by identifying its general parts and the characteristics of a digital twin. In addition, it proposes how digital twins can empower enterprises by contributing to a system-of-systems approach. The notion of a digital twin presented in this paper enables future research to delve deeper for corporate digital twins.

## **4.2 Digital twins for security**

Based on the foundations created within paper **DT-F**, the following papers tackle the first perspective on digital twin security. This perspective aims at strengthening security by the deployment of digital twins. Thereto, it is necessary to examine how the digital twin can potentially enhance security in general (Paper **DT-Sec 1**). On this basis, the subsequent papers show how digital twins simulations can be used together with corporate strategies like the SOC (Paper **DT-Sec 2**) or CTI (Paper **DT-Sec 3**), and finally, how digital twin replications can be employed in the field of digital forensics (Paper **DT-Sec 4**).

### **DT-Sec 1: Unleashing the Digital Twin's Potential for ICS Security**

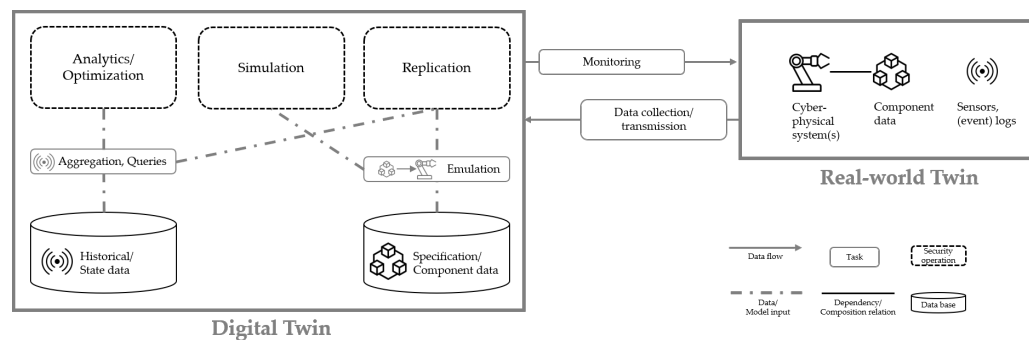
On top of the twin foundations laid out in Paper **DT-F 1**, this Paper **DT-Sec 1** investigates the potential of digital twins for enhancing security. While the subsequent Papers **DT-Sec 2-4** focus on different aspects in this regard, this paper establishes the ground work for the deployment of digital twins for security.

<sup>17</sup>In a metaverse of digital twins, digital twins cooperate and interact with each other instead of their physical counterparts [59].



In terms of security, digital twins are currently almost neglected: It has only been about four years ago since few works addressed this research niche (see Figure 16). Moreover, a generic investigation on how digital twins might benefit corporate security has not been considered yet. Hence, this paper addresses the question how security operations with digital twins are enabled in order to contribute to corporate security (see **RQ2**, Section 2). Thereby, the paper introduces the novel opportunities, entailed by the digital twins' characteristics, which can be exploited for advancing security.

To achieve this goal, the foundations laid in Paper **DT-F 1** are examined in detail and put into a security context. Compounds of the digital twin paradigm (see Figure 4) of Paper **DT-F 1** are further detailed. For instance, from the module "analysis, intelligence, simulation and other services" three operation modes are carved out. Also, the module "data" is further categorized. The exact interplay between the data, the operation modes as well as data production at the real-world twin and its interaction with the digital twin are further presented. Figure 6 shows the results of this study and in an overall model. These findings comply with and complement the few works present on digital twin security (esp. [19, 20]).



**Figure 6:** Digital twin security operation modes (dashed lines) and overall model [Paper DT-Sec1]

The model shows the security operation modes of a digital twin in correspondence with the data and tasks required. It suggests three main operation modes: *Analytics/Optimization*, *Simulation* and *Replication*. Each of these modes rely either on *Historical/State Data*, on *Specification Data* or on both. With the latter, an *Emulation* of the real-world twin can be built, which serves for the two operation modes simulation and replication. *Aggregation, Queries* can be laid upon historical state data, which can then serve the analytics operation mode or for replication. The data flow between the real-world and its digital twin is further detailed into *Data Collection*, indicating the flow from the real-world towards the digital twin, and *Monitoring*, indicating the data flow in the opposite direction. The real-world twin shows which data is produced – the symbols suggesting the corresponding data base in the digital twin.

The paper further outlines, how the security operation modes could have prevented the industrial attacks Stuxnet [37] and Triton [47]. Next to the security operation modes, the second half of the paper proposes how each of these modes can contribute in different areas including *Lifecycle Security*, *Security by Design*, *Digital Forensics* and *Security and*

*Safety*. Finally, the current and future challenges of digital twin security are highlighted.

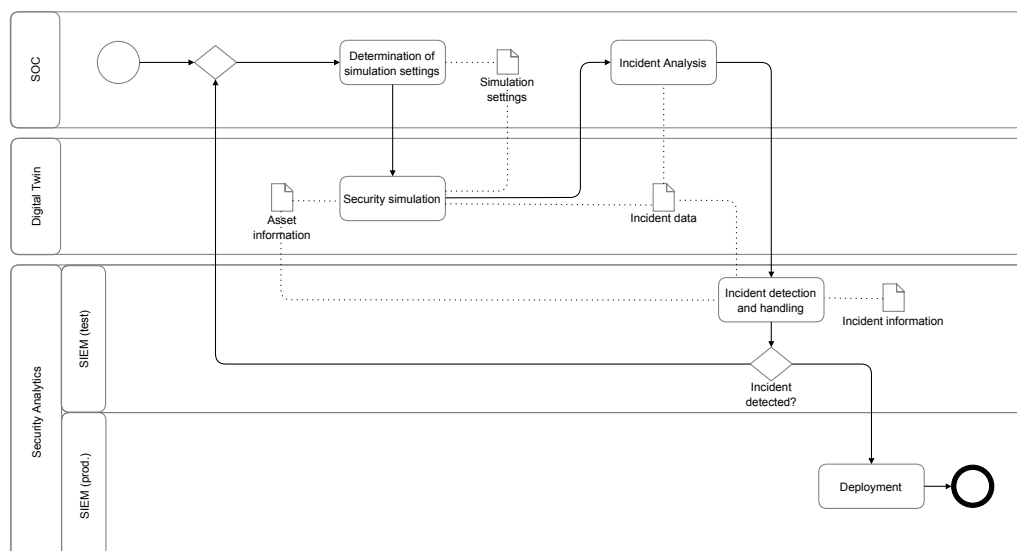
### Contribution of DT-Sec 1:

This work proposes a general view on employing digital twins for security. It presents three main security operation modes that provide potential to enhance corporate security. Different areas of security importance and how these might benefit from the digital twins deployment are further detailed. This paper provides the basis for subsequent works on using digital twins to strengthen security.

## DT-Sec 2: Integrating Digital Twin Security Simulations in the Security Operations Center

Relying on the foundations on corporate digital twins as presented in in **DT-F 1**, Paper **DT-Sec 2** deepens the aspect of integrating simulations for security purposes, which was previously proposed in **DT-Sec 1**. Thereby, it addresses **RQ3** (see Section 2). Next to a conceptual approach, Paper **DT-Sec 2** provides proof-of-concept by implementation.

To maintain security, enterprises currently establish a Security Operations Center (SOC), which is technically supported by Security Information and Event Management (SIEM) systems [69]. Although such systems provide a variety of functionalities, the integration of industrial systems is commonly not realized to date. Digital twins are not only able to digitally represent such industrial systems, they can also contribute to corporate security. Especially, the operation mode *simulation* might provide valuable information (e.g., log data) for detecting attacks, which can be integrated in SIEM systems. This paper is the first to propose an approach for integrating digital twin security simulations into corporate security structures (SOC and SIEM).



**Figure 7:** Process-based security management framework integrating digital twin security simulations [Paper DT-Sec2]

The proposed approach consists of a Business Process Model and Notation (BPMN)<sup>18</sup>-modelled process, which serves as framework for integrating digital twin security simulations into corporate security. Thereto, formal requirements are stated and a use case for demonstration is proposed. A final evaluation of the approach and use case is realized by prototypical implementation.

The process-based security management framework (see Figure 7) consists of three main lanes: the *SOC*, the *Digital Twin* and *Security Analytics*. The latter can be broken down into a test and a productive SIEM. To make use of the digital twin's simulation capabilities for security, the SOC first determines the *Simulation settings*. These and the *Asset information* are subsequently used to build the digital twin security simulation. The output of this activity is *Incident data*, which is analyzed by SOC employees. The incident detection and handling activity is conducted by the test SIEM system, where the data items produced by the digital twin contribute to create *Incident information*. If the determined incident is detected, the detection mechanism (e.g., the SIEM rule) can be deployed in the productive SIEM system. Otherwise, the process has to start anew – with altered simulation settings.

Each of these activities is formally stated to specify the process requirements. In order to evaluate the process-based approach, a use case representing an industrial filling plant is proposed (see Subfigure 8(b)). Furthermore, a mirco-service architecture, consisting of a digital twin simulation of the proposed use case and a SIEM system, is implemented (see Subfigure 8(a)). On the digital twin side, MiniCPS<sup>19</sup> is used for simulation of the industrial filling plant, while Ettercap<sup>20</sup> serves to mimic the attack setting. Filebeat<sup>21</sup> transfers the produced log data towards the SIEM infrastructure. This data is then processed by the SIEM correlation engine DSiem<sup>22</sup>, which builds upon Logstash<sup>23</sup>, Elasticsearch<sup>24</sup> and Kibana<sup>25</sup>. The prototypical implementation is made publicly available at GitHub<sup>26</sup>. Each of the proposed approach's process steps is carefully tackled for demonstration. The security simulation with the digital twin is carried out by conducting a Man-in-the-Middle (MITM) attack, which results in an overflow of the bottles of the filling plant. The log data produced serve for the incident analysis step, and the deduction of the SIEM rules to be implemented. In the evaluation it is shown how each process step can produce the required output data. Finally, it verifies that digital twin security simulations can provide a vital tool for incident detection with SIEM systems in the SOC.

Looking into practice, this paper's central proposition seems to hit the bullseye: As of 2021, today's industrial digital twins are integrated in traditional corporate structures like ERP software (e.g., from SAP) [58]. For security this would mean the integration

---

<sup>18</sup><https://www.bpmn.org/>

<sup>19</sup><https://github.com/scy-phy/minicps>

<sup>20</sup><https://www.ettercap-project.org/>

<sup>21</sup><https://www.elastic.co/de/beats/filebeat>

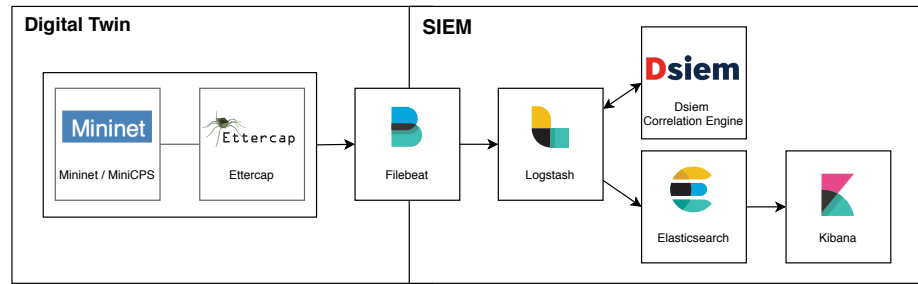
<sup>22</sup><https://www.dsiem.org/>

<sup>23</sup><https://www.elastic.co/de/logstash>

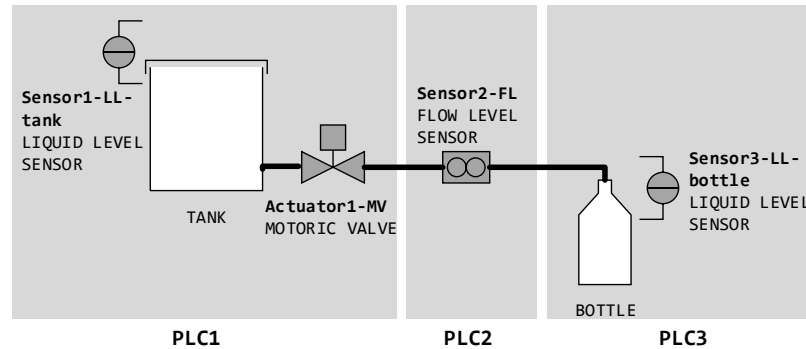
<sup>24</sup><https://www.elastic.co/de/elasticsearch/>

<sup>25</sup><https://www.elastic.co/de/kibana>

<sup>26</sup><https://github.com/FrauThes/DigitalTwin-SIEM-integration>



(a) Implementation: Miro-service architecture



(b) Digital twin use case: Filling plant

**Figure 8:** Implementation of digital twin use case and SIEM infrastructure [Paper DT-Sec2]

into traditional systems like SOC and SIEM – as realized in this work.

### Contribution of DT-Sec 2:

This paper proposes a process-based approach for integrating digital twin security simulations in existing enterprise security structures, namely the SOC and SIEM systems. Instead of focusing on a stand-alone incident detection with the digital twin, the approach is the first to suggest the integration of digital twins in traditional security solutions. The presented approach is evaluated by use case demonstration and prototypical implementation.

### DT-Sec 3: Harnessing Digital Twin Security Simulations for systematic Cyber Threat Intelligence

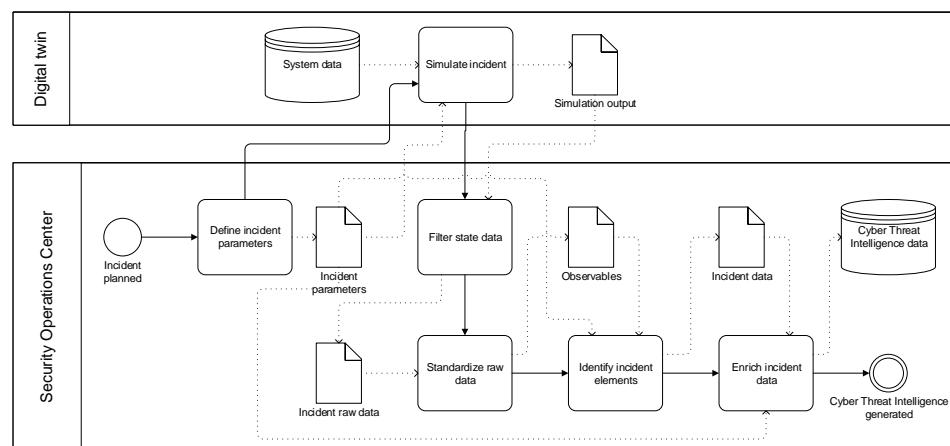
This Paper **DT-Sec 3** investigates how digital twin simulations can support corporate security structures (see **RQ3**, Section 2). It thereby relates to the previously discussed Paper **DT-Sec 2**. It further bases on the Papers **DT-F 1** and **DT-Sec 1**. The focus of this work is set on combining digital twin security simulations and CTI formats.

Security threats and incidents can be mitigated and sometimes even prevented by applying CTI. Representing vital information sources for enterprises to maintain their security, CTI reports contain details about security threats, attacks and incidents. Over the last years and due to the advancement of Industry 4.0, industrial systems have become the target of cyber attacks. This is especially problematic as those systems are designed

to last long-term and have neglected security to date. Thus, CTI referring to industrial systems might help to improve security. The digital twin, which is often used to reflect an industrial asset virtually, provides a valid threat information source as it can simulate attacks in a setting close to the real-world.

To address these issues, this paper proposes a framework using digital twins to generate valuable CTI data. It introduces the activities necessary to achieve a structured CTI report from digital twin security simulations. In order to receive relevant output data for CTI generation from digital twin simulations, it must be ensured that the simulation is as close to the digital twin's counterpart as possible. Thereto, a formal model including elementary definitions is developed. Furthermore, we select the STIX2<sup>27</sup> format to represent and structure CTI. Our framework is evaluated by the implementation of a digital twin use case with a simulated attack. Moreover, utility tools are developed to assist the generation of CTI from the digital twin simulation output. This works results in validating the framework by creating useful CTI reports of the provide use case.

Figure 9 shows the central artifact of this paper: The process-based framework relies on the BPMN modelling language and combines simulations carried out by the *Digital twin* and the CTI generation conducted by the *Security Operations Center*. Besides



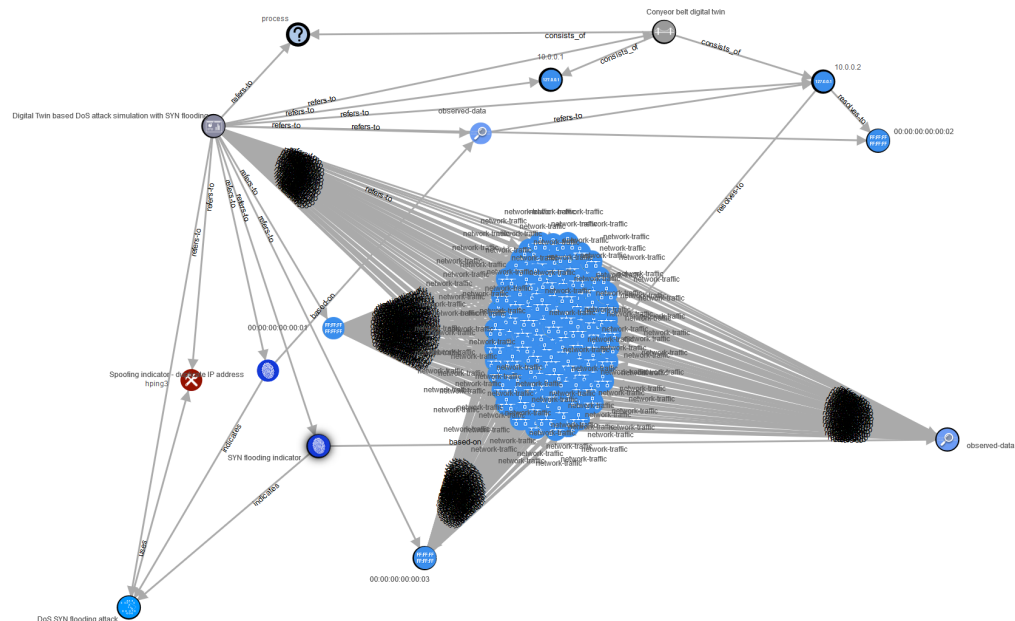
**Figure 9:** BPMN-based process for creating CTI from digital twin security simulations [Paper DT-Sec3]

*System data*, *Incident parameters* defined by the SOC are used to simulate a security incident. This digital twin simulation leads to *Simulation output*. Onward, the SOC conducts filtering of state data to extract *Incident raw data*. The latter is subsequently standardized into *Observables*. Observables and Incident parameters serve to identify elements of the incident (*Incident data*), which can be enriched by the parameters or other sources to create *Cyber Threat Intelligence data* that can form a structured report.

To evaluate the proposed framework, each step of the process is demonstrated. An experimental setup is developed: The use case relies on a digital twin representing a conveyor belt that is the target of a Denial of Service (DOS) attack. The implementation

<sup>27</sup><https://docs.oasis-open.org/cti/stix/v2.1/stix-v2.1.html>

of the digital twin simulation and the attack is made publicly available at GitHub<sup>28</sup>. The implementation relies on the MiniCPS technology<sup>29</sup>, the attacker tool Wireshark<sup>30</sup> and hping3<sup>31</sup>. The simulation output data consists of system log data as well as network



**Figure 10:** Visualized CTI report of a simulated DOS attack on a digital twin conveyor belt [Paper DT-Sec3]

traffic of the attacked system. To assist security analysts within the SOC to implement the framework steps, we provide utility tools for filtering and standardizing simulation output. These are also made publicly available at GitHub<sup>32</sup>. After running the prototypical simulation and using the implemented CTI utility tools, CTI is generated. By following the proposed steps of the process-based framework, a CTI report with relevant information about the use case can be created. The visualized CTI report of the conveyor belt use case is shown in Figure 10. It represents the main elements of the infrastructure and the attack (nodes) as well as their relationships (arrows).

### Contribution of DT-Sec 3:

This work shows how digital twin security simulations can serve to generate CTI. The process-based framework introduces the necessary steps to transform digital twin simulation output into useful CTI. By implementation of a digital twin security use case and CTI generation utility tools, the proposed process steps are validated. To conclude, this work shows how digital twin simulations provide a powerful data source for CTI.

<sup>28</sup><https://github.com/FrauThes/DigitalTwin-ConveyorBelt>

<sup>29</sup><https://github.com/scy-phy/minicps>

<sup>30</sup><https://www.wireshark.org/>

<sup>31</sup><https://tools.kali.org/information-gathering/hping3>

<sup>32</sup><https://github.com/DanielSchlette/CTI-DT-utilities>

#### **DT-Sec 4: Enhancing Industrial Control System Forensics Using Replication-based Digital Twins**

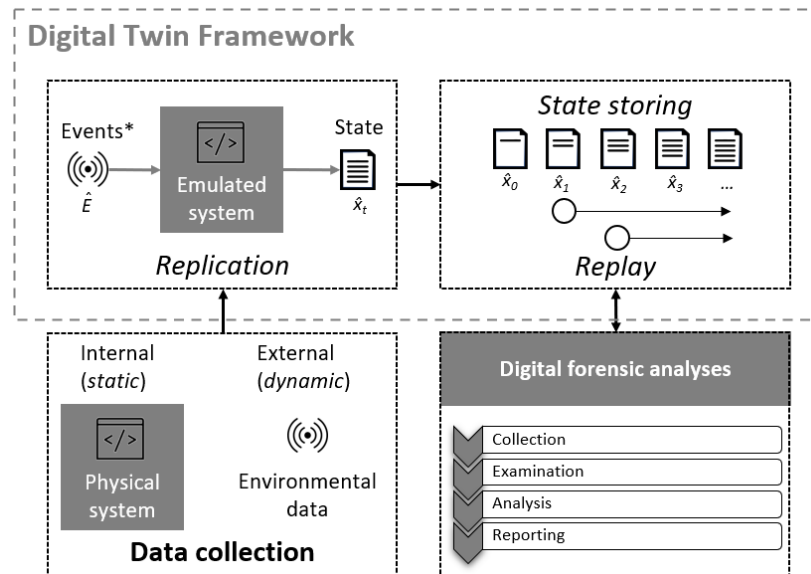
Based on the already established foundations (Papers **DT-F 1** and **DT-Sec 1**), this Paper studies the digital twin operation mode *replication* in detail to further establish foundations on digital twin security operations (see **RQ3**, Section 2). It further provides an approach for applying digital twins in this mode for forensic analyses.

Digital forensics comprises manifold tasks concerning the acquisition, analysis and interpretation of digital evidence after cyber security incidents. This aims at gaining insights to attacks, especially for court room prosecutions. However, since industrial systems are insufficiently secured and increasingly targeted these days, the digital investigation of such devices is problematic due to availability requirements: their continuous operation must be ensured and cannot be disturbed. Moreover, due to the heterogeneous system architectures, it often remains unclear which digital forensic tool is the best fit for extracting digital traces from industrial systems. A digital twin can address these issues as it is able to mirror a real-world system virtually and thereby provides an excellent investigation environment for digital forensics – without hampering the operation of the real-world system.

This work proposes the application of a replication-based digital twin to mirror an industrial system and serve for digital forensic analyses. Two existing works with reference to digital twin replication [19, 24] serve as a baseline for the creation of the five theorems, which are central for enabling the digital twin replication mode. On this basis, a profound conceptual framework for a replication-based analysis is developed (see Figure 11). A final prototypical implementation demonstrates proof of the proposed concept.

The main replication-based concept is illustrated in Figure 11 and consists of the *Data collection* phase, the *Digital Twin Framework* including replication, state storing and replay functionalities and the *Digital forensic analyses* steps.

The data collection task is focused on gathering sufficient detailed data to represent the desired physical system virtually with the digital twin. This includes receiving data from the physical system (internal) as well as from its environment (external). After the collection, this data is incorporated into the digital twin, where the *Replication* further relies upon it. From the internal data of the physical system (mostly static data about configurations etc.), the replication mode can build an emulation, which can behave like the real-world system. In order to mirror the states from the real-world system in the digital twin, it is necessary to induce the same triggers (dynamic Events, e.g., network traffic). The digital twin framework is further enhanced by including storing of the subsequent system states in historical order (*State storing* functionality) and by allowing to replay the sequences of states from specific points in time (*Replay* functionality). Next to the basic replication of its real-world counterpart, the two functionalities especially serve the digital forensic analyses. The digital twin framework can support each of the forensic investigation phases. For instance, the *Analysis* phase might rely on the state



**Figure 11:** Replication-based digital twin concept for digital forensics [Paper DT-Sec4]

storing functionality to investigate the states of the system to check for anomalies based on malicious activities.

An implementation of a replication-based digital twin prototype is carried out to evaluate the proposed concept. Thereto several tools are used to achieve the desired functionalities: In order to induce the real-world network traffic captured in data collection phase to the emulated system, Polymorph<sup>33</sup> is used. The emulation of the system is carried out with OpenPLC<sup>34</sup> that can integrate real-world PLC code and mimic its behavior. To enable state storing and replay, an extended version of a continuous data protection software (SauvegardeEX<sup>35</sup>) is applied.

For the forensic analyses, it is shown that the states resulting from a simple PLC program changing sensor values can be saved and re-analyzed with our prototype. In addition, all changes to files or values can be successfully recorded and analyzed. In the future, this can serve as a baseline for selecting adequate forensic investigation tools.

#### **Contribution of DT-Sec 4:**

In this paper the foundations of replication-based digital twins for a digital forensic investigation are elaborated. Moreover, a sophisticated framework is developed that shows how the replication mode, enhanced with two profound functionalities (state storing and replay), can serve digital forensic analyses. In addition, a prototype is provided to demonstrate proof-of-concept. In summary, the paper concludes that replication-based digital twins can be realized with current technologies and shows that it benefits digital forensic investigations.

<sup>33</sup><https://github.com/shramos/polymorph>

<sup>34</sup><https://www.openplcproject.com/>

<sup>35</sup>[github.com/LudwigEnglbrecht/sauvegardeEX](https://github.com/LudwigEnglbrecht/sauvegardeEX)



### 4.3 Security for digital twins

Based on the foundations created within the first research pillar, the second perspective on digital twin security is addressed in the following papers. The general aim lies in strengthening digital twins to prevent malicious activities. Especially since digital twins contain unified knowledge about an enterprise asset, the securing of this data is of utmost importance. The two articles of this research pillar investigate how digital twin data management can be secured and answer to **RQ4** (see Section 2). Thereby, Paper **Sec-DT 1** identifies the requirements, and proposes a first conceptual approach for secure digital twin data management by applying the blockchain technology. This approach is further developed, implemented and extensively evaluated in Paper **Sec-DT 2**.

#### **Sec-DT 1: A Distributed Ledger Approach to Digital Twin Secure Data Sharing**

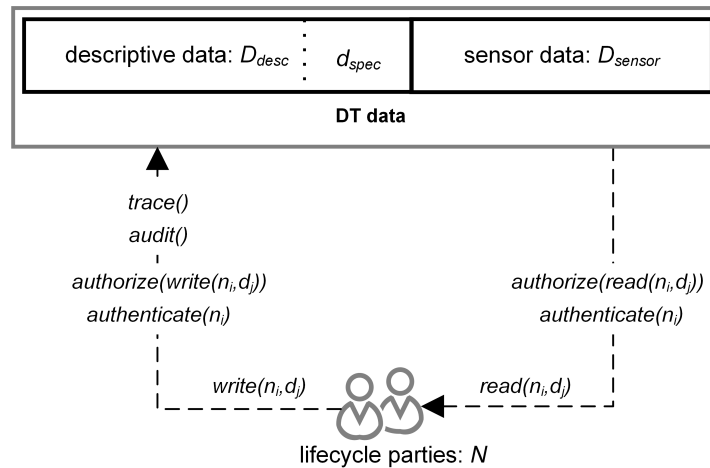
While Paper **DT-F 1** lays the foundations for corporate digital twins, Papers **DT-Sec 1-4** tackle the first angle of digital twin security. This Paper, **Sec-DT 1**, constitutes the starting point for the second angle: It investigates how digital twins, especially the management of their data, can be secured. A powerful enabler for that is the usage of the distributed ledger/blockchain technology.

A main characteristic of digital twins represents storing unified data about its real-world counterpart. Furthermore, it accompanies its real-world asset among its whole lifecycle. During this lifecycle, multiple parties are involved (e.g., manufacturer, owner, maintainer) that deliver and consume data about the asset. More precisely, those parties do not only contribute their data to the digital twin, but they might also obtain certain data from the digital twin (feedback loop). For instance, the asset's manufacturer might require data about maintenance tasks in order to optimize the manufactured product in terms of longevity. Each of the involved lifecycle parties follows its own business goals, which might contradict the goal of another involved party. Hence, trust cannot be established. Instead, sufficient security, especially in terms of confidentiality and integrity, is required. Confidentiality ensures that the parties involved have restricted access to the data elements in the digital twin. Integrity is required as data should not be manipulated and has to be reliable, otherwise it will lead to misinformation and poor decisions. Before the publication of this research paper, no other works have considered these security aspects for digital twin data sharing.

In Paper **Sec-DT 1** the problems stated above are addressed by designing a data exchange of digital twins in a secure way. To this end, a formal basis is established (see Figure 12) and requirements are derived. On this basis, a solution architecture is conceptualized (see Figure 13).

As shown in Figure 12, digital twin data can be categorized into (a) *descriptive data*, which is rather static and describes the asset (e.g., information about the configuration, its composition and network topology), and (b) *sensor data*, which is produced dynamically and informs about the asset's state and environment. A special item of descriptive data is the *asset specification* that provides general information about the asset (e.g.,

via standardized data formats like AutomationML<sup>36</sup> or SysML<sup>37</sup>). The lifecycle parties



**Figure 12:** Control flows for a single digital twin [Paper Sec-DT 1]

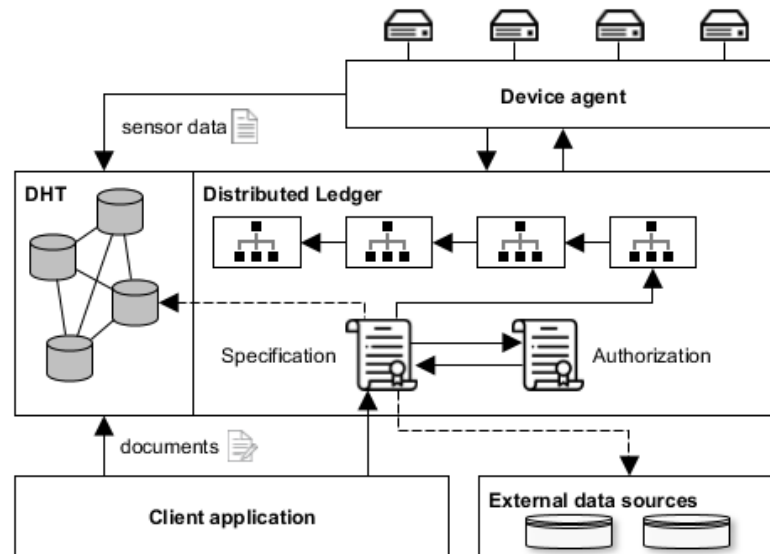
interact with digital twin data by *read* and *write* operations. However, to ensure data integrity and confidentiality, security mechanisms are required: Confidentiality is ensured by installing access control mechanisms (*authentication* and *authorization*). Integrity is achieved by *auditability* and *traceability* mechanisms in write operations: Auditability allows only authorized users to write or update data, while other operations cannot transform data. Traceability enables the chaining of transformations of data items, so that the history of changes of a data element can be traced. In conclusion, five requirements are identified to enable secure digital twin data sharing, including: (1) *multi-party sharing*, (2) *data variety support*, (3) *data velocity support*, (4) *data integrity and confidentiality mechanisms* and (5) *read and write operations*.

Before the solution architecture is developed, the paper discusses if a distributed ledger – respectively a blockchain – approach is suitable. Thereby, it follows the framework by Wüst and Gervais [73] with the result that either public or private permissioned blockchain can be applied. Figure 13 illustrates the conceptualized solution architecture based on the blockchain/distributed ledger technology.

At the core of the architecture lies the *distributed ledger* and a distributed hash table (*DHT*), both managing the digital twin data. Via a *client application*, lifecycle parties can write and read data, whereby they commonly share descriptive data (documents). A *device agent* manages the incorporation of data (especially sensor data) from the digital twin’s physical real-world counterpart. This sensor data and the documents are stored in the DHT. However, to ensure all benefits from the distributed ledger, this data is linked in the ledger. The ledger consists of two smart contracts, the specification and the authorization contract. The authorization contract contains information about the data items and the users that have access to them, and therefore, enables access control. The specification contract is built upon the specification data (see Figure 12) that describes the

<sup>36</sup><https://www.automationml.org/o.red.c/home.html>

<sup>37</sup><https://sysml.org/>



**Figure 13:** Blockchain-based architecture for secure digital twin data sharing [Paper Sec-DT 1]

digital twin’s counterpart. For instance, this description contains the composition of the represented asset, i.e. its subparts, sensors. Whenever document or sensor data is written to the DHT, an entry in the distributed ledger is created. This entry points to the data element it refers to and thus, to the corresponding storage place in the DHT. Furthermore, *external data sources* that contain information about the asset can be integrated by adding a reference to the specification contract.

By means of a theoretical use case, this work further shows how the solution architecture supports the incorporation of a new sensor and the resulting data. Moreover, it shows how all five requirements are met by the proposed approach.

#### **Contribution of Sec-DT 1:**

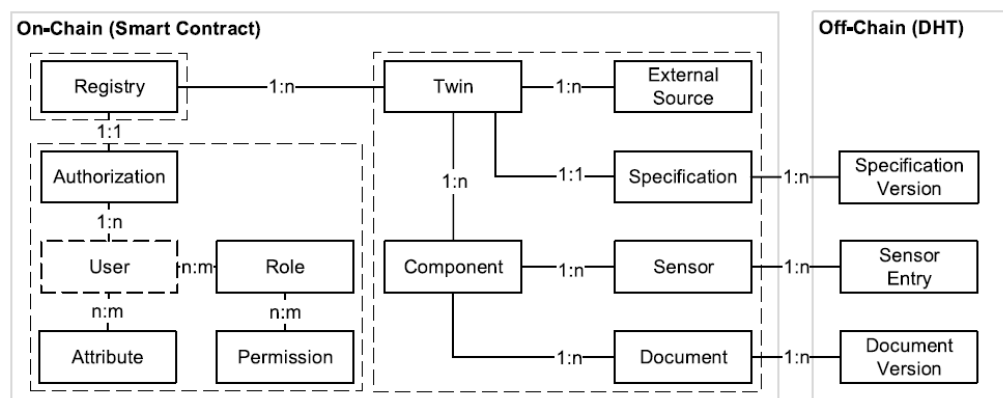
This article proposes the first blockchain-based approach to enable secure digital twin data sharing in research. Thereto, it identifies the requirements and provides a formal basis for the secure sharing. Moreover, this work explains why the blockchain technology is applicable by following a standard method. The result constitutes a solution architecture, which functionality is demonstrated by a theoretical use case. The proposed architecture fulfills the identified requirements.

#### **Sec-DT 2: EtherTwin: Blockchain-based Secure Digital Twin Information Management**

While Paper **Sec-DT 1** proposed a first approach to tackle digital twin data sharing in a secure way, Paper **Sec-DT 2** builds upon the previous work and elaborates a concrete design of a blockchain-based decentralized application (DApp) to securely manage digital twin data. Moreover, the proposed DApp is prototypically implemented (EtherTwin) and extensively evaluated.

The digital twin covers the lifecycle of its physical counterpart. Thereby, accruing data of the lifecycle phases are provided by the different parties involved and incorporated in the digital twin. However, this information sharing and management requires sufficient security as the lifecycle parties contribute to digital twin data, but might not trust each other per default. Thus, the storage and exchange of digital twin data should be secured to be fit for practice. Moreover, the complex ecosystem of interacting lifecycle parties and digital twins requires a new approach. Although digital twins require strong security, fully implemented solutions do not exist yet. To date, prototypical implementations have been either neglected or only partially accomplished.

This work addresses these issues by designing and fully implementing digital twin data management. At the beginning, the digital twin basics in terms of lifecycle phases, potentially involved parties and data are summarized. Following the ten-step decision path by Pedersen et al. [51], it is verified that the blockchain is suitable for secure digital twin data management. Afterwards, the system model is developed. Moreover, the underlying Entity Relationship Model (ERM) of the targeted DApp is introduced (see Figure 14): Three main modules *Registry*, *Authorization* and *Twin Data* are represented

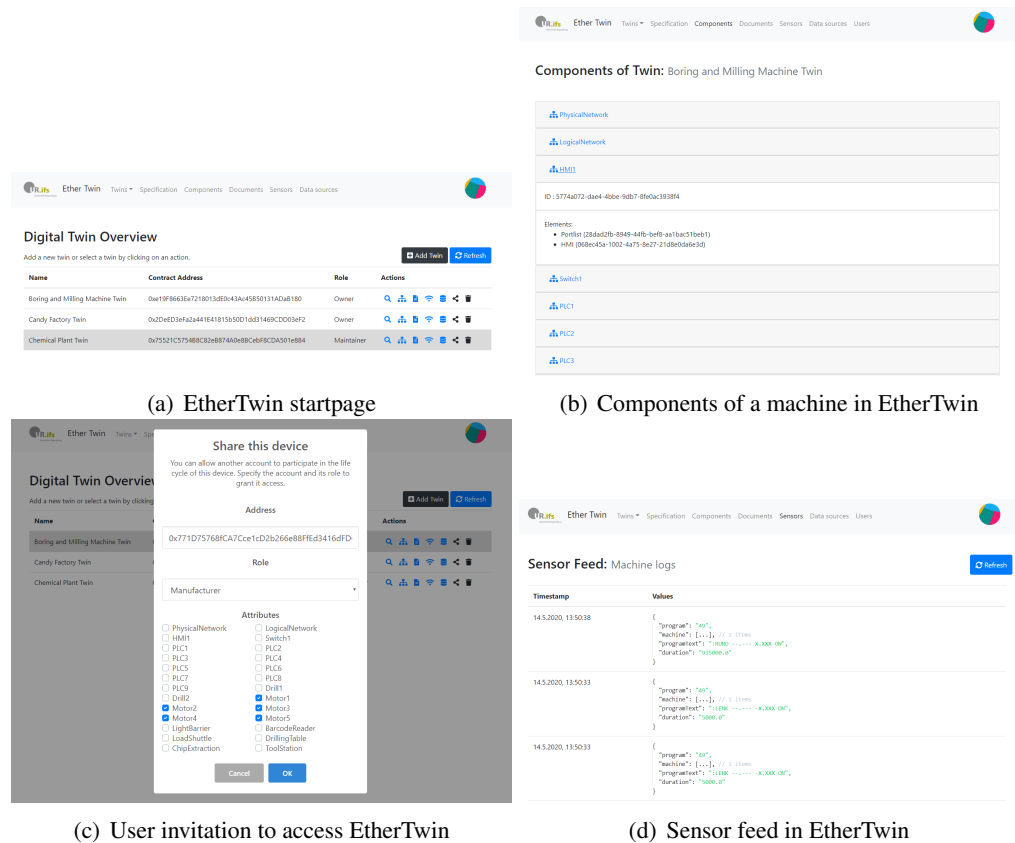


**Figure 14:** Entity Relationship Model of the DApp [Paper Sec-DT 2]

on-chain with the smart contract (indicated by dashed borders in Figure 14). The first manages the incorporated twins, while the Authorization module comprises the proposed access control model, which includes users, roles, attributes and permissions. The digital twin data module relies on the involved digital twin data presented in **Sec-DT 1**. Thereby, each twin has one general specification and consists of several components. Moreover, external data sources might be integrated. The entities off-chain (represented by the DHT) are the specification version, sensor entries and document version. Here, the complete data is contained and a link to the on-chain entities exists to indicate the storage of the data items on the DHT.

On the basis of the system model, technologies for implementation are selected. Furthermore, the design patterns and data flows are specified. Finally, the proposed access control mechanism is designed. The mechanism includes five roles concerning the lifecycle parties, the mapped permissions to documents, sensor data and specification, which are further restricted by the individual attributes. These attributes are generated

individually from the specification data of the system at hand. The proposed system model design is implemented and results in the EtherTwin prototype. This implemented DApp relies on the technologies Ethereum<sup>38</sup> for the blockchain part, Erebos<sup>39</sup> and Swarm<sup>40</sup> for distributed data management and storage, NodeJS<sup>41</sup> for the automatically managing devices (device agent) and a VueJS<sup>42</sup> framework for the user interface. Figure 15 shows parts of the EtherTwin’s user interface. On the *EtherTwin startpage* (see



**Figure 15:** User interface of the EtherTwin DApp [Paper Sec-DT 2]

Subfigure 15(a)) an overview of the available twins, the user’s role and the possible operations is given. Subfigure 15(b) presents the *components* of an exemplary machine with its subcomponents. This view shows the complexity of the presented asset. An example for the implemented fine-grained access control mechanism is given in Subfigure 15(c). It shows the sharing of a twin, respectively the *user invitation* to a digital twin data management space. Thereto, the role of the new user must be specified and the attributes, respective components, the user is involved with. In this example, the invited user represents the manufacturer of all five motors of the machine. An exemplary *sensor feed* is captured in Subfigure 15(d). The EtherTwin DApp provides several more functionalities, including uploading, downloading, deleting and versioning of documents,

<sup>38</sup><https://ethereum.org/en/>

<sup>39</sup><https://erebos.js.org/>

<sup>40</sup><https://swarm.ethereum.org/>

<sup>41</sup><https://nodejs.org/en/>

<sup>42</sup><https://vuejs.org/>

viewing and renewing machine specification, adding external data sources, changing roles and attributes of a user, removing a user, showing user account details, representing an overview of a digital twin’s involved users, the deletion of a twin and the creation of a new twin.

The prototype is further evaluated by the conduction of technical experiments concerning performance (e.g., cost, latency, user quantity). These results show that the transaction costs of a private blockchain might be a more cost-effective alternative as they do not depend on the exchange rates of the public Ethereum mainnet. Moreover, it is proven that an implementation of an industrial use case is feasible. The evaluation shows how the prototype can be used by managing digital twin data concerning a real-world machine. This machine data is gathered from the SISSeC project<sup>43</sup> and describes a boring and milling machine. Finally, semi-structured expert interviews are carried out, including a total of ten experts from six different enterprises. The experts commented on the EtherTwin prototype. Especially the developed user interface appealed to the experts, access control and encryption were also highly valued. In regard to the application of the blockchain technology, the opinion was mixed. However, the experts agreed that the EtherTwin prototype provides a solid basis for digital twin data sharing in practice.

#### **Contribution of Sec-DT 2:**

This paper is the first work that designs and implements a secure data management approach for digital twins. It relies on a decentralized application (DApp) based on the blockchain technology, thereby providing unique features such as data integrity and confidentiality while avoiding a trusted third party. Confidentiality is further ensured by the proposed fine-grained access control model for digital twin data sharing and encryption of digital twin data. The work results in the implemented prototype EtherTwin, which is extensively evaluated by performance measurements, an use case from industry and semi-structured expert interviews. The resulting EtherTwin prototype successfully shows how digital twin data management and lifecycle party involvement can be practically carried out in a secure way.

#### **4.4 Further publications**

In addition to the research papers presented above, further publications have originated during this dissertation. Although they do provide contributions to the topic of this dissertation, they do not present its central research works.

Table 2 gives an overview of these papers. It shows the full citation of the publication, its submission status, and the publication type – (C) for conference, (J) for journal and (W) for workshop.

**Paper 1** describes an anomaly detection technique for detecting outliers in data streams and was presented at the *DEXA 2018* conference [13]. It is loosely related to digital twins as it develops a technique that can be incorporated in digital twins to foster

<sup>43</sup>The ZIM SISSeC project (<https://www.it-logistik-bayern.de/produktionslogistik/projekt-sissec>) is supported under contract by the German Federal Ministry for Economic Affairs and Energy (16KN085725).

No.	Publication	Status	Type
1	DIETZ, M. AND PERNUL, G. Big Log Data Stream Processing: Adapting an Anomaly Detection Technique. In <i>Database and Expert Systems Applications. Lecture Notes in Computer Science</i> , vol. 11030. Springer, Cham (2018), pp. 159-166.	published	C
2	VIELBERTH, M., GLAS, M., DIETZ, M., KARAGIANNIS, S., PERNUL, G. AND MAGKOS, E. A Digital Twin-based Cyber Range for SOC Analysts. In <i>Data and Applications Security and Privacy XXXIV. Lecture Notes in Computer Science</i> , vol. 12840. Springer, Cham (2021), pp. 293-311.	published	C
3	BÖHM, F., DIETZ, M., PREINDL, T., AND PERNUL, G. Augmented Reality and the Digital Twin: State-of-the-Art and Perspectives for Cybersecurity. <i>Journal of Cybersecurity and Privacy I</i> , 3 (2021), 519-538.	published	J
4	DIETZ, M., HAGEMANN, L., VON HORNING, C., AND PERNUL, G. Employing Digital Twins for Security-by-Design System Testing. In <i>ACM Workshop on Secure and Trustworthy Cyber-Physical Systems</i> . ACM, New York, NY, USA (2022).	accepted	W

**Table 2:** Overview of complementary research papers

security. Thereby, real-world data streams from the DINGFEST<sup>44</sup> project are incorporated and analyzed.

In **Paper 2**, the research described in Section 4.2 was enhanced with the concept of cyber ranges: The micro-service architecture comprising a digital twin of a filling plant as well as a SIEM system (see Paper DT-Sec 2) form the data and back-end part of the newly developed cyber range.

**Paper 3** studies the combination of Augmented Reality and digital twins in terms of cyber security in theory, and outlines its potentials.

Finally, **Paper 4** proposes an approach to utilize digital twins for system testing. Thereby, it implements an academic simulation technology to mimic potential attacks. The work aims at rendering systems secure-by-design. In practice, by comparison, the Mindsphere software supports planning and design when there is no physical counterpart yet available [58]. However, its focus is currently not widened to security purposes.

## 4.5 Roots of this thesis and related works

Regarding literature and research on digital twins, which started only about five years ago (see Figure 1), this dissertation is one of the pioneer works tackling the combination of digital twins and security. Below scientific works until December 2021 are summarized and brought into context with the works originating from this dissertation.

<sup>44</sup>The BMBF DINGfest project (<https://dingfest.ur.de>) was supported by the Federal Ministry of Education and Research, Germany.

**Digital twin foundations: DT-F** Table 3 summarizes important works on digital twin foundations, which are detailed in the following.

Year	Publications	Content	Focus	Citations <sup>45</sup>
2017	[49]	review	industry	839
2018	[35]	review	industry	866
2019	[22]	taxonomy	general	56
2019	[66]	review	industry	769
2020	[30]	review	industry	334
2020	[48]	survey	IoT	92
2020	[14]	survey	enterprise	36
2020	[68]	taxonomy	general	48
2021	[42]	review	industry	168

**Table 3:** Publications on digital twin foundations<sup>46</sup>

The first published review on digital twins focuses on CPS in manufacturing and appeared in 2017 [49]. It provides an overview on the different definitions of the digital twin in an industrial context. Another work differentiates the digital twin from the related terms "digital shadow" and "digital model" [35]. Thereby, the underlying categorical literature review centers on digital twins in manufacturing. Tao et al. review the state-of-the-art of industrial digital twins [66]. Jones et al. extract the characteristics of digital twins with the help of a literature review [30]. A recent survey provides foundations of the digital twin in an IoT context [48]. By means of a literature review, Enders and Hossbach extrapolate a taxonomy for digital twin applications [22]. Another recent work tackles digital twin foundations by building a general taxonomy [68]. A relatively new and already well-cited work reviews the digital twin literature in terms of concepts, technologies, and industrial applications dependent on the twin's lifecycle phase [42]. Paper DT-F 1 of this thesis complements this research by studying digital twins for enterprise usage [14].

**Digital twins for security: DT-Sec** Table 4 shows existing publications<sup>47</sup> on this perspective, which are described in the following.

The first two works focusing on digital twin security implement a virtual environment for cyber-physical systems (CPS) [19, 20]. This virtual environment is used to simulate a Man-In-The-Middle (MITM) attack, and provides rules to detect such attacks [20]. In another work, stimuli of real-world events are identified, and the virtual environment reproduces these stimuli in order to mirror the real-world counterpart's network traffic states subsequently [19]. Another work from 2018 tackles the cost-security/fidelity

<sup>46</sup>Number of citations according to Google Scholar on March, 21 2022.

<sup>46</sup>Own work is grayed out. Please note that this overview does not claim completeness. Instead, it shows an excerpt of influential works on digital twin foundations.

<sup>47</sup>Publications on digital twins for security until 2021.

<sup>48</sup>Own work is shown grayed out. Three works on digital twin security from 2018 ([19, 20, 7]) and one work from 2021 ([12]) did not appear in the literature search (see Figure 1) due to not mentioning security in their abstract and have been added as a result of the back- and forward-search [71].



Year	Publication(s)	Content	Focus
2018	[19, 20]	attack detection	industry
2018	[7]	cyber range	industry
2018	[5]	cost-security tradeoff	industry
2019	[21]	state-of-the-art	industry
2019	[36]	security test	robotics
2019	[9]	attack detection	smart grid
2020	[15]	security operations	industry
2020	[2]	security test	smart grid
2020	[17]	integration	industry
2020	[56]	attack detection	smart grid
2020	[26]	security controls	enterprise networks
2021	[28]	attack detection	smart grid
2021	[11]	anomaly detection	smart grid
2021	[70]	cyber range	industry
2021	[61]	security architecture	industry
2021	[60]	security architecture	smart grid
2021	[25]	security test	nuclear power plant
2021	[12]	digital forensics	industry

**Table 4:** Publications on employing the digital twin for security until 2021<sup>48</sup>

tradeoff by proposing a method on how to determine minimal costs for sufficient security [5]. In another work, adequate controls for security are developed with the help of digital twins [26]. Three other works focus on a digital twin to detect attacks or anomalies in grid security [9, 11, 28, 56]. Sellitto et al. propose a digital twin-based security architecture in the smart grid area [60] as well as for industry in general [61]. Regarding security testing with digital twins, works propose concepts for the smart grid area [2], for nuclear power plants [25] as well as a prototypical realization in robotics [36]. A pioneer approach for utilizing the digital twin replication mode to conduct digital forensics is rooted in this dissertation [12]. Instead of focusing on the digital twin as a standalone technology, the twin is integrated with existing enterprise structures to enhance security in another work of this dissertation [17]. A first concept to combine cyber ranges and digital twins in Industry 4.0 is proposed by Becue et al. [7], while we realize an implementation of a digital-twin based cyber range [70]. A state-of-the-art summary on digital twins for industrial security is given in Eckhart et al. [21]. An overview of potential digital twin operation modes and applications for security is given in our first work regarding digital twins for security (Paper DT-Sec 1) [15].

**Security for digital twins: Sec-DT** Regarding security for digital twins, Table 5 summarizes published works<sup>49</sup> in terms of their content and area of focus. The very first publication on security for digital twins [16] originates from this dissertation. It provides pioneer work by proposing the usage of the *blockchain* technology to share and store digital twin data among various parties along the lifecycle of the twin’s real-world counter-

<sup>49</sup>Publications on security for digital twins until 2021.

part. From there on, various works have combined blockchain with digital twins: Kanak et al. propose a model to store digital twin operations on a blockchain [31]. Another blockchain framework intends to store healthcare information of digital twins in smart cities [3]. Also originating from this dissertation, the very first implementation of a blockchain-based digital twin data management solution is presented by Putz et al. [53]. Jiang et al. harness the blockchain for use in digital twin networks [29]. The concept of Digital-Twin-as-a-Service is enabled by a permissioned blockchain in the work of Liao et al. [39]. In order to secure the communication digital twin edge networks, Lu et al. provide a blockchain-empowered federated learning scheme [43]. In another work, the authors apply their concept to digital twin industrial networks [44]. A similar approach is provided for digital twin vehicular edge networks [41]. Also, digital twins in power grids need to be secured: Danilczyk et al. thereto propose to use the blockchain for integer storing of sensor values [10]. Another blockchain framework to strengthen data integrity of digital twins is provided by Dong et al. [18]. Concerning digital twins in manufacturing, a blockchain-based concept and prototypical implementation is provided in another work [63].

Year	Publication(s)	Content	Focus
2019	[16]	blockchain, storage	enterprise
2019	[31]	blockchain, DT ecosystem	general
2020	[24]	synchronization	industry
2020	[64]	cloud, storage	general
2020	[3]	blockchain, storage	smart city
2021	[53]	blockchain, data management	industry
2021	[29]	blockchain, DT ecosystem	general
2021	[39]	blockchain, DT-as-a-service	smart city
2021	[74]	synchronization	AV
2021	[8]	DT ecosystem	general
2021	[1]	safety, security framework	AV
2021	[43, 44]	blockchain	IoT, industry
2021	[75]	privacy, queries	healthcare
2021	[10]	blockchain, storage	smart grid
2021	[32]	security risk, threats	general
2021	[18]	blockchain, storage	general
2021	[23]	encryption	general
2021	[63]	blockchain	manufacturing
2021	[40]	spoofing protection	AV
2021	[41]	blockchain	AV

DT= digital twin; AV = autonomous vehicles

**Table 5:** Publications on providing security for digital twins until 2021<sup>50</sup>

Other works tackle securing digital twins by regarding different aspects. For instance, the impact of cloud security for digital twin data storage is studied [64]. Similarly, securing and privacy-preserving querying the cloud storage of digital twin healthcare

<sup>50</sup>Own work is shown grayed out.

data is tackled in another work [75]. Chen et al. develop an interaction engine to secure digital twin networks [8]. Safety and security issues and potential solutions for vehicular digital twins are identified by Almeaibed et al. [1], while another work studies risks and proposes countermeasures concerning digital twin components in general [32]. Feng et al. propose an encryption mechanism for digital twin networking communication, which is evaluated in a jamming attack [23]. An anti-spoofing method for digital twin vehicular networks is presented by Liu et al. [40]. A first concept to securely *synchronize* digital twins and their counterparts is provided and implemented by Gehrman and Gunnarsson [24]. To ensure secure information synchronization of digital twins in the automotive area, Jing et al. propose an authentication protocol [74].

## 5 Conclusion and Future Work

This dissertation contributes to the state-of-the-art by studying digital twins in respect to security. Thereby, research is conducted threefold: At first, the foundations of digital twins in enterprises are laid. On this basis, two perspectives of digital twin security are studied. On the one hand, the digital twin is a young digitization concept requiring security. On the other hand, it can enhance (corporate) security. At the beginning of this dissertation, no works concerning digital twin security were published. During the course of this dissertation, various scholars tackled this vital research domain. Thus, the research papers originating from this cumulative dissertation represent pioneer works opening up a new and important perspective on digital twins.

Before diving into security, the foundations of digital twins in enterprises are addressed. Here, a framework describing the digital twin paradigm is established. Thereby, several building blocks contributing to the digital twin are identified from literature. Simulation and other analyses (i.e., operation modes) are one vital part. Another element of the digital twin paradigm are the counterpart's lifecycle and the parties involved therein. Both these aspects will play a greater role in the constructive security works. The work on digital twin foundations also proposes that interacting digital twins can create a so-called system-of-systems to empower current enterprises. It is the first work to take such a view and aligns with the recent prediction to establish a metaverse of digital twins in the future [59].

The security part of this dissertation firstly focuses on enhancing security with the help of digital twins. Thereto, three main operation modes are identified: *historical state analysis*, *simulation* and *replication*. Furthermore, various security applications including security testing and digital forensics are proposed. With focus on the second operation mode, security simulations with the digital twin are integrated into existing enterprise structures, precisely into the SOC. This approach ensures the acceptance of digital twin solutions in enterprises, where completely new technologies commonly face resistance. Thus, by using traditional structures and combining them with novel concepts like the digital twin, the chance for adoption is advanced. This aligns with the current integration of digital twins with traditional information systems (e.g., ERP) [58]. To

this end, a prototypical implementation of an attack on a digital twin of a filling plant is developed in this dissertation. The integration of the digital twin simulation with traditional security tools of the SOC proves applicability of the proposed approach. From this follows that digital twin simulations may also be used to provide threat information and generate CTI reports. Based on another industrial use case, it is shown that CTI can be generated from digital twin simulation outputs. Moreover, digital twins provide the advantage of providing insights for digital forensic investigations. Storing the states of machines and devices and replaying past events can be realized with the *replication* mode of digital twins. This can significantly facilitate digital forensic investigations as these functionalities provide ground for testing hypotheses and tools without destroying or tampering with data on the real-world device. A first prototypical approach combines various technologies to realize these functionalities.

The second security perspective focuses on providing security for the digital twin. Scholars consider adequate security a key requirement for digital twins. This dissertation contributes to the state-of-the-art by providing the very first concept of securing digital twin data management (including data storage and exchange) with the distributed ledger technology. More precisely, a blockchain is introduced in combination with a DHT and an advanced access control model. The latter orientates on the involved lifecycle parties and the components (e.g., sensors, actuators and field devices) of the digital twin's counterpart. Furthermore, the very first implementation of blockchain-based digital twins originates from this work. The outcome of the expert evaluation and the performance tests prove that the developed prototype is indeed suited for practice.

Digital twins represent a highly practical topic. In contrast to other research areas (e.g., databases, distributed ledgers), theoretical foundations and methods are rather neglected. To counteract this issue, most of the research papers published in the course of this cumulative dissertation provide a formalism and rely on rigorous methods. Nevertheless, it should be noted that this dissertation still provides propositions for practical use. In terms of evaluation, several limitations present themselves while researching digital twin security. For instance, benchmark studies could not be conducted at this point in time as this dissertation's research papers are among the very first publications on digital twin security. Another issue concerns the real-world counterpart of a digital twin. Generally speaking, it is almost impossible for researchers to study an (industrial) system and create a digital twin thereof. The reasons lie in the costliness and complexity of these systems, which are usually found at a company's industrial site but not in research labs of a university. In this dissertation, this problem is alleviated as far as possible by expert input and data input from research projects with industrial firms (like the SisseC<sup>51</sup> project). Industrial companies may advise researchers on industrial systems, however, they do not share the detailed composition of their systems. Nevertheless, this way a digital twin of a real-world counterpart can be created – without actually having the counterpart present for research. These limiting aspects will become easier in future studies. With time, public data will be provided, research labs will be created, and benchmarks will

<sup>51</sup><https://www.mobilitylogistics.de/logistics/produktionslogistik/projekt-sissec>

surely be developed. Despite tackling enterprises and security, many of the papers in this dissertation focus on industrial enterprises. Since digital twins originate from and dominate in industry, the focus naturally shifted in this direction. Nevertheless, traditional enterprises might use the digital twin as well. Whether they might employ digital twins is dependent on their use case and the balance between the twins' value and their cost. Future research could extrapolate the potential of corporate digital twins and study which incentives may improve digital twin adoption. It would also be interesting to provide a theoretical basis for digital twin fitness for use along with a maturity model and digital twin development stages. Finally, the implementation of a system-of-system or metaverse of digital twins might be an interesting subject for future studies.

## Part II

# Research Papers

## 1 Digital Twin: Empowering Enterprises Towards a System-of-Systems Approach

---

Current status:	Published
Journal:	Business & Information Systems Engineering, Volume 62, April 2020
Date of acceptance:	June 28, 2019
Full citation:	DIETZ, M. AND PERNUL, G. Digital Twin: Empowering Enterprises Towards a System-of-Systems Approach. <i>Business &amp; Information Systems Engineering</i> 62, 2 (2020), 179–184.
Authors contributions:	Marietheres Dietz      90% Günther Pernul      10%

---

**Journal Description:** BISE (Business & Information Systems Engineering) publishes scientific research on the effective and efficient application of information systems. During the last years, the double-blind peer reviewed journal has transformed from the flagship journal of the German information systems community to one of the leading international journals in this field.



## CATCHWORD

# Digital Twin: Empowering Enterprises Towards a System-of-Systems Approach

Marietheres Dietz · Günther Pernul

Received: 8 February 2019 / Accepted: 28 June 2019  
© Springer Fachmedien Wiesbaden GmbH, ein Teil von Springer Nature 2019

**Keywords** Digital twin · Virtualization · System-of-systems · Enterprise asset management · Lifecycle management · Internet of Things

## 1 Introduction

It is common knowledge that the management of enterprise assets (e.g. plants, IT systems, staff and machineries) contributes to value creation in today's organizations. Moreover, efficient asset management significantly enhances corporate performance. The Digital Twin (DT) is an asset's virtual counterpart that enables enterprises to digitally mirror and manage an asset along its lifecycle. This asset can be tangible as well as non-tangible – ranging from turbines to services. In order to represent the asset's life and behavior virtually, a DT incorporates all kinds of data related to the asset and continuously provides the enterprise with information on the asset's condition. In fact, in some asset-centric organizations, especially those with critical infrastructures, losses due to significant downtime of the asset (e.g. a power plant) involve risks beyond the company's financials. Here, a DT can play an important role to mitigate or even avoid these risks by comprehensively informing about the real-world asset's

status, history and its maintenance needs. Moreover, some DTs even provide a direct interaction with the asset.

Although the first vision of a DT dates back more than a decade (Grieves 2002), it has only recently obtained increased research interest within multiple domains. The main reason for this lies in the fact that central technological enablers, such as the Internet of Things (IoT), have only recently reached the maturity to be deployed profitably in economic environments. Within the different interested communities, the term has evolved leading to at least two different viewpoints of a DT nowadays (Negri et al. 2017). The first defines the DT merely as the simulation of the physical asset itself and is mostly used by engineering scholars. However, beyond the scope of simulation, the second perspective refers to a DT as a model which constitutes the basis for simulations, analyses and the like. The latter perspective is currently the most adopted view on DTs and thus the focused viewpoint of this work.

Besides, a multitude of terms exist describing similar phenomena. For instance, the term “product avatar” emerges from product lifecycle management (PLM) research and refers to a concept which is similar to a DT for a product. However, the focus of these works lies on the availability of user-oriented product information in social networks and web pages (Ríos et al. 2015). Moreover, erroneously, the term “digital shadow” is often used interchangeably with the Digital Twin – despite its mainly referring to a digital footprint.

As DTs allow enterprises to gain an in-depth understanding of their assets, corporate optimization and business transformations can benefit from their unique knowledge. We believe that the DT can contribute to value creation in asset-centric companies due to its power to combine previously separated data from different domains

---

Accepted after two revisions by Ulrich Frank.

---

M. Dietz (✉) · G. Pernul  
Department of Information Systems, University of Regensburg,  
Universitätsstrasse 31, 93053 Regensburg, Germany  
e-mail: marietheres.dietz@ur.de

G. Pernul  
e-mail: guenther.pernul@ur.de

Published online: 25 November 2019

Springer

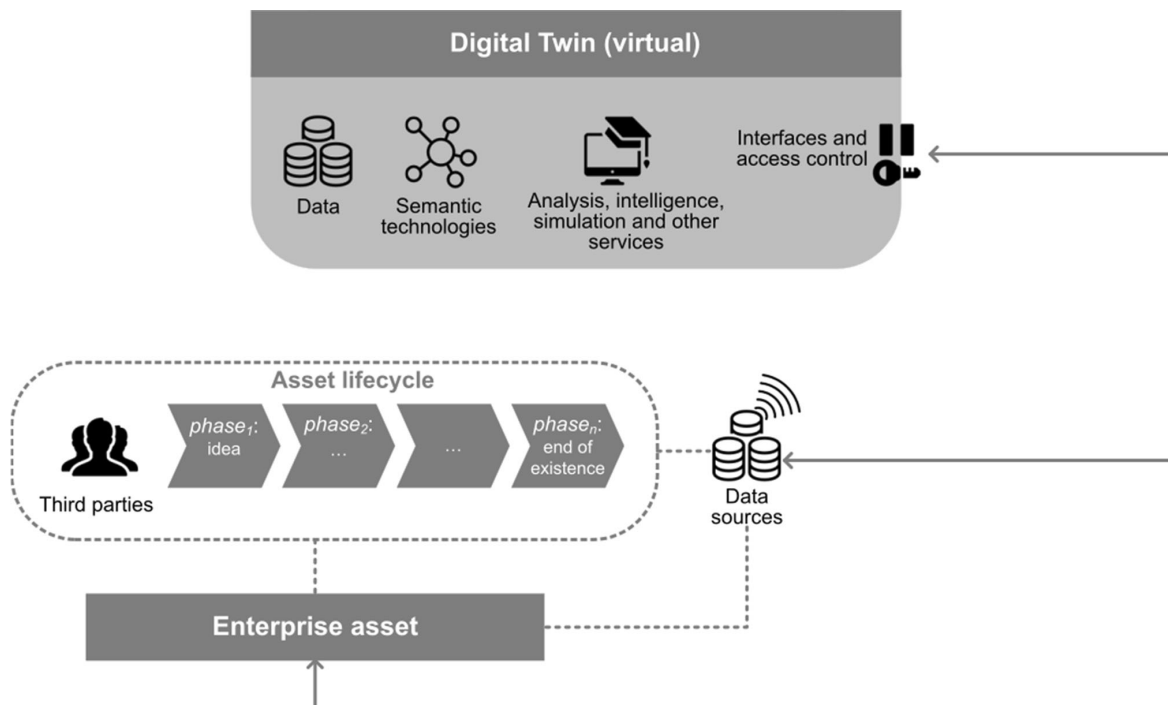
along the asset's lifecycle (see Sects. 2, 3). The DT is, in fact, a proactive digital approach introducing a next step in digitalization. Additionally, BISE scholars recently elected the DT as a central and important technological trend for the community (van der Aalst et al. 2018). However, at present, the concept is still in its infancy as various challenges have yet to be mastered (see Sect. 4) in order to put the DT meaningfully into practice. Hence, to provide profound ideas for practical operation, business and information systems engineering (BISE) scholars need not only to consider this paradigm, but also better understand its key components, its underlying mechanisms and the challenges entailed. Thus, this catchword article aims to pave the road for research by defining key characteristics of a DT, presenting the DT as a paradigm enabling a system-of-systems, demonstrating its potential application fields as well as future research challenges. Although the DT paradigm can also be applied in societal and private contexts, the following primarily concentrates on DT potentials in a business-centric view.

## 2 Key Characteristics of a Digital Twin

To infer the key characteristics of the DT paradigm, a DT's competences are considered. In general, DTs are capable of monitoring, and can be further enhanced with control, over

optimization to autonomy capabilities. To enable these competences, DTs require specific building blocks and need to be integrated into corporate environments. Thus, Fig. 1 illustrates the paradigm of a DT. In the following, we first describe the general parts involved in the DT concept. Afterwards, we proceed to explain the building blocks as the inner part of a DT.

- *Enterprise asset*: An object, subject, system (tangible), or process (non-tangible) relevant to the enterprise, commonly contributing to corporate benefit. Each asset evolves along its lifecycle. A tangible enterprise asset might not physically exist in its first lifecycle phases as well as in the last phase of its lifecycle.
- *Asset lifecycle*: The lifecycle the enterprise asset evolves along. The number of lifecycle phases varies from asset to asset. Thereby, the first phase generally represents the idea of creating the asset and the final phase constitutes the end of its existence. Each lifecycle phase produces relevant information about the asset and therefore, third parties, such as partners, can be involved.
- *Data sources*: The providers of data about the enterprise assets. Data sources can be of any type (e.g. sensors, enterprise systems etc.). Thereby, they can differ between the lifecycle phases, and even belong to an involved third party.



**Fig. 1** The Digital Twin paradigm



- *Digital Twin (virtual)*: The virtual counterpart to the specific enterprise asset. The DT might exist even after the asset has ceased to exist for documentation purposes. It enables data integration and sharing across the lifecycle phases, resulting in a continuous learning process. It may even be embedded in the physical asset itself or deployed on the cloud or an edge computer (Boschert et al. 2018).

DTs are no monolithic data models, but include different aspects of digital representations, functionalities and even interfaces. The following four building blocks make up a DT (see Fig. 1):

- *Data*: Data in various forms (static, dynamic, functional, behavioral, environmental, sensor-based, from handbooks/manuals etc.) relevant to virtually represent the asset.
- *Semantic technologies* (Schroeder et al. 2016; Boschert et al. 2018): Technologies describing the relations between data elements to infer their context, understand their meaning and thus, derive utility for later analyses.
- *Analysis, intelligence, simulation and other services* (Ríos et al. 2015; Boschert et al. 2018): Software enabling search, supporting different analyses, intelligence and other services ranging from simple monitoring to autonomy. Thereby, the extent of the functionalities varies. For monitoring purposes, 3D models (e.g. CAD models) are commonly incorporated and sensor data is often visualized in dashboards for control of the real-world asset. Analyses of past situations, simulations<sup>1</sup> of possible alternatives and further predictive maintenance techniques offer asset optimization. Sometimes self-healing mechanisms are instituted that operate autonomously.
- *Interfaces and access control* (Schroeder et al. 2016; Boschert et al. 2018): Mechanisms to mediate between the virtual and the real world, enable data sharing and synchronization. Especially the bi-directional connection between the real-world asset and its twin provides a novel opportunity not only to report real-world data to the twin, but to send commands from the twin towards its real-world counterpart for its optimization.

However, there are also other perspectives on the paradigm's characteristics. For instance, in the manufacturing area Tao et al. (2018) suggest a DT to be a combination of the product (physical asset), its virtual counterpart and the connected data. In contrast, Uhlemann et al. (2017) see the

<sup>1</sup> To impart a common understanding, a DT can include emulation and simulation functions. For instance, the DT in the work of Eckhart and Ekelhart (2018) emulates a real-world counterpart and provides a simulation environment for testing safety and security rules of the real-world counterpart.

DT as an enabler to realize Cyber Physical System (CPS) that is divided into system layer, data layer, and information and optimization layer. The latter points towards another characterization of the concept that focuses on the data lifecycle: In DTs data is mostly gathered by sensors and the assets they describe (e.g. production systems with CPSs). It is then commonly transferred via IoT technologies and processed by fine-granular and real-time capable simulations, data analytics and the like (Uhlemann et al. 2017). As the DT is an asset-centric concept, the described key characteristics put the emphasis on its representational character, and considering this view, we derived the building blocks as given in Fig. 1.

To conclude, the term “Digital Twin” may be conflated. According to the Oxford dictionary the term “twin” refers to “Something containing or consisting of two matching or corresponding parts” (Oxford University Press 2019). While it fits in terms of the DT being the corresponding virtual counterpart to the real-world asset, the digital part may nevertheless exist after the end of existence of the real-world part and of course, be of different capabilities and granularity by contrast with its real-world counterpart. Therefore, the term “twin” might on the one hand be useful to catch this phenomenon in a metaphorical way. On the other hand, it can be quite misleading as one might expect twins to be rather identical.

### 3 With Digital Twins Towards a System-of-Systems Approach

As stated above, DTs may have different dimensions depending on their context of application. Therefore, they can be divided into multiple perspectives. Note that this does not contradict the universal principle of having only one digital counterpart per asset. At first, a DT exhibiting the characteristics shown in Fig. 1 can be referred to as *basic DT*. *Complex DTs* embed sub-DTs and thus represent a composition consisting of basic DTs or other complex DTs. For instance, consider a simple conveyor belt (*complex DT*) that consists of smaller components (e.g. motor, PLC) that themselves are represented by *basic* or *complex DTs*. Furthermore, DTs may be categorized into different types, i.e. “moving”, “outdoor”, “engine” (*typed DT*). In doing so, a DT referring to a human is referred to as a special type, the *personal DT*. Every *typed DT* can thereby either be a *complex* or *basic DT*. Likewise, similar DT instances might be derived from a *reference-DT* (fleet management<sup>2</sup>). An example is a wind park consisting of

<sup>2</sup> Considering that the asset can carry the digital representation of its type, the establishment of suitable reference systems (models) in which important types are described would be beneficial.

multiple, similar windmills. Here, each windmill produces its own data and thus needs to be monitored by its own DT. However, a *reference-DT* that describes a windmill in general (blueprint) might exist, from which each windmill instance is derived. Expanding beyond a single company's scope towards a notion of a network of companies, the foundation for autonomously *cooperating DTs* that originate from different home domains is laid.

Generally, by connecting an object with further related objects, their functionality is enhanced and a system originates (Wortmann and Flüchter 2015). Furthermore, the consolidation of multiple, previously discontinuous systems offers a system-of-systems approach, which provides the opportunity to fade out company boundaries and promote networking, in order to overthrow competitive dynamics (Porter and Heppelmann 2014). By linking corporate DTs, systems emerge which in turn can be linked to establish a system-of-systems approach. Figure 2 illustrates this idea exemplarily. In this example we consider the system of a power plant, which includes various *complex DTs* (e.g. windmills) and their sub-DTs (e.g. wind turbines). Moreover, this system is connected to other systems, such as to the supply chain system delivering materials and the like, or to the power distribution system, containing assets like transmission towers. By combing these systems through their DTs, a system-of-systems approach is realized. On a more detailed level, DTs or their sub-DTs can be connected in a spatio-temporal manner (Canedo 2016) to indicate a real-world connection. For instance, consider a car being filled up at a gas station, where the DT of the car is virtually connected with the fuel dispenser of the gas station. As soon as the car is filled up, the relation disappears. In a nutshell, DTs manage assets along their lifecycle and thus support the management of the system-of-systems containing these. It is further vital to highlight the relations between the assets, which boost efficiency as optimization can be performed globally at a system-of-systems level. Thus, the DT constitutes more than just a novel technology – it may become a real game changer.

Certainly, the DT paradigm applies in multiple domains. The industrial domain, including Industrial IoT, smart factories and Industry 4.0, is not only a very suitable area for the DT concept, but also the most advanced domain regarding its realization. For instance, General Electric (GE) already counts about 551,000 DTs referring to products, part of products, processes and systems in late 2017 (Saracco 2018). It also offers the “world's first digital wind farm”<sup>3</sup> including DT technology. Another industrial example is Tesla, which applies the DT paradigm to its

cars: every car reports its experience on a daily basis, which further serves simulation in the DT to detect anomalies and propose corrective measures (Saracco 2018). Also, commerce is an area where DTs can achieve efficiency gains. The digitalization of a shop floor, for example, allows the enterprise to manage sub-parts (e.g. shelves) in correspondence with the global system of the shop floor and even the interdependence to other systems such as the supply chain. Moreover, the real-world counterpart must not necessarily be tangible, also processes can be assets monitored by DTs such as shown in Meroni and Plebani (2018). Furthermore, DT implementation can foster social and governance areas equipped with IoT, such as smart cities (Saracco 2018). Besides, DTs can also support individuals (*personal DT*). For instance, the trend “quantified self”<sup>4</sup> refers to individuals gathering as much quantitative information about themselves and their daily lives as possible – mostly with the help of technology. A real-world example is the “most connected human on earth”<sup>5</sup>, for whom up to 700 sensors daily gather his physical condition, activities etc. A DT could combine these data sources and manage the physical condition virtually. This leads also towards medical healthcare, where currently a lot of effort is put into DTs. For instance, a major German engineering company established a blueprint of a DT representing a human heart by using MR, ECG measures and massive data sets next to complex algorithms to enable planning, prediction of recovery of medical procedures.<sup>6</sup> In the future, the creation of an individual's own digital heart could be based upon that blueprint. Moreover, in the long-term view, the complete human body with its organs, its inner cellular constitution etc. will be represented by a DT. Hereby, again *reference-DTs* can deliver the general structure of the organs, from where the individual's organs are derived, enhanced with data of the individual (e.g. ECG, medication plan, diseases) and further customized. In addition to the representation of individuals, DTs can also support the optimization of medical and organizational workflows, or even entire hospitals. Combining these two applications towards a system-of-systems approach, a variety of scenarios can be simulated and their effects on process efficiency can be presented without great expense. Especially the U.S. market is adopting the technology, e.g. through introduction of DT technology in a new facility at the Medical University of South Carolina (MUSC) Shawn Jenkins Children's Hospital and at the Pearl Tourville Women's Pavilion to predict workflows, propose optimizations, to forecast the impact of changes and health

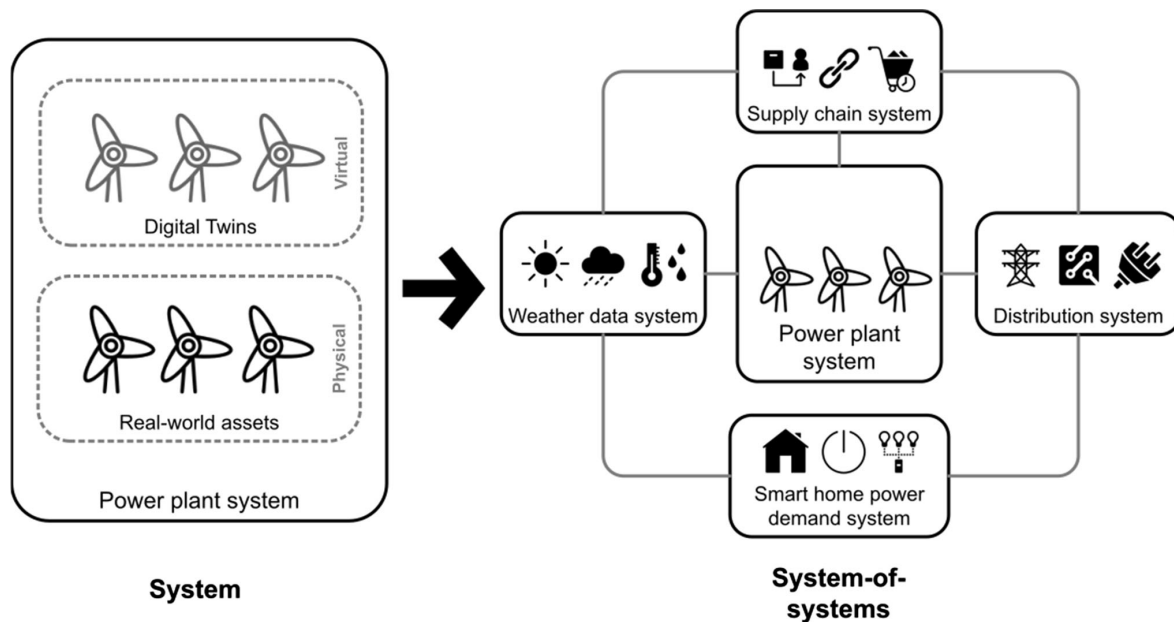
<sup>4</sup> <http://www.quantifiedself.com/>.

<sup>5</sup> <http://www.chrisdancy.com>.

<sup>6</sup> <https://www.siemens-healthineers.com/de/press-room/press-videos/im-20181204001shs.html>.

<sup>3</sup> <https://www.ge.com/renewableenergy/wind-energy/technology/digital-wind-farm>.

M. Dietz and G. Pernul: Digital Twin: Empowering Enterprises Towards a System-of-Systems Approach, Bus Inf Syst Eng



**Fig. 2** A system-of-systems approach by networking systems and their DTs (based on Porter and Heppelmann 2014)

innovations.<sup>7</sup> Clearly, domains of all verticals and among different use cases can benefit from the DT.

#### 4 Challenges and Future Directions

DTs create substantial economic and organizational power for firms. However, while the paradigm broadens its application domains in the business world, challenges from technical as well as business perspective emerge. The following illustration of challenges will hopefully initiate fruitful discussions among BISE researchers in order to solve some fundamental problems in the area of DTs. Thereby, the issues are categorized in technical and corporate challenges. Finally, prospective research areas are listed.

DTs require substantial *technical efforts* from firms. Especially the incorporation of heterogeneous data requires further progress. For instance, the standardization in data acquisition needs to be accelerated (Uhlemann et al. 2017). Also, issues with real-time data have to be focused on. For instance, the manual acquisition of real-time data has to be automated to enable the collection of full historical data instead of snapshots (Uhlemann et al. 2017). Another great technical burden is the current decentralization. Interfaces, connections and the like need to be developed to enable a

more holistic approach. Moreover, security concerns have to be addressed. This includes version management and compatibility checks of the DT versions to ensure data integrity as well as access management to allow third parties to access (parts of) the virtual twin.

At different levels, *corporate challenges* remain. While operational challenges mostly overlap with the technical challenges, the DT potentially implies major economic and organizational transformations at a strategic level. The strategic decision of DT implementation focuses on but is not limited to the degree of asset-centralization of an organization. The study of Klostermeier et al. (2018) shows that the DT paradigm qualifies for different enterprises in multiple application scenarios to various implementation degrees. Hence, the degree of reflecting the real-world asset varies depending on the use case. Consequently, this poses yet another corporate challenge in terms of depth of detail and granularity. Moreover, the study indicates that at present, the term “Digital Twin” can be misleading as it is applied for slightly different phenomena in different areas (Klostermeier et al. 2018). Entangled with the complexity of DTs, a further strategic decision bases on the tradeoff between the cost of an asset versus the cost of the DT and its data granularity (e.g. by including sensors). However, with the general decrease in storage and sensor cost, the implementation of DTs tends to become more attractive. Moreover, the emergence of novel business models, including such where the organization does not possess the asset but provides the service of

<sup>7</sup> <https://www.wallstreet-online.de/nachricht/10822802-musc-and-siemens-healthineers-form-strategic-partnership-to-disrupt-and-reshape-health-care-delivery/all>.

establishing a DT, will entail major privacy issues – including the question of ownership of the virtual counterpart.

Although first assumptions about the effects of DTs exist, *future research* should analyze the concrete technical and corporate implications. One impact of DTs concerns the improvement in interconnectivity, especially within supply chains, where research should not be limited to a single firm's perspective but rather take a system-of-systems approach. Another potential for enterprises is the closure of feedback loops along and the coherent linkage between the lifecycle phases of the asset. Here, the gain of new insights can lead to a win-win situation among the involved parties. Currently, very little is known about potential data-driven business models, their power of digital disruptiveness and pricing strategies for DT services. Research can contribute by identifying the strategic role of DTs for firms and its position in digital transformation, not only from the perspectives of companies owning the DTs but also from third parties contributing to the resources. Nevertheless, first evidence indicates that DTs can give a cutting edge for next-generation virtual asset management.

**Acknowledgements** The authors would like to thank Ulrich Frank for his fruitful comments and suggestions on earlier versions of this manuscript. We also thank our reviewers for their helpful remarks.

## References

- Boschert S, Heinrich C, Rosen R (2018) Next generation digital twin. In: Horvath I, Suarez Rivero J, Hernandez Castellano P (eds) Proceedings of TCME, pp 209–217
- Canedo A (2016) Industrial IoT lifecycle via digital twins. In: 2016 international conference on hardware/software codesign and system synthesis (CODES+ISSS), pp 1–1
- Eckhart M, Ekelhart A (2018) Towards security-aware virtual environments for digital twins. In: Proceedings of the 4th ACM workshop on cyber-physical system security, ACM, New York, NY, USA, CPSS'18, pp 61–72
- Grievies M (2002) Conceptual ideal for PLM. Presentation for the Product Lifecycle Management (PLM) center, University of Michigan
- Klostermeier R, Haag S, Benlian A (2018) Digitale Zwillinge—Eine explorative Fallstudie zur Untersuchung von Geschäftsmodellen. HMD Praxis der Wirtschaftsinformatik 55(2):297–311
- Meroni G, Plebani P (2018) Combining artifact-driven monitoring with blockchain: analysis and solutions. In: Matulevičius R, Dijkman R (eds) Advanced information systems engineering workshops, CAiSE. Springer, Cham, pp 103–114
- Negri E, Fumagalli L, Macchi M (2017) A review of the roles of digital twin in cps-based production systems. Procedia Manufacturing 11:939–948. In: 27th international conference on flexible automation and intelligent manufacturing, FAIM2017, 27–30 June 2017, Modena, Italy
- Oxford University Press (2019) Definition of *twin* in English. <https://en.oxforddictionaries.com/definition/twin>. Accessed 13 May 2019
- Porter M, Heppelmann J (2014) How smart, connected products are transforming competition. Harv Bus Rev 92:11–64
- Ríos J, Hernández JC, Oliva M, Mas F (2015) Product avatar as digital counterpart of a physical individual product: literature review and implications in an aircraft. In: Curran R, Wognum N, Borsato M (eds) Transdisciplinary lifecycle analysis of systems. Advances in transdisciplinary engineering, vol 2. IOS Press, Amsterdam, Netherlands, pp 657–666. <https://doi.org/10.3233/978-1-61499-544-9-657>
- Saracco R (2018) The rise of digital twins. <http://sites.ieee.org/futuredirections/2018/01/16/the-rise-of-digital-twins/>. Accessed 13 May 2018
- Schroeder G, Steinmetz C, Pereira C, Espindola D (2016) Digital twin data modeling with automationml and a communication methodology for data exchange. IFAC-PapersOnLine 49(30):12–17. In: 4th IFAC symposium on telematics applications TA 2016
- Tao F, Cheng J, Qi Q, Zhang M, Zhang H, Sui F (2018) Digital twin-driven product design, manufacturing and service with big data. Int J Adv Manuf Technol 94(9):3563–3576
- Uhlemann T, Lehmann C, Steinhilper R (2017) The digital twin: realizing the cyber-physical production system for industry 4.0. Procedia CIRP 61:335–340. In: The 24th CIRP conference on life cycle engineering
- van der Aalst WMP, Becker J, Bichler M, Buhl HU, Dibbern J, Frank U, Hasenkamp U, Heinzl A, Hinz O, Hui KL, Jarke M, Karagiannis D, Kliewer N, König W, Mendling J, Mertens P, Rossi M, Voss S, Weinhardt C, Winter R, Zdravkovic J (2018) Views on the past, present, and future of business and information systems engineering. Bus Inf Syst Eng 60(6):443–477
- Wortmann F, Flüchter K (2015) Internet of things: technology and value added. Bus Inf Syst Eng 57(3):221–224

---

## 2 Unleashing the Digital Twin's Potential for ICS Security

---

Current status:	Published
Journal:	IEEE Security & Privacy, Volume 18, July-August 2020
Date of acceptance:	December 19, 2019
Full citation:	DIETZ, M. AND PERNUL, G. Unleashing the Digital Twin's Potential for ICS Security. <i>IEEE Security &amp; Privacy</i> 18, 4 (2020), 20–27.
Authors contributions:	Marietheres Dietz      90% Günther Pernul      10%

---

**Journal Description:** IEEE Security & Privacy aims at stimulating and tracking advances in security, privacy, and dependability regarding information systems. The peer reviewed articles feature both, a practical and research perspective, in the field of security and privacy. The magazine targets a broad cross-section of the professional community — ranging from scholars to industry practitioners.

Utility of Synthetic Data ■ A Cybersecurity Terminarch ■ The Security Midlife Crisis

IEEE

# SECURITY & PRIVACY

BUILDING DEPENDABILITY, RELIABILITY, AND TRUST

## Unleashing the Digital Twin's Potential for ICS Security



Reliability Society

IEEE COMPUTER SOCIETY

IEEE

July/August 2020  
Vol. 18, No. 4

# Unleashing the Digital Twin's Potential for ICS Security

Marietheres Dietz and Günther Pernul | University of Regensburg

**With the advent of Industry 4.0 (I4.0), traditional operational-technology and IT systems converge, entailing novel attack vectors. However, I4.0 technologies might also contribute to industrial security. Digital Twins are virtual entities of physical industrial systems that offer opportunities for security, such as simulation and replication of system behavior.**

**T**raditional industrial infrastructures include operational technology (OT) such as supervisory control and data acquisition (SCADA) systems, human-machine interfaces (HMIs), programmable logic controllers (PLCs), and other field devices. These systems can also be characterized as industrial control systems (ICSs), as they are able to command processes within industrial environments. With the advent of the Industry 4.0 (I4.0) concept, the traditional systems are increasingly integrated with general-purpose IT systems, including enterprise-resource planning systems, mostly by establishing connections via Ethernet, TCP/IP, and wireless technologies, such as Bluetooth. Whereas integrating IT and OT systems certainly benefits communication and the creation of collaborative business processes, it introduces novel attack vectors into industrial ecosystems.

It was the beginning of a new era when the industrial world became aware of the Stuxnet worm, whose primary target was the Natanz nuclear-enrichment lab in Iran.<sup>1</sup> The lab's OT systems were maintained by IT-engineering systems. By using known Windows exploits, attackers infected the IT systems and their maintenance OT devices. Lateral movement

via (virtual) private networks became possible, and additional OT systems (especially PLCs and SCADA systems) were targeted.<sup>1</sup> However, unlike previous malware, only OT systems that had specific configurations to manage certain industrial processes were attacked.<sup>1,2</sup> In particular, Stuxnet targeted frequency-converter drives (PLCs) that controlled the speed of centrifuges: If the drives were connected, the malicious code changed the centrifuges' speed to harmful values. Moreover, those changes were concealed by yet another code. All of these aspects indicated a deep knowledge of system behavior and targeted exploitation, which represented a novel kind of security threat, now commonly known as the *advanced persistent threat (APT)*.<sup>3</sup>

In 2017, another industrial APT was discovered: Triton. This one targeted industrial safety systems at a petrochemical plant in Saudi Arabia and, undesignedly, caused plant-process shutdowns. The attacker remained in the target network for approximately one year before entering the safety system's engineering workstation. By analyzing the workstation's communication (IT) with its safety systems, the attacker gained knowledge of the hardware controllers (OT).<sup>4</sup>

These examples show that advanced ICS attacks generally leverage traditional IT systems to reach their final targets. Standards and guidelines to reinforce industrial ecosystems' security [for example,

Digital Object Identifier 10.1109/MSEC.2019.2961650  
Date of current version: 22 January 2020

International Electrotechnical Commission (IEC) 62443 and National Institute of Standards and Technology SP 800-82] include measures that encompass security testing and monitoring as well as attack detection. Moreover, many works address various fields of I4.0 security (for example, Kargl et al.<sup>5</sup>). However, industrial security measures can also benefit from I4.0 concepts,<sup>6</sup> one of which involves creating a virtual entity that represents IT/OT components: the digital twin (DT).

In general, a DT refers to a digital representation of any real-world counterpart during its whole lifecycle. At most times, however, it represents an industrial enterprise asset. The core building blocks of the DT paradigm are asset-specific data items, often enhanced with semantic technologies, as well as analysis/simulation environments to digitally explore the real-world asset.<sup>7</sup> One of the DT's main purposes is to manage the real-world asset, with a scope ranging from simple monitoring to autonomy. DTs are meant to be applied in asset-centric organizations, especially those with critical infrastructures where losses due to significant downtime (for example, at a power plant) may involve risks beyond a company's financials. DTs can play an important role in mitigating and avoiding these risks by providing comprehensive information about the asset's status, history, and maintenance needs. Some DTs even provide direct interaction with the asset.

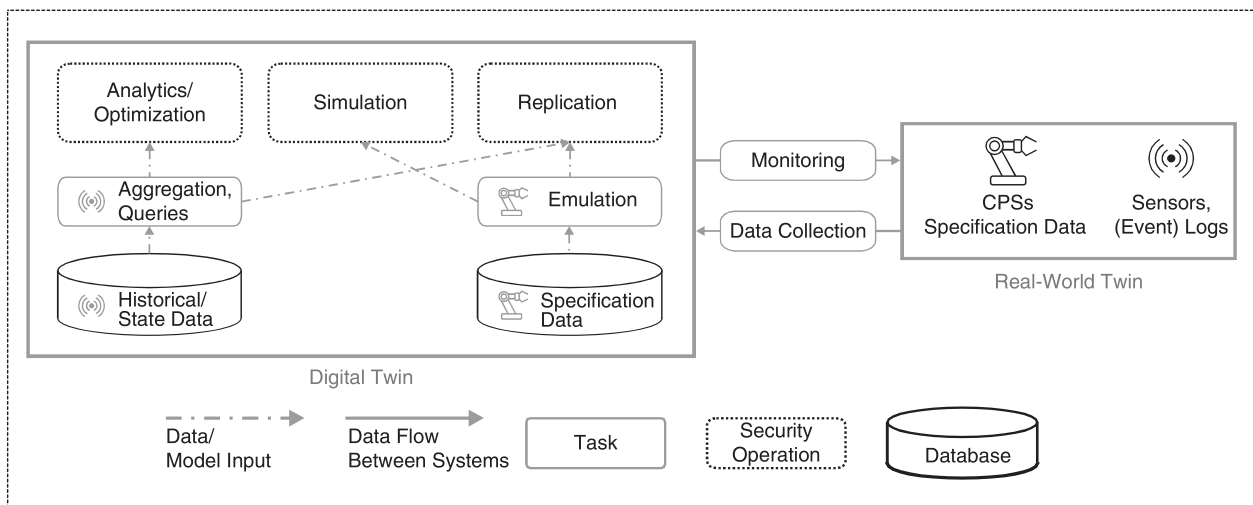
In the Stuxnet example, the frequency-converter drives could have had DTs monitoring their states and sending commands to maintain the necessary centrifuge speed as soon as changes occurred. In terms of security, operations can be carried out in the virtual, isolated environment of the DT without risking negative impact on its real-world counterpart. The direct

connection between the twins further supports an immediate reaction to security incidents.

**Security-Operation Modes**

Most of the time, a security system operates in one mode. For instance, security-information and event-management systems analyze current event logs that are retrieved from the investigated systems (historical data analytics). In contrast, DTs provide the opportunity to operate in more than one mode; namely, the historical data analytics/optimization, simulation, and replication. Figure 1 depicts the modes and simplistically illustrates the data that are used as the basis for the operations. (Note that we do not claim to have included every such mode. However, we are confident that we have chosen the most relevant ones, as the majority of the DT data-processing capabilities can be assigned to them.)

In Figure 1, the real-world twin's state data are captured by the device's sensors or (event) logs and submitted to the DT during the data-collection task. The state data are added to the historical/state database where they can be aggregated and used for ad hoc querying. The specification database includes information about the real-world twin's features, such as the cyberphysical system (CPS) and its composition, including its subparts and program logic. Such IT/OT-system specification is often represented in Automation Markup Language (AML) and used to emulate the real-world twin. The specification database can include other knowledge, including security and safety rules that apply mostly to the program logic. Such rules can contain thresholds and consistency checks and are usually defined in advance.<sup>8</sup> For instance, safety rules might impose upper- and lower-speed thresholds on a motor.<sup>8</sup>



**Figure 1.** The security operations (dashed lines) in a DT.



An example of a security rule involves a speed-variable consistency check of an HMI interface and PLC to make sure that the number entered in the HMI is transferred to the PLC without risking a man-in-the-middle (MITM) attack.<sup>8</sup> Connected to its real-world counterpart, the DT can immediately release commands to monitor its double's current state, if necessary.

### Historical Data Analytics and Optimization

In addition to the on-the-fly analysis based on the asset's state data, the current and stored historical data comprise the input for various analytical purposes and optimizations; among them, behavioral analytics and predictive maintenance.<sup>7</sup> The outcome of those operations can be used to optimize the real-world counterpart; for example, by issuing commands to alter its state.

Analytics and optimization rely on techniques including machine learning, artificial intelligence, and the like. The historical data can serve as baseline and training data for normal instances. Behavioral analytics, for instance, can then "learn" the behavior of the asset's environment, which, in turn, serves as valuable input for detecting anomalies and optimizing processes. Another analytics technique involves the detection of changes and trends through time. Also, detecting events and patterns is an essential historical data analytics technique. These methods enable typical I4.0 use cases, such as predicting maintenance needs, anticipating manufacturing throughput, and forecasting system health.

Most of these analyses can be harnessed for security purposes, which are increasingly deployed. For instance, by monitoring the state through time, outliers and changes that could be the result of malicious activities can be detected. Analyzing the relationships between variables through time enables the detection of possible safety/security-rule violations. The state-of-the-art industrial-ecosystem security includes only a few components that search for anomalies in the behavior of I4.0-essential protocols, such as the OPC Unified Architecture. The analytics mode of a DT can help to remedy I4.0 security issues.

In the Stuxnet example, outliers in the network traffic between the WinCC SCADA system and PLC controllers would have indicated the attack. The DT's analytics mode would have shown anomalous traffic between the infected SCADA system and controllers that requested data beyond the legitimate PLC blocks during a 5-s time span.<sup>1</sup> Similarly, Triton caused unusual traffic across its victim's network. DT analyses of the network traffic during the attack would have indicated inbound and outbound connections to and from nonstandard ranges as well as anomalous secure-shell sessions.<sup>4</sup>

### Simulation

In contrast to the simple analyses of historical data, simulations of a specified asset present a variety of novel security opportunities. Simulations differ from analyses and optimizations because they are established on a model of a real-world asset, which, in turn, is based on specification data. Simulations are run with user-specified settings and parameters. After the evaluation of the simulation outcome, the real-world counterpart can be adapted to the desired setting by transmitting the respective commands and parameters.

Simulations engender the following generic advantages: First, test repetitions are realized by resetting the model and rerunning the simulation. Second, in contrast to tests on real-world assets, the simulation time interval is compressed. Third, simulations reveal a system's behavior through a broad range of specified conditions, which supports the comprehension of emergent behavior. Despite all of these benefits, simulations are mostly neglected in current security analyses.

In terms of security monitoring and operations, these opportunities can be harnessed for a variety of purposes and during the whole asset's lifecycle. For instance, security testers can provide feedback about shortcomings and potential vulnerabilities during the design phase. The setup can also be tested against specified rules for security and safety to detect misconfigurations. New components can be modeled in the virtual environment and tested for security without influencing the production system. Additionally, to cope with missing information, the simulation parameters can be modified. This approach limits the damage to the physical asset and the costs for prototyping and certainly introduces security by design (a concept that will be addressed in the "Security by Design" section). More importantly, security testers can induce security-relevant conditions, such as MITM and insider attacks, and examine the simulated environment.

### Replication

The replication mode resembles simulation in the way that it operates on the specification data to emulate the real-world twin. Additionally, stimuli from the real world are induced by incorporating the current state data and reproducing the events from the physical environment. Although the replication mode comes close to historical data analytics, it differs in two ways. The input knowledge (for example, events and resulting activities) must be known to produce the same stimuli in the digital world, whereas historical data analytics needs a large enough data baseline to train effectively. The specification data play a greater role in the replication mode. This provides the benefits of instantly recognizing affected components, tracing

and reproducing activities, and providing customized countermeasures.

This mode is essential for properly detecting threats and attacks. Comparing the state of the DT's replication mode against the real-world state should deliver the same results (consistency) since the physical environment is mirrored to its virtual counterpart by its specification and current events. Consequently, identified divergences represent a possible attack or failure. For instance, hashing the blocks of the real-world PLC should produce hash values that are identical to those of the DT.

In the Stuxnet example, a divergence in the SCADA/PLC network traffic between the original set and its virtual replica could have revealed the attack. The package count in the replication mode would have been drastically different from the real state, indicating unusual behavior. Moreover, the attack could have been revealed only on the basis of the PLC state versus the state of its digital counterpart. If the PLC operates under certain configurations, the Stuxnet code will be injected into the controller.<sup>2</sup> The Stuxnet infection of certain PLC organization blocks (OBs) is based on increasing the size of the original block. Although the original code of the concerned blocks (that is, OB1 and OB35) remains, the malicious Stuxnet code is prepended.<sup>2</sup>

In conclusion, an examination of the OB1 and OB35 blocks of an infected PLC would result in an unusual block size as well as an unknown code segment at the beginning. In terms of the DT application, predefined rules for protecting the integrity of the PLC code blocks could be stated in the PLC's specification (see Figure 2); this would generate an alert as soon as a block size was altered. Figure 2 presents a pseudocode for detecting mismatching block sizes. The regular sizes of each PLC code block (*regular\_plc\_blocksizes*) as well as the current block sizes (*current\_plc\_blocksizes*) are retrieved. The current block size is compared to the regular one, which produces an alert in case of a mismatch.

In the Triton example, a divergence between the DT and real-world traffic would have indicated unusual behavior as well. Moreover, differences in the file number between the real-world entity and its virtual twin would have alerted the victim. Although Triton's malicious files were named like legitimate ones (for example, "Kb77846376.exe" to mimic Microsoft update files),<sup>4</sup> the compromise would be uncovered by comparing these hashes with the known legitimate ones.

### From Differences to Synergies

Table 1 highlights the differences between the operation modes in terms of the database, techniques, and applications and further introduces the resulting (dis)advantages. It shows that the combination of these modes covers a broad range of security analyses.

The interplay of these operation modes empowers users when they apply the DT for security purposes. For instance, the data outcomes of an attack that is induced during the simulation mode can serve as comparison input for the replication mode to detect the attack in the real world. An example would be to simulate the Stuxnet MITM attack on a virtual OT system that subsequently produces certain data outputs. These outputs can be compared with the real-world twin's current state data to detect whether such an attack was happening. Additionally, simulations can detect system weaknesses, such as missing security rules that can be formulated, added to the specification database, tested, and transmitted to the real-world system (ad hoc monitoring).

The historical data analytics mode can provide information about future system states through its trend analyses. Those predictions can serve as input parameters for simulations. Moreover, the historical data analytics mode supports event identification and provides information about incidents and their occurrences (pattern detection). This serves as the basis for determining the stimuli input for the replication mode.

### Related Concepts

Since some ICS security approaches are conceptually similar, we highlight the commonalities and differences between them and the DT security concept. Testbeds are commonly employed to evaluate the functionality of critical infrastructures and assess the security level by, for example, spotting vulnerabilities and conducting security attacks. However, in contrast to DTs and their operation modes, testbeds usually combine virtual and physical components. For instance, the Electrical Power and Intelligent Control testbed<sup>9</sup> provides an active, realistic environment for researchers to explore a system's vulnerability to attacks and the effectiveness of attack-detection algorithms. Another example is the Secure Water Treatment (SWaT) testbed,<sup>10</sup> which provides an ICS testbed to conduct security research and training. Purely virtual testbeds are even closer to DTs; however, they are commonly deployed during an

```

GET regular_plc_blocksizes AS ARRAY
GET current_plc_blocksizes AS ARRAY
FOREACH index IN current_plc_blocksizes,
    regular_plc_blocksizes
    IF current_plc_blocksize.index !=
        regular_plc_blocksize.index
        ALERT
    ENDIF
ENDFOREACH

```

Figure 2. The security rule that protects the integrity of the PLC blocks based their size.

## DIGITAL TWIN

Table 1. A comparison of the DT security-operation modes.

Operation mode	Advantage	Disadvantage	Aspect	Manifestation
Analytics/ optimization	<ul style="list-style-type: none"> <li>AS-IS state analysis</li> <li>Alerts for current security incidents</li> <li>Broad user base through the prominence of the techniques</li> </ul>	<ul style="list-style-type: none"> <li>Temporal dependencies</li> <li>Analysis of potential future conditions not possible</li> <li>Database size influences the functionality of most of the algorithms</li> </ul>	Database	State data
			Techniques	Statistical analyses, machine learning, and so on Data queries
			Applications	Network-traffic analysis, outlier detection, and so forth
Simulation	<ul style="list-style-type: none"> <li>Time independence</li> <li>Reproducibility</li> <li>Repeatability</li> <li>Analysis of potential future conditions</li> <li>Security-by-design support</li> </ul>	<ul style="list-style-type: none"> <li>Hypothetical conditions</li> <li>AS-IS state not known</li> <li>Isolated view of system</li> <li>Complexity requires professional users</li> </ul>	Database	Specification data
			Techniques	Emulation, simulation
			Applications	Vulnerability analysis, system-security testing
Replication	<ul style="list-style-type: none"> <li>AS-IS state analysis</li> <li>Alerts for current security incidents</li> <li>Digital tracing of real-world stimuli/events</li> </ul>	<ul style="list-style-type: none"> <li>Temporal dependencies</li> <li>Stimuli/events to be known in advance</li> <li>Complexity requires professional users</li> </ul>	Database	State data, specification data
			Techniques	Emulation, stimuli reproduction, differential algorithms
			Applications	Attack and threat detection

industrial asset's design and planning, while the DT covers the whole asset lifecycle.

Cyberranges are virtual environments that provide training opportunities, including tools that help to improve the security, stability, and performance of infrastructures and systems. For example, SWaT Security Showdown<sup>11</sup> gamifies the testbed<sup>10</sup> to provide an enhanced training environment. It targets academia as well as industry ICS security professionals and was designed as a capture-the-flag event. In contrast to cyberranges, DTs focus on the asset, its operation, and its security. However, they might play a role for education and cyber-range challenges in the future.

Deception technologies (for example, honeypots) are effective countermeasures against cyberattacks and APTs. They aim to draw attackers away from real systems, usually by simulating the systems' behavior. While DTs certainly have a different aim than deception technologies, their simulation mode may provide a helpful basis for system simulation with honeypots.

### ICS Lifecycle Security

In common practice, an industrial asset's lifecycle phases are handled separately by the participating parties. This leads to duplicate data and information islands and results in wasted resources and inefficient data sharing. The DT paradigm tackles this issue by anchoring all of the lifecycle phases and arranging information

continuously across the asset's life. The DT evolves with its real-world twin and incorporates the retrievable and relevant information during the lifecycle. In terms of security, this aspect has advantages including the seamless inclusion of protection measures throughout the entire asset's lifespan, beginning with the planning and design phases. (Note that lifecycle security in the scope of this article is tackled via the asset's lifecycle since DTs are asset centric, not by the lifecycle of security incidents, such as kill/attack chains, which are attack centric.)

A unified model of asset lifecycle phases does not exist since the phases strongly depend upon their assets, which can manifest as products, processes, or systems. The most common phases, in practice, are the idea, planning, and design (early phases); operation (medium phases); and demolition and end of existence (final phases). In the following, we discuss two aspects that focus on an asset's lifecycle: security by design and digital forensics. We describe how the DT can contribute to security during an asset's early phases as well as during its later lifecycle phases (operation and end of existence).

### Security by Design

Within numerous IT/OT systems, operational functionality outweighs security. Therefore, security is, if addressed at all, often added during later lifecycle phases (for example, the asset's operation) instead of being considered beforehand. The security-by-design principle

introduces safety mechanisms before an asset's implementation, ideally during the design phase.<sup>12</sup> The lifecycle aspect of a DT complies with the security-by-design principle in the following way: Because a DT aims to incorporate every lifecycle phase, it originates during an asset's planning and design, thereby offering the opportunity to focus on security before operation. This could include stating security and safety rules as input knowledge (specification data; see Figure 1) to ensure the presence of mechanisms to avoid program-logic misuse and attacks. Another way of introducing security by design through DTs includes simulating a prospective environment to detect vulnerabilities, which supports the creation of a security-aware asset.

Introducing security by design helps to lower security and incident-response costs. Since safety mechanisms are installed at the beginning, later lifecycle phases are less prone to security incidents. This reduces the expense of maintaining security, as opposed to spending vast amounts to recover from an attack.

### DTs for Digital Forensics?

Most security requirements should be considered during an asset's design. Nevertheless, each subsequent lifecycle phase brings its own security demands. While the real-world twin must not exist during the early phases, it will presumably operate and evolve during later ones. As well as implementing-by-design, the DT paradigm provides an opportunity to focus on security during later lifecycle phases. For instance, adding novel functionalities during the operation phase of an OT system might require the formulation of new security rules for connecting IT devices to avoid misuse and attacks.

When attacks occur during a real-world asset's operation phase, digital forensics can play a decisive role. The primary digital-forensics advantage that DTs present concerns the fact that the real-world counterpart is not modified by attack investigations. Hence, the digital evidence in the real-world counterpart is not contaminated. DTs support the conservation of evidence and presumably lead to higher-quality digital forensics.

Digital forensics can be divided into live and post-mortem categories. Live forensics are performed during an attack to provide a clear overview of what is happening. The replication mode of the DT can help to detect the attack by mirroring the current events. Live forensics also target the attack's point of origin and subsequent lateral movement. The DT's historical database contains the states as they were before and during the attack. Historical data analytics can identify the time of the attack, while, combined with simulation, the replication mode can replay the malicious activities. If the stimuli are not clear, the simulation mode can replicate different versions of the attack by trying different input

parameters until an output is produced that mirrors the one in the historical database. These modes enable the back-tracing of different ways to enter a system and the subsequent attack dissemination, which is vital for live and postmortem forensics.

A minor limitation is that the DT data might be illegitimate evidence. However, the DT helps to select the procedure for finding traces of the attack. Moreover, the DT may store legacy data (in its historical data; see Figure 1) that might have been deleted from the real-world counterpart. That information may contribute to post-mortem digital-forensic investigations. DTs can be kept for documentation purposes after the real-world twin's demolition, providing relevant data for postmortem digital forensics that continue long after an attack.

### ICS Security and Safety

IT and OT are not the only elements that converge in ICSs. Security and safety are also incorporated to mitigate the outcomes of malfunctioning components (weaknesses) and malicious activities (threats). While safety and security techniques may be applied separately, their merger will inevitably be called for during later lifecycle phases. The effort to align security and safety solutions increases considerably during later lifecycle stages.<sup>13</sup> DTs, with their focus on an asset's lifecycle, may help by addressing the security and safety aspects.

Generally, the works on ICS security and safety can be categorized as combined security and safety approaches, security-informed safety approaches, and safety-informed security approaches.<sup>13</sup> A combined approach consists of the six-step model and information-flow diagrams.<sup>14</sup> First, it models the system's functions and structure. Afterward, the system's corresponding information flows are specified. The subsequent steps identify potential failures and weaknesses, safety countermeasures, possible attacks, and security counteractions.

This example reveals the closeness of the DT operations, especially the simulation mode. The system's structure and functions can be gathered from the specification/component data, which helps to build the emulation. System failures, the exploitation of weaknesses, and attacks can be simulated and further explored, and potential countermeasures can be tested.

### Challenges Ahead

With every novel concept there are challenges to solve. Because the DT paradigm is in its infancy, various obstacles need to be addressed before its effective and efficient operation for security purposes can take place. The major issues are data-related, which influences security operations. For instance, the control logic (of a PLC, for example) provides important security

information. If it is missing, has poor detail perception, or is of bad quality, wrong real-world emulations might occur, strongly impacting security operations.

Although the ICS was identified in 2014, its proprietary nature and a lack of standards represent difficulties.<sup>5</sup> However, the increasing adoption of AML, as specified in IEC standard 62714, moves in the right direction. There are works that aim for the automated creation of DTs by building a software-defined network (SDN) from the AML description of CPSs.<sup>8</sup> Moreover, ICS can be found in various industries, including manufacturing, energy, and transportation. This introduces the question of whether differences remain in the DT concepts between these domains. Currently, it seems that the DT paradigm can be generalized, as various industrial sectors tackle the technology.

Another issue arises with the simulation and replication setup. It requires experience with the applied tools and expert knowledge of the data that are used to establish the model. Otherwise, missing or wrong influence factors or improper coarsening can lead to misleading results on which security decisions and actions are based. Knowledge from related works, such as SDNs, testbeds, and cyberranges, should be incorporated. For instance, the works on cyberranges could provide valuable input to train security specialists (for example, pen testers) and institutions (such as security operations centers) about ICS by the security operations of DTs. The works demonstrate a variety of ways to gamify the security practice by, for instance, providing capture-the-flag challenges to identify and defend attacks in the virtual ICS.<sup>11</sup>

To detect and respond to APTs, future research needs to develop innovative techniques that enhance the traditional approaches.<sup>3</sup> Consider anomaly identification in a network topology that has varying criticality. Combining graph theory and structural controllability concepts can provide a dynamic decision mechanism that detects the topological abnormality.<sup>3</sup> DTs may provide relevant feedback and uncover improvement possibilities: By incorporating the techniques in the simulation environment and launching an APT attack, the effectiveness of these novel approaches can be tested.

ICS interconnection remains an open issue. CPSs are physically or digitally connected and often depend on each other.<sup>9</sup> Thus, one CPS's security incident or weakness may impact the security of another. Future work should focus on these interconnections, especially to facilitate studying the effect of multiple, simultaneous attacks.<sup>10</sup> Whether this remains an issue for DTs depends on the asset a DT represents. If the CPS components within the twin's counterpart are interconnected, the digital or physical connections can be modeled within the DT. If, however, ICSs that are represented by different DTs are interconnected, the question

arises: How can dependencies between separate DTs be managed? To find a solution, works on developing a system-of-systems could provide a starting point.<sup>7</sup>

Covering an asset's whole service span should include all of the lifecycle parties that can provide information about it. However, this multiparty use of a DT can impact the information consistency. For example, the specification of a rotor could include a maximum allowable speed, whereas the specification of the PLC could predefine rules for protecting the integrity of the firmware. The multiparty use of a DT also affects the main security principles, confidentiality, integrity, and availability as well as access-control issues. A DT must ensure that only authorized parties can access, read, and write certain resources. Integrity requires mechanisms that provide data and model consistency among the different domains and lifecycle parties. The availability of a DT is a security requirement that can be tackled by a trusted third party or by applying a decentralized approach, as discussed in Dietz et al.<sup>15</sup>

**A**lthough a number of security mechanisms exist for industrial ecosystems (for example, firewalls and air gaps for network segregation), they are usually not sufficient to reach a proper security level. By compromising the Windows engineering systems of contractors that had physical access to the targeted systems, Stuxnet vividly demonstrated how to overcome air gaps.<sup>1</sup> Mechanisms beyond conventional information security are typically required to achieve the desired security level. The DT incorporates knowledge from multiple domains that provide the information-security view and a perspective of the control systems and their subparts and program logic. Through the operation modes and their interplay that the DT provides, opportunities to detect novel industrial attack vectors emerge. Moreover, commands from the DT to its real-world counterpart enable immediate defense action.

More importantly, OT devices feature similar configurations and are usually supplied by a limited number of vendors. The Stuxnet example showed that on the ICS level, the abused exploits represent legitimate product features. Therefore, the current approach to secure industrial assets and especially OT systems starts with the identification of physical vulnerabilities and protection systems, which provide safety but generally fall short of security.<sup>1</sup> Consequentially, future research should address the combination of DT data with known IT/OT-system vulnerabilities to detect weaknesses. DTs include configuration knowledge (such as system and software versions) that could be exploited to search for physical vulnerabilities and the

presence of protection systems that are prone to security exploits.

Future studies should address the current challenges in DT research, as outlined in the previous section. While the growing implementation of sensors in I4.0 generally leads to more information and solves a lot of data-related issues, awareness and training for identifying the relevant data remains an important requirement. Researchers and practitioners should work hand in hand to develop standards for I4.0 and I4.0 security beyond what is currently available. ■

## References

1. R. Langner, "To kill a centrifuge: A technical analysis of what Stuxnet's creators tried to achieve," The Langer Group, Hamburg, Germany, Nov. 2013. [Online]. Available: <https://www.langner.com/wp-content/uploads/2017/03/to-kill-a-centrifuge.pdf>
2. N. Falliere, L. O. Murchu, and E. Chien, "W32.Stuxnet Dossier, Version 1.4," Symantec, Cupertino, CA, 2011. [Online]. Available: <https://nsarchive2.gwu.edu//NSAEBB/NSAEBB424/docs/Cyber-044.pdf>
3. J. E. Rubio, C. Alcaraz, and J. Lopez, "Preventing advanced persistent threats in complex control networks," in *Computer Security – ESORICS 2017*, S. N. Foley, D. Gollmann, and E. Sneekenes, Eds. New York: Springer-Verlag, 2017, pp. 402–418.
4. S. Miller, N. Brubaker, D. Kappelmann Zafra, and D. Caban, "Triton actor TTP profile, custom attack tools, detections, and ATT&CK mapping," FireEye, 2019. [Online]. Available: <https://www.fireeye.com/blog/threat-research/2019/04/triton-actor-ttp-profile-custom-attack-tools-detections.html>
5. F. Kargl, R. W. van der Heijden, H. König, A. Valdes, and M. C. Dacier, "Insights on the security and dependability of industrial control systems," *IEEE Security Privacy*, vol. 12, no. 6, pp. 75–78, Nov./Dec. 2014. doi: 10.1109/MSP.2014.120.
6. J. E. Rubio, R. Roman, and J. Lopez, "Analysis of cybersecurity threats in Industry 4.0: The case of intrusion detection," in *Critical Information Infrastructures Security*, G. D'Agostino and A. Scala, Eds. New York: Springer-Verlag, 2018, pp. 119–130.
7. M. Dietz and G. Pernul, "Digital twin: Empowering enterprises towards a system-of-systems approach," *Bus. Inform. Syst. Eng.*, 2019. doi: 10.1007/s12599-019-00624-0.
8. M. Eckhart and A. Ekelhart, "Towards security-aware virtual environments for digital twins," in *Proc. 4th ACM Workshop on Cyber-Physical System Security (CPSS '18)*, June 2018, pp. 61–72. doi: 10.1145/3198458.3198464.
9. S. Adepu, N. K. Kandasamy, and A. Mathur, "EPIC: An electric power testbed for research and training in cyber physical systems security," in *Computer Security*, S. K. Katsikas et al., Eds. New York: Springer-Verlag, 2019, pp. 37–52.
10. A. P. Mathur and N. O. Tippenhauer, "SWaT: A water treatment testbed for research and training on ICS security," in *Proc. Int. Workshop on Cyber-Physical Systems for Smart Water Networks (CySWater)*, Apr. 2016, pp. 31–36. doi: 10.1109/CySWater.2016.7469060.
11. D. Antonioli, H. Ghaeini, S. Adepu, M. Ochoa, and N. O. Tippenhauer, "Gamifying ICS security training and research: Design, implementation, and results of S3," in *Proc. Workshop on Cyber-Physical Systems Security and Privacy (CPS '17)*, Nov. 2017, pp. 93–102. doi: 10.1145/3140241.3140253.
12. C. Tankard, "The security issues of the Internet of Things," *Comput. Fraud Secur. Bull.*, vol. 2015, no. 9, pp. 11–14, Sept. 2015. doi: 10.1016/S1361-3723(15)30084-1.
13. E. Lisova, I. Šljivo, and A. Čaušević, "Safety and security co-analyses: A systematic literature review," *IEEE Syst. J.*, vol. 13, no. 3, pp. 2189–2200, Sept. 2019. doi: 10.1109/JSYST.2018.2881017.
14. G. Sabaliauskaite and S. Adepu, "Integrating six-step model with information flow diagrams for comprehensive analysis of cyber-physical system safety and security," in *Proc. 2017 IEEE 18th Int. Symp. High Assurance Systems Engineering (HASE)*, Singapore, pp. 41–48. doi: 10.1109/HASE.2017.25.
15. M. Dietz, B. Putz, and G. Pernul, "A distributed ledger approach to Digital Twin secure data sharing," in *Data and Applications Security and Privacy XXXIII*, S. N. Foley, Ed. New York: Springer-Verlag, 2019, pp. 281–300.

**Marietheres Dietz** is a research assistant in the Department of Information Systems and a doctoral student at the University of Regensburg, Germany. Dietz studied business-information systems at the University of Regensburg and Linnaeus University, Växjö, Sweden, and received her M.S. with honors. Her research focuses on the digital twin paradigm, with a special interest in the security perspective. Contact her at [marietheres.dietz@ur.de](mailto:marietheres.dietz@ur.de).

**Günther Pernul** is a professor in the Department of Information Systems at the University of Regensburg, Germany. Pernul received his diploma and Ph.D. (with honors) in business informatics from the University of Vienna, Austria. Previously, he held positions at the University of Duisburg–Essen, Germany; University of Vienna; University of Florida, Gainesville; and College of Computing at the Georgia Institute of Technology, Atlanta. His research interests focus on data and information-security aspects, data protection and privacy, data analytics, and advanced datacentric applications. He is a Member of the IEEE. Contact him at [guenther.pernul@ur.de](mailto:guenther.pernul@ur.de).

### 3 Integrating Digital Twin Security Simulations in the Security Operations Center

---

Current status:	Published
Conference:	Availability, Reliability and Security (ARES) - 15th International Conference, Virtual Event, Ireland, August 25-28, 2020
Date of acceptance:	June 08, 2020
Full citation:	DIETZ, M., VIELBERTH, M. AND PERNUL, G. Integrating Digital Twin Security Simulations in the Security Operations Center. In <i>Proceedings of the 15th International Conference on Availability, Reliability and Security</i> . ACM, New York, NY, USA (2020), pp. 1–9.
Authors contributions:	Marietheres Dietz 45% Manfred Vielberth 45% Günther Pernul 10%

---

**Conference Description:** The International Conference on Availability, Reliability and Security (ARES) brings together researchers and practitioners in the field of IT security and privacy. In particular, the conference addresses the reciprocal action between the basic principles and practical issues of security. ARES is published by ACM under the International Conference Proceedings Series (ICPS).

# Integrating Digital Twin Security Simulations in the Security Operations Center

Marietheres Dietz  
marietheres.dietz@ur.de  
Department of Information Systems  
University of Regensburg  
Germany

Manfred Vielberth  
manfred.vielberth@ur.de  
Department of Information Systems  
University of Regensburg  
Germany

Günther Pernul  
guenther.pernul@ur.de  
Department of Information Systems  
University of Regensburg  
Germany

## ABSTRACT

While industrial environments are increasingly equipped with sensors and integrated to enterprise networks, current security strategies are generally not prepared for the growing attack surface that resides from the convergence of their IT infrastructure with the industrial systems. As a result, the organizations responsible for corporate security, the Security Operations Center (SOC), are overwhelmed with the integration of the industrial systems.

To facilitate monitoring the industrial assets, digital twins represent a helpful novel concept. They are the virtual counterparts of such assets and provide valuable insights through collecting asset-centric data, analytic capabilities and simulations. Moreover, digital twins can assist enterprise security by simulating attacks and analyzing the effect on the virtual counterpart. However, the integration of digital twin security simulations into enterprise security strategies, that are mainly controlled by the SOC, is currently neglected.

To close this research gap, this work develops a process-based security framework to incorporate digital twin security simulations in the SOC. In the course of this work, a use case along with a digital twin-based security simulation provides proof of concept. It is demonstrated how a man-in-the-middle attack can be performed in a simulated industry setting and how it affects the systems. Moreover, we show how the resulting system logs can support the SOC by building technical rules to implement in Security Information and Event Management (SIEM) systems.

## CCS CONCEPTS

• **Computer systems organization** → **Embedded and cyber-physical systems**; • **Security and privacy** → **Virtualization and security**.

## KEYWORDS

digital twin, security operations center, security information and event management, security framework

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 ACM 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](mailto:permissions@acm.org).

ARES 2020, August 25–28, 2020, Virtual Event, Ireland

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8833-7/20/08...\$15.00

<https://doi.org/10.1145/3407023.3407039>

## ACM Reference Format:

Marietheres Dietz, Manfred Vielberth, and Günther Pernul. 2020. Integrating Digital Twin Security Simulations in the Security Operations Center. In *The 15th International Conference on Availability, Reliability and Security (ARES 2020)*, August 25–28, 2020, Virtual Event, Ireland. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3407023.3407039>

## 1 INTRODUCTION

The increasing integration of industrial assets in corporate networks leads to a convergence of operational technology (OT) and information technology (IT). While this phenomena entails various benefits, e.g. enhanced monitoring and predictive maintenance, the attack surface eventually increases.

In current enterprises, the Security Operations Center (SOC) upholds corporate security. On the technical side, the SOC is commonly supported by Security Information and Event Management (SIEM) systems. The capabilities of such systems range from managing security-relevant data, over security analyses to deduce security rules or patterns as well as to check the adherence of security rules. However, contemporary SIEM systems mainly monitor the IT infrastructure, while the incorporation of industrial systems is pressing.

A novel paradigm, currently considered an essential milestone in companies, especially in those pursuing Industry 4.0, is the digital twin. It refers to a virtual representation of an enterprise asset – at most an industrial one. It relies, amongst others, on simulation technology to analyse potential outcomes of physical processes and to determine machine fatigue. Moreover, some digital twins even employ prescriptive maintenance that provides maintenance solutions on top. Next to optimization of production, digital twins may contribute to corporate security. For instance, simulating an attack on an industrial asset with digital twins might provide information about system weaknesses and behavior under attack. Thus, it provides potential assistance for SOC and SIEM systems by delivering novel security insights on industrial systems that are currently neglected.

This paper tackles this issue by proposing an effective integration of the digital twin paradigm to SOC and SIEM systems that provides novel potentials for security. The main contributions can be summarized as follows:

- development of a process-based security framework and its formal requirements
- proposition of a use case for demonstration
- evaluation by prototypical implementation

The remainder of this work is organized as follows. Section 2 provides the foundations for our research, outlines related works and



ARES 2020, August 25–28, 2020, Virtual Event, Ireland

Dietz et al.

the research gap. Subsequently, Section 3 proposes a process-based security framework to integrate digital twin security operations with SIEM and SOCs and states formal requirements to achieve this integration. In Section 4, we evaluate our framework by a use case and implemented prototype. Finally, a conclusion of our work is drawn and future work is stated in Section 5.

## 2 BACKGROUND

The following sections lay the foundation of this work and introduce the respective related work. The first two focus on the SOC and SIEM concept. The next section presents the digital twin paradigm, while the subsequent section addresses works employing simulations in the security field and points out the addressed research gap.

### 2.1 Security Operations Center

The Security Operations Center (SOC) represents an *organizational aspect* of a security strategy in an enterprise by providing *procedures, technologies and people* [17, 25]. It is usually not seen as a single entity or system, but rather as a complex structure to manage and enhance the overall security posture of an organization, whereby the core purpose of a SOC is the protection of the organization’s system infrastructure. Thereto, it integrates, monitors and analyses all security-relevant systems and events in a central point. In general, the activities within a SOC can be classified as reactive and proactive, although these cannot always be clearly separated. An integral task of the SOC is to handle alerts and take countermeasures to protect data and applications. Furthermore, it provides governance and compliance as a framework, in which people operate and to which processes and technologies are tailored. When installed and operated correctly, a SOC improves an organizations security posture, creates situational awareness, mitigates the exposed risks, and helps to fulfill regulatory requirements [14]. Since people play an essential role in the security of companies, this is also an important part of a SOC. From the SOC manager to the analyst, a variety of roles can be defined, whereby a SOC must take care of staffing and recruitment. Furthermore, the security awareness of employees can also be assigned to a SOC. To realize the technical side for security operations, SOCs commonly employ, amongst others, SIEM systems as central tools.

### 2.2 Security Information and Event Management

A key aspect of today’s SIEM systems is that it provides a holistic and centralized view on all security relevant systems of an organization, whereas other systems (such as Network Intrusion Systems) only take a limited perspective on selected systems or functions.

A pattern describing SIEM abstractly was proposed in [28]. Compared to the anatomy of SIEM introduced by [18], it describes SIEM in more detail by considering relevant interfaces and a broader breakdown of relevant components. In general, a SIEM provides means for collecting data (such as log data or network flows) from various heterogeneous sources and reports about incidents by humans [27]. To improve its value for further processing, the collected data gets enriched with context data and is normalized into a uniform format. The core of a SIEM is the correlation and analysis

module, which interconnects the gathered data and deduces possible security incidents or abnormalities. Most SIEM systems provide interfaces for sharing the gained threat intelligence externally with other systems on the one hand, and interfaces for human or expert interaction on the other hand. In case of an incident, measures for incident response can be taken either automatically or supported by staff that gets informed by alerts.




Figure 1: Black box view of SIEM embedded in a SOC.

Moreover, recent works propose to deploy SIEM systems to handle the industrial security by incorporating OT systems next to IT systems [23]. In the course of this paper the in- and outputs of a SIEM are of special importance, as these provide interaction options with digital twins. In Figure 1, these are illustrated by a black box view of a SIEM.

### 2.3 Digital Twin

Generally, a digital representation of any real-world counterpart such as a system, product, process or other enterprise asset over its lifecycle can be referred to as digital twin [2]. The digital twin comprises asset-specific data and typically adds context to the data by semantic technologies. Based on the semantically linked data, analyses such as predictive maintenance can be conducted [5]. Moreover, the data allows to model the real-world counterpart virtually in order to conduct simulations [12]. Overall, simulations play a decisive role in digital twin research [20]. Dependent on the specific use case, digital twins are capable of asset management ranging from simple monitoring to autonomy. Figure 2 illustrates this paradigm. It highlights the simulation aspect of the paradigm which presents the focus of this work in respect to digital twin research.




Figure 2: The Digital Twin paradigm.

Various authors mention the importance of security in the digital twin concept (e.g. [10], [24] and [26]), which can also be considered from two different angles: At first, security mechanisms may support the digital twin such as presented in [7]. The second perspective proposes the application of the digital twin concept to enhance security, e.g. [8] and [9]. For the latter, the digital twin provides versatile potentials to enhance current security analyses [6], such as analyses on an asset's historical data, emulated asset-environments to simulate cyber-attacks and transmission of an asset's current state (e.g. by sensor data). Our work focuses on how the simulation of security-relevant incidents within a digital twin can provide helpful insights for SOCs and support the enhancement of SIEM systems.

#### 2.4 Simulations and their Use in Cyber Security

In contrast to traditional intrusion detection mechanisms, simulations of cyber security incidents (e.g. attack simulation) can lead to completely new ways for security monitoring. Simulations differ from simple analyses as they are based on a model, e.g. of a real-world asset. Moreover, simulations depend on user-specified settings and parameters. Consequently, security can benefit from cyber security simulations in the following ways [6]: Simulations enable repeatability and offer to compress the time interval. Additionally, simulations can show a system's behavior under a broad range of specified configurations like during a security incident, which supports the comprehension of emergent as well as prospective behavior. Most importantly, simulations run in a standalone-virtual environment and therefore do not affect the physical environment. Moreover, whenever historical data is missing, resulting data of simulations might deliver important input. Despite all these benefits, simulations are mostly neglected in current security approaches. However, during recent years, some authors have approached this area as follows:

*Testbeds, cyber ranges and honey pots.* The generation of testbeds is commonly employed for critical infrastructure testing and represents an approach that usually combines virtual and physical components. For instance, the Opnet module [19] is applied in numerous testbeds (e.g. [13], [3]) to integrate physical network devices with the virtual part of the testbeds. Testbeds are often used to test planned infrastructures in terms of functionality, but also to assess the level of security e.g. by spotting vulnerabilities. Cyber ranges are virtual environments to develop IT systems or infrastructures [11]. Their main purpose is to provide a training environment with tools that help improve the security as well as stability and performance of IT infrastructures and systems [22]. Honey pots are commonly employed to attract attackers by emulating real-world systems and simulating their behavior. Their usage mainly aims at extracting attack methodologies including the attackers' tactics, techniques and procedures (TTP). To create a more realistic system for the attacker than virtualized machine environments, honey pots perform better when automation hardware is integrated [23].

Our approach extends work of [9], where a digital twin of a cyber-physical system (CPS) is generated to run a man-in-the-middle (MITM) attack simulation on it. However, the security functionalities only lie within their digital twin solution, while our approach crosses borders beyond the digital twin through the inclusion of

SOCs and SIEM systems. To the best of our knowledge, a holistic approach to employ simulations to enhance corporate security management is missing. Therefore, we suggest applying simulations of security incidents (e.g. threats, attacks) in the virtual counterparts of assets (digital twins) and to combine the gathered information with SOC and SIEM systems that protect these assets.

### 3 PROCESS-BASED SECURITY FRAMEWORK

This chapter introduces the proposed approach to holistically address security by combining the digital twin concept, SOC and SIEM systems. The resulting process-based framework is illustrated in Figure 3, which applies the Business Process Model and Notation (BPMN) modelling technique to describe the process, its corresponding actors, systems and artifacts. Each process activity is explained in the individual sections that refer to the respective BPMN swim lane. The first part addresses the SOC activities, the next part explains the process steps within the digital twin and afterwards the activities conducted by the security analytics tools are presented. Finally, the formal requirements of this framework are summarized.

#### 3.1 The Role of the Security Operations Center

In order to detect possible security breaches or system security weaknesses, a large amount of data often must be collected and processed – usually handled by SIEM systems. However, SIEM systems also require domain knowledge about threats etc. and thus, highly depend on cognitive processes [4]. For instance, although SIEM systems allow setting and monitoring security rules, these rules must be created in advance by experts, who we assume to be organized within a SOC.

*Determination of simulation settings.* This first activity of the process is conducted by the SOC. Within this activity, the security experts decide the purpose of the subsequent simulation in the digital twin. This includes deciding which security incident to be simulated as well as which parameters and settings to use. For instance, they could decide to run an attack on the communication between a host to another host using an MITM attack. Another security simulation purpose could be to test if specified security rules apply – preventing the exploitation of vulnerabilities. The output of this activity is the *Simulation settings*-artifact.

*Incident Analysis.* Subsequent to the security simulation by the digital twin, the SOC studies the simulation output. Thereby, it analyzes the *Incident data* produced by the digital twin simulation. The outcome of this analysis is the *Incident information*-artifact. This information is further used in the subsequent incident detection and handling of the test SIEM to verify the derived SIEM logic/patterns in order to catch the security incidents.

#### 3.2 Security Simulation with the Digital Twin

To ensure the security of the system, attack detection mechanisms and rules require constant review and development. The Digital Twin can serve this purpose. Thus, a central part of our framework lies in simulating security incidents within a digital twin.

In contrast to the related work introduced in section 2.4, the digital twin provides a *sheer digital approach* with *strong focus on the asset* instead of the attack [6]. Moreover, it allows to model modular

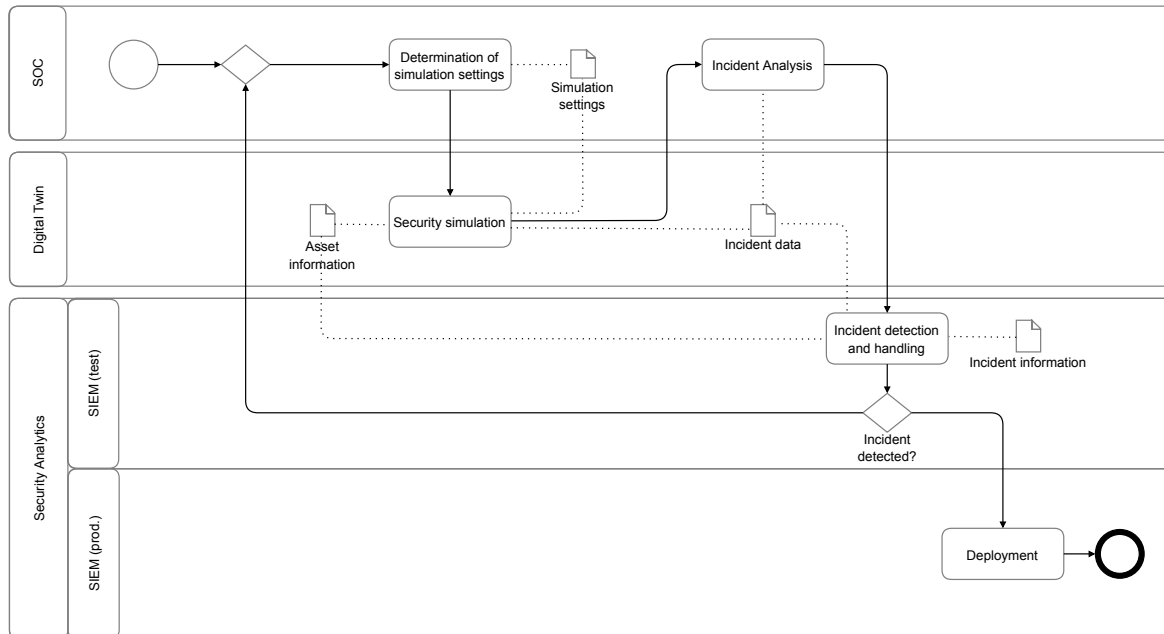


Figure 3: Process-based Security Management Framework integrating Digital Twin security simulations.

components of an asset as well as their network and interplay. Another advantage presents the *bi-directional connection* of the real-world asset with its digital twin. Therefore, subsequent to the evaluation of the simulation outcome, the real-world counterpart can be adapted to the desired setting by transmitting the respective commands and parameters. Possible benefits of this approach are damage limitation of the physical asset, lower prototyping costs, reduction of asset downtime and certainly the promotion of a high security level.

*Security simulation.* At first, the digital twin virtually represents the real-world asset with the help of the *Asset information*-artifact. The previous activity of the SOC sets the purpose of the simulation and its related conditions. Thereby, the purpose of the simulation can be distinguished between:

- **System security testing:** The simulation investigates security configurations of the emulated systems to investigate potential misconfigurations or vulnerabilities. It further provides feedback whether the weakness is present. For instance, in order to reveal security weaknesses, specified security rules of the system might be tested.
- **Pentesting:** While system security testing involves the assessment and reporting of security weaknesses, pentesting attempts to exploit these vulnerabilities in order to verify the feasibility of malicious activity. In short, it assesses the system's sensitivity to security incidents.
- **Attack simulation:** The simulation of an attack reveals the behavior of the system under attack. Its subsequent analysis

can derive patterns. These can be signature-based, anomaly-based or specification-based [16].

Note that the presented simulation types are not absolutely distinct from each other and might exist additional as well as hybrid forms. By setting the system's conditions and details (*Simulation settings*-artifact), the simulation can be created. Each security simulation produces certain data, e.g. logs reporting occurred events. The produced *Incident data* of the digital twin security simulation is analyzed in the succeeding activity by the SOC.

### 3.3 Security Analytics with SIEM

The benefit of a SIEM is its holistic view on multiple IT systems. In order to detect security incidents, it correlates events and analyses them. Thereby, various techniques for attack detection are applied. The most common method is rule based detection [18] but is expanded by more advanced methods in modern SIEMs. To avoid influencing the production system negatively during testing a test SIEM can be set up in the test environment, where it can be linked to the digital twin simulation.

*Incident detection and handling.* Based on the *Incident information*-artifact, the SIEM ideally detects the incident. The incident should be detected in the *Incident data* produced by the digital twin security simulation. In case the incident detection was not successful, either the simulation settings or the analysis of the incident by the SOC were erroneous or incomplete. Thus, these steps must be reevaluated and repeated. *Incident information* can for example be detection rules, which were deduced manually by SOC analysts or

automatically derived patterns by machine learning methods such as anomaly detection. To determine the risk of the incident, *asset information* such as the asset value is needed.

*Deployment.* If the incident was detected successfully in the test environment, the security rules or security settings, such as a trained detection model, are transferred to the production SIEM so that the incident can be detected upon occurrence in the production environment. Therefore, the utilized technology in the production SIEM should be the same as in the test SIEM in order to avoid side effects during deployment.

### 3.4 Formal requirements

On the basis of the explanations of the framework’s activities, the requirements to be met are formally stated in the following.

**REQUIREMENT 1 (DETERMINATION OF SIMULATION SETTINGS).** A set of simulation settings  $S = \{s_1, s_2, \dots, s_n\}, n \in \mathbb{N} \setminus \{0\}$  is determined.

**REQUIREMENT 2 (SECURITY SIMULATION).** Asset information  $I_{Asset}$  linked with simulation settings  $S$  are required to build the simulation  $SIM$ , which produces incident data  $D_{Inc}$ :

$$I_{Asset} \bullet S \rightarrow SIM \rightarrow D_{Inc}$$

**REQUIREMENT 3 (INCIDENT ANALYSIS).** Incident information  $I_{Inc}$  is deduced from data  $D_{Inc}$ :

$$D_{Inc} \models I_{Inc}.$$

$I_{Inc} = \{I_{Det}, I_{Rel}, I_{Pri}\}$ , whereby  $I_{Det}$  is information to incident detection (such as rules),  $I_{Rel}$  is the detection reliability and  $I_{Pri}$  the priority or severity of the incident.

**REQUIREMENT 4 (INCIDENT DETECTION + DEPLOYMENT).** Incident information  $I_{Inc}$  implemented in a functional SIEM that detects the incident:

$$SIEM(I_{Inc}, D_{Inc}) \rightarrow Inc$$

Determination of risk  $I_{Ris}$ , based on  $I_{Asset}, I_{Rel}$  and  $I_{Pri}$ :

$$I_{Ris} = I_{Rel} \bullet I_{Pri} \bullet I_{Asset}$$

Each of these requirements corresponds to at least one activity of the process-based framework (Figure 3). Moreover, all artifacts of the framework, which serve as inputs resp. outputs of the activities, can be found formally stated:

- *Asset information:*  $I_{Asset}$
- *Incident data:*  $D_{Inc}$
- *Incident information:*  $I_{Inc}$

In regard to the swimlanes, the simulation carried out in the *Digital Twin* and the *SIEM* system are formalized into *SIM* and *SIEM*.

## 4 EVALUATION

To evaluate the proposed approach a security simulation for a digital twin of an industrial filling plant is implemented. Furthermore, it is demonstrated how this output can serve the SOC to build SIEM rules. This is tested by implementing a SIEM tool that directly receives the logs produced by the digital twin simulation. In the following, the use case (an industrial filling plant) and the tools used for the prototypical implementation are introduced. Afterwards, we concentrate on the conceptual setting of the attack. Afterwards, the obtained results are summarized. We review each activity of

our process-based security framework (Figure 3) in our use case, and show how the formal requirements can be met. The evaluation is concluded by a discussion of the use case and the proposed approach.

### 4.1 Use Case and Tools

The physical process and the relevant components of our use case are illustrated in Figure 4. The industrial filling plant consists of a tank that contains some liquid and an actuator, e.g. a motor-driven valve (MV), which controls the outflow of the tank. The liquid flows through the pipe into a bottle. The tank, the bottle and the pipe each have a sensor that reads the liquid level (LL) resp. the flow level (FL). Three programmable logic controllers (PLCs) monitor the sensors and the actuator. In the use case setting, PLC2 controls the sensor measuring the flow level in the pipes (Sensor2-FL) and PLC3 controls the liquid level of the bottle that is to be filled (Sensor3-LL-bottle). Meanwhile, PLC1 gets hold of the sensor value of the tank’s liquid level (Sensor1-LL-tank) and follows a control strategy for the motor-driven valve (Actuator1-MV). To accomplish the control of the actuator, PLC1 receives the other sensor values controlled by PLC2 and PLC3. The network communication is realized over Ethernet/IP (ENIP) and organized by ENIP tags that store the sensor values in the respective PLCs. For instance, the "Sensor3-LL-bottle"-tag is stored in PLC3, and can be requested and received by PLC1.

Figure 6 shows the use case’s network infrastructure. Next to the PLCs, a Human-Machine Interface (HMI) allows direct control the actuator (open/close). Together, the Figures 4 and 6 can be referred to as *Asset information-artifact* as presented in our process-based framework.




Figure 4: Use case – An industrial filling plant.

The following tools are utilized for prototypical implementation of the filling plant simulation within the digital twin. Commonly, Mininet<sup>1</sup> is employed for simulations with a digital twin [8, 9]. To fit our use case, we chose a technology relying on Mininet, namely MiniCPS<sup>2</sup>, which enables the simulation of industrial assets and originates from research [1]. MiniCPS is tailored for industrial settings: It simulates of traditional industrial systems like PLCs, HMIs and common industrial network communication over ENIP or Modbus. Moreover, an underlying database supports the simulation of the physical process by storing the current states, e.g. the liquid

<sup>1</sup><http://mininet.org>

<sup>2</sup><https://github.com/scy-phy/minicps>

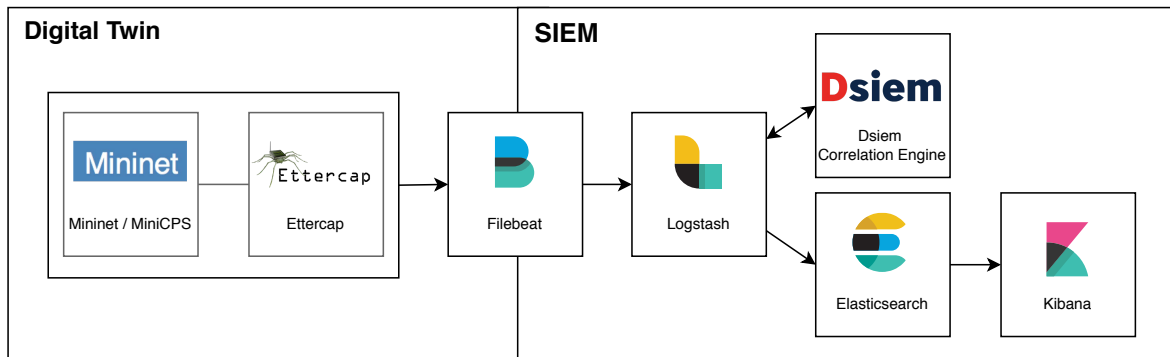


Figure 5: Implemented micro-service architecture for digital twin and SIEM deployment – managed by the SOC.

levels, from which the simulated sensors get their values. Moreover, each PLC in MiniCPS can be implemented with a own ENIP server that can store tags [1]. For instance, in the use case setting, PLC2 stores the value of Sensor2 as ENIP tag ("SENSOR2-FL") in its ENIP server. Additionally, we added a logging strategy to capture the respective system's events (system logs) to each of the MiniCPS PLCs main function. Currently, they are stored in the logs folder of the project, from where they can be directly transferred to the SIEM system.

To simulate the attackers behavior, the attacker node in our setting currently makes use of tools like Ettercap<sup>3</sup> to capture the network traffic and start a MITM attack. The respective source code with installation details, further description of the use case implementation and the resulting logs are available for the public and can be found at Github<sup>4</sup>.

The technical side of the incident analysis within the SOC is supported by the SIEM tool Dsiem<sup>5</sup>, which builds upon Elasticsearch, Logstash and Kibana<sup>6</sup>. These tools are under open source licence to a large extent, and are commonly favored in research for their in-depth comprehensibility.

The overall architecture of the tools and their interaction is shown in Figure 5. From the SIEM system side, Filebeat is responsible for the collection of the log data that is transmitted by the digital twin simulation into the files of the logs folder. Within Filebeat, the log files are monitored and in case new lines are added, these are transmitted to Logstash. Logstash normalizes the data by parsing the logs line by line and transforming them into semi-structured JSON-documents in order to enable and facilitate further processing and readability. Dsiem is a correlation engine, which detects incidents based on rules and offers the possibility to trigger an alarm. If an alarm is triggered, it is pushed back to logstash to pass it along the pipeline. Elasticsearch provides the data storage and query execution. Finally, Kibana visualizes the data, displays alarms and offers analytics capabilities, which enables experts to manually analyze incidents.

<sup>3</sup><https://www.ettercap-project.org/>

<sup>4</sup><https://github.com/FrauThes/DigitalTwin-SIEM-integration>

<sup>5</sup><https://www.dsiem.org/>

<sup>6</sup><https://www.elastic.co/>

The total environment is realized in the form of a micro-service architecture, where each component is deployed within a docker container. This facilitates the later transfer into the production system and simplifies the use of the framework for research in order to build upon or extend it. Furthermore, individual components are easily replaceable if more suitable ones for the respective environment emerge.

#### 4.2 Attack Setting

*SOC: Determination of simulation settings.* The first activity of process-based framework (see Figure 3), defines the attack scenario. In our use case, the SOC determines the attacker to be present in the network of the filling plant as shown in Figure 6.




Figure 6: The attacker inside the network topology of the industrial filling plant.

Moreover, the SOC lets the attacker start an ARP spoofing MITM attack between PLC1 that monitors the actuator (open/close motor-driven valve) and PLC3, which sends the liquid level of the bottle to PLC1. More precisely, the attacker only shortly sniffs the network traffic between PLC1 and PLC3 and stops the attack. This is repeated about every two minutes. The goal is to not completely stop the network traffic for a longer time period, so that the physical process seems to be running regular and is not stopped at all. This attack strategy proceeds similarly to the spying phase of Advanced Persistent Threats (APTs). With this procedure, the SOC tests if

their log management strategy is sufficient to detect this simple attack.

In terms of requirements, our use case produces the set of settings  $S = \{attacker\ in\ network, attack, targets, repeat, attack\ duration\}$ , whereby  $attacker\ in\ network = true$ ,  $attack = MITM$  ARP spoofing,  $targets = \{PLC1, PLC3\}$ ,  $repeat = \{true, 2\ min\}$ ,  $attack\ duration = 15\ s$  that fulfills REQUIREMENT 1.

### 4.3 Results

*Digital Twin: Security simulation.* The purpose of the performed simulation as dictated by the SOC, can be summarized as a hybrid of attack simulation and pentesting as it reveals how the attack affects the systems and their response. To build the simulation  $SIM$ , the general information about the asset  $I_{Asset}$ , i.e. the information about the network topology and the physical process, is linked with the simulation settings  $S$  (REQUIREMENT 2). To start the attack, the attacker node of the simulation executes the shell commands for the MITM attack.

```
INFO 03/16/2020 13:31:59 10.0.0.1 main_loop
Liquid level (SENSOR 3) under
BOTTLE_M['UpperBound']: 0.84 < 0.90
-> open mv (ACTUATOR 1).
INFO 03/16/2020 13:32:01 10.0.0.1 main_loop
Flow level (SENSOR 2) under
SENSOR2_THRESH: 2.45 < 3.00
-> leave mv status (ACTUATOR 1).
WARNING 03/16/2020 13:32:06 10.0.0.1 main_loop
Liquid level (SENSOR 3) is not received.
Program is unable to proceed properly
INFO 03/16/2020 13:32:08 10.0.0.1 main_loop
Flow level (SENSOR 2) under
SENSOR2_THRESH: 2.45 < 3.00
-> leave mv status (ACTUATOR 1).
```

**Listing 1: System logs of PLC1**

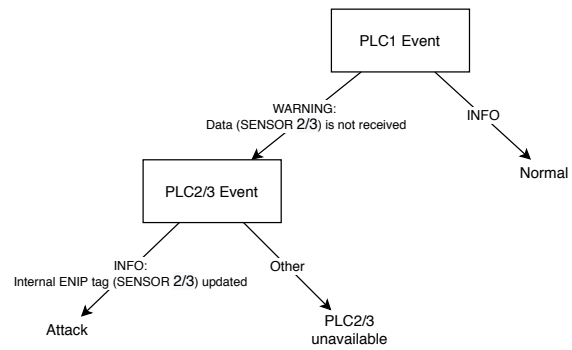
The output of the security simulation  $SIM$  is the incident data  $D_{Inc}$  (REQUIREMENT 2), which, in this case, are several log files, i.e. the system logs of PLC1, PLC2 and PLC3. This security simulation output of the digital twin as referenced in Figure 3 (*Incident data-artifact*), can be found at GitHub<sup>7</sup>. Listing 1 shows an extract of the logged system events of PLC1 (10.0.0.1). As can be seen from the logs, the control strategy of PLC1 works fine until the value of Sensor3 that is managed by PLC3 cannot be received any more. This causes the program to produce error messages and leads to the actuator being left at the state as it was at last. In the case of Listing 1, it remains open. The physical result would be a bottle overflow, which would only be stopped when the tank reaches its lower-bound threshold or after the attacker stops the MITM attack.

*SOC: Incident Analysis.* The analysis of simulation output  $D_{Inc}$  by the SOC suggest that the ARP poisoning MITM attack allows to place the attacker successfully between PLC1 and PLC3 and results in a denial-of-service (DOS) of network communication between PLC1 and PLC3. Thereby, the pattern of the ARP MITM attack can be deduced from  $D_{Inc}$ , and the *Incident information-artifact*  $I_{Inc}$  is created (REQUIREMENT 3): The artifact mainly consists of multiple correlation rules  $I_{Det}$  to detect the attack. Therefore, it can

<sup>7</sup><https://github.com/FrauThes/DigitalTwin-SIEM-integration/tree/master/example-logs>

be detected in two stages. The first stage detects, that the sensor data of PLC3 was not received. If the condition of the first stage is fulfilled, the second stage waits for log data within a specific time window, that shows, that PLC3 is operating normally and that there must be a communication problem between PLC1 and PLC3, which in turn indicates a potential MITM attack. Figure 7 exemplary shows the generalized detection rule in the form of a decision tree. The actually JSON-formatted and implemented correlation rules can be found at Github<sup>8</sup>.

Furthermore, parameters for risk calculation are preset during this step. First, the priority of the incident  $I_{Pri}$  is set. This parameter determines the severity of the incident. In the case of Dsiem the priority is represented as a value between 1 and 5. Since the severity of the attack is estimated as medium,  $I_{Pri} = 3$  is assumed. Second, the reliability  $I_{Rel}$  of the detection is set to a value between 1 and 10. This is based on the two-stage rules as explained above. If it was recognized that no more data is received from PLC3,  $I_{Rel} = 3$  is predefined. The reliability of the first stage is rather low, as the warning might also indicate a general connection problem of PLC3. However, if the second stage rule (PLC3 is operating normally) applies in addition,  $I_{Rel} = 8$  is predefined, since an attack is more likely, but there is still the possibility of a false positive.



**Figure 7: Deduced detection rules for MITM Attack as decision tree**

*SIEM: Incident detection and handling.* To enable incident investigation and the derivation of detection rules or patterns a SIEM system is utilized. To instantly detect the attack within the produced incident data  $D_{Inc}$  and generate an alert, the security rule is implemented:  $SIEM(I_{Inc}, D_{Inc})$ . The physical result of Listing 1 would be a slight overflow of the bottle. The SIEM alert can inform the SOC that an incident  $Inc$  is detected. The SOC, in turn, can react as quickly as possible or even install new mechanisms that completely prevent the overflow scenario. During this step, the risk of the incident  $I_{Ris}$  is calculated based on  $I_{Rel}$ ,  $I_{Pri}$ , and  $I_{Asset}$  (more specifically the value of the asset). In the case of Dsiem the formula  $I_{Ris} = \frac{I_{Rel} \cdot I_{Pri} \cdot I_{Asset}}{25}$  is applied. This results in a risk value  $I_{Ris} = \frac{8 \cdot 3 \cdot 2}{25} = 1,92$  which triggers a low risk alarm. The calculated risk  $I_{Ris}$  and the detected incident  $Inc$  fulfill REQUIREMENT 4.

<sup>8</sup>[https://github.com/FrauThes/DigitalTwin-SIEM-integration/blob/master/deployments/docker/conf/dsiem/configs/directives\\_dsiem-digital\\_twin.json](https://github.com/FrauThes/DigitalTwin-SIEM-integration/blob/master/deployments/docker/conf/dsiem/configs/directives_dsiem-digital_twin.json)

ARES 2020, August 25–28, 2020, Virtual Event, Ireland

Dietz et al.

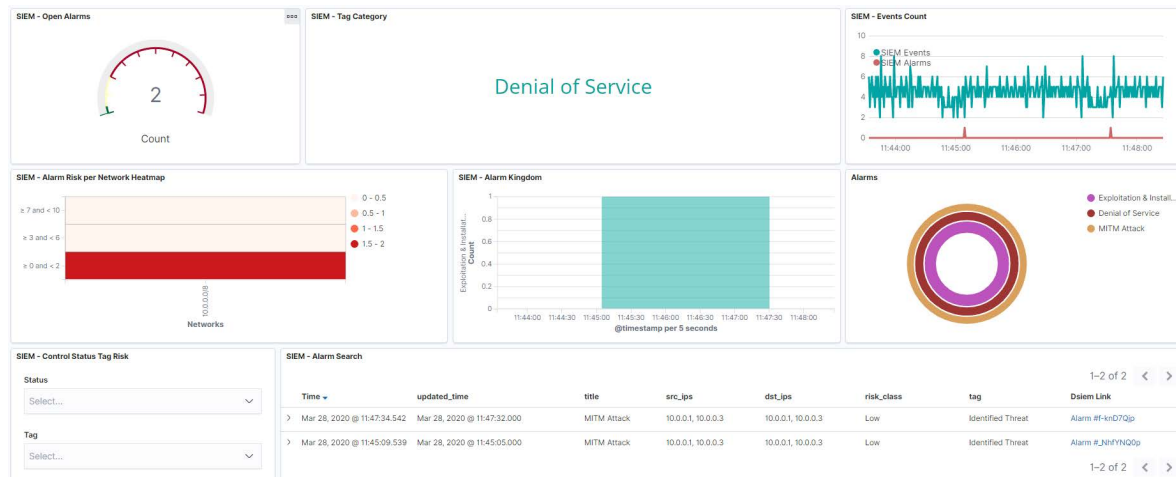


Figure 8: Screenshot of the Dsiem incident detection.

Figure 8 shows the visualization of the prototypical recognition of the incident by Dsiem. In the upper right corner a line chart shows the log data stream (turquoise) and the detected attacks (red peaks). The heatmap on the left side shows the risk level of the detected incidents and allows a quick overview of the threat level. If the incident is successfully detected, one can proceed to deploy the implementation in the production SIEM.

*SIEM: Deployment.* As in the use case, the asset is in the Planning & Design phase, no actual physical asset exists. Therefore, the deployment is currently accomplished in the virtual environment (cf. further explanation in the following section) by carrying out the deduced security rule.

#### 4.4 Discussion

In order to reach significant results, the use case and its simulation requires to be realistic. To achieve a setup as close as possible to reality, the filling plant use case is oriented on the input of a security expert of an industrial company<sup>9</sup> that plans, develops and manufactures machines and complete lines for the fields of process, filling and packaging technology. Additional statistics support the use case settings. For instance, a recent study shows that the communication's over Industrial Ethernet annual growth is at 20%, whereby ENIP is applied the most [21]. Moreover, regarding the current states of OT solutions, software patching remains a slow process [15], representing easy entry points for the attackers. So, the attacker already being in the network topology as illustrated in Figure 6 is a realistic setting indeed. While the ARP spoofing represents also a realistic attack against ICS, future attack simulations could cover more advanced and more affecting scenarios (e.g. ransomware attacks).

Regarding the process-based framework, it might be questionable to what extent other simulation technologies for cyber security can be applied instead of the digital twin. While the concept might be

<sup>9</sup><https://www.krones.com/en/>

realizable with technologies such as testbeds, the digital twin is however the most fitting solution for the following reasons: Digital twins cover the whole lifecycle of their physical counterparts. This means that they can be created in the Planning & Design phase, where no physical counterpart might exist yet, but data can already be gathered. In terms of security, this enables creating an asset by following security-by-design paradigm as the digital twin might already simulate a planned industrial asset and security-related vulnerabilities can be found before the real-world assets exists and operates [6]. This goes hand in hand with the digital twins being mere virtual technologies, and, esp. in comparison to testbeds, no physical assets are mandatory. The main argument for employing digital twins in our framework is that digital twins contain by far more data and analytic capacities than simulations alone (see Figure 2), such as asset-centric data (e.g. context data, domain- and expert-knowledge). This data presents additional important input for SOC and SIEM systems (cf. Figure 1).

## 5 CONCLUSION AND FUTURE WORK

In this work, a process-based framework for integrating digital twin security simulations in SOCs is developed. By prototypical implementation and use case demonstration, it is shown that the suggested approach can be accomplished in practice. As the use case's digital twin security simulation repeats the MITM attack only about every two minutes and sniffs the occurring traffic for only a few seconds, the attack wouldn't be noticed without logs. This is reminiscent of the espionage phase of APTs, which often target industrial environments. However, in industrial systems currently many proprietary, but no standardized, log management solutions are pursued. In this work, we highlight how the digital twin security simulation can support the SOC in their security strategies (e.g. log management and monitoring).

The presented approach is the first of its kind to combine the research areas digital twin and SOC and thus, requires demonstration

Integrating DT Security Simulations in the SOC

ARES 2020, August 25–28, 2020, Virtual Event, Ireland

of feasibility first. In the scope of this work, an evaluation in terms of efficiency or performance is not carried out yet. However, future work will address the performance-based evaluation as well as the creation of more and different attacks to simulate, and the subsequent processing of the simulation output in the SOC and SIEM systems. As addressed above, future work could further extend the framework, e.g. by passing on further security-relevant data from the digital twin to the SOC.

Moreover, future research should tackle to integrate even further – towards cyber threat intelligence (CTI). For instance, the structured description of attacks (e.g. STIX format) could be applied to the simulation data. This includes the report of found vulnerabilities, e.g. via lists like the Common Vulnerabilities and Exposures (CVE). Seeing the even bigger picture, the presented approach could work together with further attack knowledge to simulate the attacks as close to reality as possible, e.g. by integrating knowledge from honeypots to get current methods of the attackers.

Regarding SIEM systems, one major research challenge lies in the complexity of creating detection and correlation rules. Since this requires the definition of multiple parameters and since this syntax for rule definition is designed for SIEM experts, SOC analysts struggle with creating rules. Future work should address this problem by designing approaches, which are targeted at a broader security audience and offer a higher usability – while still providing the ability to define rules for detection of complex incidents. Lowcode approaches, tailored to the special demands within a SOC, could present promising solutions to this problem.

## ACKNOWLEDGMENTS

This work is partly performed under the ZIM SISSEC project (<https://www.it-logistik-bayern.de/produktionslogistik/projekt-sissec>), which is supported under contract by the German Federal Ministry for Economic Affairs and Energy (16KN085725).

We further thank Andreas Reisser from the Krones Group (<https://www.krones.com/en/>) for the helpful remarks to the industrial use case and the professional support from the industry.

## REFERENCES

- [1] Daniele Antonioli and Nils Ole Tippenhauer. 2015. MiniCPS: A Toolkit for Security Research on CPS Networks. In *Proceedings of the First ACM Workshop on Cyber-Physical Systems Security and/or Privacy (CPS-SPC '15)*. ACM, New York, NY, USA, 91–100.
- [2] Stefan Boschert, Christoph Heinrich, and Roland Rosen. 2018. Next Generation Digital Twin. In *Proceedings of the 12th International Symposium on Tools and Methods of Competitive Engineering (TMCE 2018)*, 209–217.
- [3] B. Chen, N. Pattanaik, A. Goulart, K. L. Butler-purry, and D. Kundur. 2015. Implementing attacks for modbus/TCP protocol in a real-time cyber physical system test bed. In *2015 IEEE International Workshop Technical Committee on Communications Quality and Reliability (CQR)*, 1–6. <https://doi.org/10.1109/CQR.2015.7129084>
- [4] Marcello Cinque, Domenico Cotroneo, and Antonio Pecchia. 2018. Challenges and Directions in Security Information and Event Management (SIEM). In *2018 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW)*. IEEE, 95–99. <https://doi.org/10.1109/ISSREW.2018.00-24>
- [5] Marietheres Dietz and Günther Pernul. 2020. Digital Twin: Empowering Enterprises Towards a System-of-Systems Approach. *Business & Information Systems Engineering* 62, 2 (2020), 179–184. <https://doi.org/10.1007/s12599-019-00624-0>
- [6] Marietheres Dietz and Günther Pernul. 2020. Unleashing the Digital Twin's Potential for ICS Security. *IEEE Security Privacy* (2020). <https://doi.org/10.1109/MSEC.2019.2961650>
- [7] Marietheres Dietz, Benedikt Putz, and Günther Pernul. 2019. A Distributed Ledger Approach to Digital Twin Secure Data Sharing. In *Data and Applications Security and Privacy XXXIII*, Simon N. Foley (Ed.), Springer International Publishing, Cham, 281–300. [https://doi.org/10.1007/978-3-030-22479-0\\_15](https://doi.org/10.1007/978-3-030-22479-0_15)
- [8] Matthias Eckhart and Andreas Ekelhart. 2018. A Specification-Based State Replication Approach for Digital Twins. In *Proceedings of the 2018 Workshop on Cyber-Physical Systems Security and Privacy (CPS-SPC '18)*. ACM, New York, NY, USA, 36–47. <https://doi.org/10.1145/3264888.3264892>
- [9] Matthias Eckhart and Andreas Ekelhart. 2018. Towards Security-Aware Virtual Environments for Digital Twins. In *Proceedings of the 4th ACM Workshop on Cyber-Physical System Security (CPSS '18)*, 61–72. <https://doi.org/10.1145/3198458.3198464>
- [10] Matthias Eckhart and Andreas Ekelhart. 2019. *Digital Twins for Cyber-Physical Systems Security: State of the Art and Outlook*. Springer International Publishing, Cham, 383–412. [https://doi.org/10.1007/978-3-030-25312-7\\_14](https://doi.org/10.1007/978-3-030-25312-7_14)
- [11] B. Ferguson, A. Tall, and D. Olsen. 2014. National Cyber Range Overview. In *2014 IEEE Military Communications Conference*, 123–128. <https://doi.org/10.1109/MILCOM.2014.27>
- [12] Michael Grieves and John Vickers. 2017. *Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems*. Springer International Publishing, Cham, 85–113. [https://doi.org/10.1007/978-3-319-38756-7\\_4](https://doi.org/10.1007/978-3-319-38756-7_4)
- [13] A. Hahn, A. Ashok, S. Sridhar, and M. Govindarasu. 2013. Cyber-Physical Security Testbeds: Architecture, Application, and Evaluation for Smart Grid. *IEEE Transactions on Smart Grid* 4, 2 (2013), 847–855. <https://doi.org/10.1109/TSG.2012.2226919>
- [14] Diana Kelley and Ron Moritz. 2006. Best Practices for Building a Security Operations Center. *Information Systems Security* 14, 6 (2006), 27–32. <https://doi.org/10.1201/1086.1065898X/45782.14.6.20060101/91856.6>
- [15] Peter Kieseberg and Edgar Weippl. 2018. Security Challenges in Cyber-Physical Production Systems. In *Software Quality: Methods and Tools for Better Software and Systems*, Dietmar Winkler, Stefan Biffl, and Johannes Bergsmann (Eds.), Springer International Publishing, Cham, 3–16. [https://doi.org/10.1007/978-3-319-71440-0\\_1](https://doi.org/10.1007/978-3-319-71440-0_1)
- [16] Hung-Jen Liao, Chun-Hung [Richard] Lin, Ying-Chih Lin, and Kuang-Yuan Tung. 2013. Intrusion detection system: A comprehensive review. *Journal of Network and Computer Applications* 36, 1 (2013), 16 – 24. <https://doi.org/10.1016/j.jnca.2012.09.004>
- [17] Afsaneh Madani, Saeed Rezayi, and Hossein Gharraee. 2011. Log management comprehensive architecture in Security Operation Center (SOC). In *2011 International Conference on Computational Aspects of Social Networks (CASoN)*. IEEE, 284–289. <https://doi.org/10.1109/CASON.2011.6085959>
- [18] David Miller, Shon Harris, Allen Harper, Stephen VanDyke, and Chris Blask. 2011. *Security information and event management (SIEM) implementation*. McGraw-Hill, New York, NY.
- [19] Sparsh Mittal. 2014. OPNET: An Integrated Design Paradigm for Simulations.
- [20] Elisa Negri, Luca Fumagalli, and Marco Macchi. 2017. A Review of the Roles of Digital Twin in CPS-based Production Systems. *Procedia Manufacturing* 11 (2017), 939–948. <https://doi.org/10.1016/j.promfg.2017.07.198>
- [21] Joakim Nideborn. 2019. Industrial network market shares 2019 according to HMS. <https://www.hms-networks.com/news-and-insights/news-from-hms/2019/05/07/industrial-network-market-shares-2019-according-to-hms> [Online; accessed 19-Mar-2020].
- [22] Cuong Pham, Dat Tang, Ken-ichi Chinen, and Razvan Beuran. 2016. CyRIS: A Cyber Range Instantiation System for Facilitating Security Training. In *Proceedings of the Seventh Symposium on Information and Communication Technology (SoICT '16)*. ACM, New York, NY, USA, 251–258. <https://doi.org/10.1145/3011077.3011087>
- [23] R. Piggan and I. Buffey. 2016. Active defence using an operational technology honeypot. In *11th International Conference on System Safety and Cyber-Security (SSCS 2016)*, 1–6. <https://doi.org/10.1049/cp.2016.0860>
- [24] Juan E. Rubio, Rodrigo Roman, and Javier Lopez. 2018. Analysis of Cybersecurity Threats in Industry 4.0: The Case of Intrusion Detection. In *Critical Information Infrastructures Security*, Gregorio D'Agostino and Antonio Scala (Eds.), Springer International Publishing, Cham, 119–130. [https://doi.org/10.1007/978-3-319-99843-5\\_11](https://doi.org/10.1007/978-3-319-99843-5_11)
- [25] Stef Schinagl, Keith Schoon, and Ronald Paans. 2015. A Framework for Designing a Security Operations Centre (SOC). In *2015 48th Hawaii International Conference on System Sciences (HICSS)*. IEEE, 2253–2262. <https://doi.org/10.1109/HICSS.2015.270>
- [26] Thomas H.J. Uhlemann, Christian Lehmann, and Rolf Steinhilper. 2017. The Digital Twin: Realizing the Cyber-Physical Production System for Industry 4.0. In *Procedia CIRP*, Vol. 61. Elsevier B.V., 335–340. <https://doi.org/10.1016/j.procir.2016.11.152>
- [27] Manfred Vielberth, Florian Menges, and Günther Pernul. 2019. Human-as-a-security-sensor for harvesting threat intelligence. *Cybersecurity* 2, 23 (2019). <https://doi.org/10.1186/s42400-019-0040-0>
- [28] Manfred Vielberth and Günther Pernul. 2018. A Security Information and Event Management Pattern. In *12th Latin American Conference on Pattern Languages of Programs (SugarLoafPLoP)*. The Hillside Group.



## 4 Harnessing Digital Twin Security Simulations for systematic Cyber Threat Intelligence


---


Current status:	Accepted for publication	
Conference:	Computers, Software, and Applications Conference (COMPSAC) - 46th Annual Conference, Virtual Event, Torino, Italy, June 27- July 1, 2022	
Date of acceptance:	April 02, 2022	
Full citation:	DIETZ, M., SCHLETTE, D. AND PERNUL, G. Harnessing Digital Twin Security Simulations for systematic Cyber Threat Intelligence. In <i>46th Annual Computers, Software, and Applications Conference</i> . IEEE Computer Society (2022).	
Authors contributions:	Marietheres Dietz	45%
	Daniel Schlette	45%
	Günther Pernul	10%


---

**Conference Description:** The annual IEEE Computer Society Signature Conference on Computers, Software and Applications (COMPSAC) targets scholars from academia, industry, and government. It welcomes discussions on innovative approaches, research advancements, emerging challenges, as well as future trends in computer and software technologies and applications. The conference proceedings are published by IEEE.

# Harnessing Digital Twin Security Simulations for systematic Cyber Threat Intelligence

Marietheres Dietz   
Chair of Information Systems  
University of Regensburg, Germany  
marietheres.dietz@ur.de

Daniel Schlette   
Chair of Information Systems  
University of Regensburg, Germany  
daniel.schlette@ur.de

Günther Pernul   
Chair of Information Systems  
University of Regensburg, Germany  
guenther.pernul@ur.de

**Abstract**—Understanding cybersecurity threats, attacks, and incidents is crucial for organizations to perform preventive or reactive measures. Nevertheless, detailed Cyber Threat Intelligence (CTI) is reluctantly shared. Digital twins, the virtual counterparts of real-world assets, offer security simulation capabilities. The simulation of attack scenarios on industrial control systems (ICS) with digital twins yields valuable threat information. In our work, we outline the systematic steps towards a structured threat report starting with digital twin security simulations: We first present the course of action and define formal requirements for framework deployment. We then conduct an attack simulation with a prototypical digital twin application to evaluate our framework. Using the STIX2.1 standard, we assist CTI generation by providing utility tools guiding through the process steps. Our experimental results show that a STIX2.1 CTI report can be systematically constructed with the opportunity to customize according to the use case at hand. Adding digital twin security simulations to the list of CTI sources can provide shareable CTI and help organizations improve their security posture.

**Index Terms**—digital twin, cyber threat intelligence, simulation, security analysis

## I. INTRODUCTION

Collaborative cybersecurity demands common standards and the sharing of information about threats, attacks, and security incidents. However, sharing such Cyber Threat Intelligence (CTI) is facing obstacles. Organizations are reluctant to share due to fear of negative repercussions and trust concerns [1]. Security simulations can overcome some of the existing limitations. Harnessing the digital twin for security simulation presents a novel opportunity to generate shareable CTI.

Over the last decade, critical infrastructures and operational environments have become the target of cyberattacks. The industrial control systems (ICS) commonly used are designed to operate for long times, with the result of outdated security [2]. This way, ICS present easy targets for attacks, which turn more cunning every year. Aside from attacks, far-reaching digitization is another current development in the industry. This development puts the concept of the digital twin into focus. Digital twins are able to virtually represent and mirror different assets, including ICS [3].

Common features of digital twins allow to analyze system states, predict machine failures or material fatigue. Besides, digital twins enable security simulations [4], [5]. For instance, penetration testing can be conducted with digital twins while their real-world counterparts continuously operate [6]. Despite

requiring security themselves, digital twins can be leveraged to simulate attack scenarios, enhancing security [7], [8]. The resulting output of digital twin simulations represents important threat information.

The sharing of CTI promises benefits for the community but requires the structured representation of threat information. Therefore, structured CTI reports base on standards (e.g., STIX) and describe attacks, including relationships between indicators. Organizations then leverage CTI for prevention, mitigation and remediation of security incidents.

Recapitulating, to utilize digital twin simulation data a structured representation is required. The integration of digital twin security simulations and CTI can tap into an additional data source and foster collaborative security. In our work we integrate digital twin security simulation and CTI (see Figure 1). We simulate an attack scenario with the digital twin and transform the output into shareable threat intelligence.




Fig. 1. Approach to integrate digital twin security simulations and CTI

Thereby, our work enhances the current state-of-the art by providing the following contributions:

- 1) attack simulation with digital twins and successive structuring into threat intelligence
- 2) definition of formal prerequisites for practical framework deployment
- 3) concept evaluation by implementing a digital twin use case and tool-based guidance on the CTI process steps

The remainder of this work is structured as follows. We first provide background on digital twins and CTI (Section II). Section III presents related works. Integration of digital twin

security simulations into CTI process steps is part of Section IV. To evaluate our framework, a practical implementation is demonstrated (Section V). Subsequently, the framework proposition and the results are discussed in Section VI. Finally, a conclusion is drawn (Section VII).

## II. BACKGROUND

Section II-A states the necessary prerequisites on digital twins. Section II-B covers the foundations of CTI.

### A. Digital Twin

Digital twins are commonly defined as virtual representations of a real-world counterpart over its lifecycle [9]. This real-world counterpart can be a system, product, process, or any other asset [10]. To manage its counterpart, a digital twin comprises asset-specific data, enhanced with semantic technologies [11], and offer capacities to virtually explore the real-world asset (e.g., by simulations [9], [12]). In our work, we use the following security-centered definition of Eckhart and Ekelhart [4]:

*Definition 1:* A digital twin refers to a virtual double of a system along its lifecycle, providing sufficient *fidelity* for *security measures* by the *consumption of (real-time) data* whether required.

The usage of digital twins is currently discussed for security as its different modes benefit security operations [5]: The *analytics/optimization* mode processes incoming data to detect malicious outliers. The *replication* mode entails direct mirroring of the real-world counterpart including all its states. The *simulation* mode bases on a model of the digital twin's real-world double and requires user-specific settings. Thereby, security experts can specify attacks on an industrial plant and examine the effect in the simulated environment.

The *security measure* focused in this work links to the simulation of attacks. To build the simulation, *data* about the real-world system's specification is required:  $S_{Sys} = \{s_1, \dots, s_n\}$  denotes the real-world counterpart's set of specifications. This specification data can be a single file based on industrial formats (e.g., AutomationML<sup>1</sup>) or separate data tuples (e.g., network topology and program code). Ideally, the system's specification data  $S_{Sys}$  can be exactly transferred into the digital twin:  $\hat{S}_{Sys} = \{s_1, \dots, s_m\}$  describes specification data to simulate the security incident. Moreover, it should manifest as subset of the referred real-world system at minimum, leading to  $\hat{S}_{Sys} \subseteq S_{Sys}$ .

While  $P_{Inc} = \{p_1, \dots, p_n\}$  defines the set of parameters of a security incident to the real system,  $\hat{P}_{Inc} = \{p_1, \dots, p_m\}$  does so for the simulation. The security settings  $\hat{P}_{Inc}$  should be realistic parameters that could occur on the real systems, e.g., details of a DoS-attack. Thus again, the incident parameters of the simulation should be the same or a subset of those potentially occurring in real-world resulting in  $\hat{P}_{Inc} \subseteq P_{Inc}$ .

A security incident  $Inc$  in the real system is denoted in Equation 1. Performing the security simulation  $f_{Sim}$  produces

the output  $O_{Sim}$ , which depends on the simulation's specification data as well as on the security settings (Eq. 2).

$$f_{Sys}(Inc) : S_{Sys} \times P_{Inc} \mapsto O_{Sys}(Inc) \quad (1)$$

$$f_{Sim}(Inc) : \hat{S}_{Sys} \times \hat{P}_{Inc} \mapsto O_{Sim}(Inc) \quad (2)$$

Sufficient *fidelity* can be obtained if the simulation output corresponds to the presumed real output denoted as  $O_{Sim}(Inc) \hat{=} O_{Sys}(Inc)$ . The simulation's fidelity also depends on the appropriate implementation. For instance, attacks affecting the network layer must simulate topologies, while physical aspects might be neglected [6].

### B. Cyber Threat Intelligence

Cyber attacks target organizations and their information systems. To defend against advanced attacks and to cope with prevalent security incidents, organizations increasingly focus on the use of Cyber Threat Intelligence (CTI) [13]. At its core, the concept of CTI refers to threat information gathered by security analysts and security systems [1]. Besides, CTI centers on the sharing of actionable knowledge in form of *Indicators of Compromise (IoC)*, *Attack Patterns* or defender's *Courses of Action (CoA)*. Consequently, CTI formats have been proposed to describe incident elements [14]. Most notably, the holistic representation of CTI has resulted in the *Structured Threat Information Expression (STIX)* in version 2.1 [15].

Throughout this work we will rely on the STIX2.1 specification and refer to its individual elements [15]. As a key concept, STIX2.1 links *STIX Domain Objects (SDO)* as well as *STIX Cyber-observable Objects (SCO)* with each other through pre-defined *STIX Relationship Objects (SRO)* or embedded relationships. Thus, threat intelligence elements form a graph  $G(v, e)$  where SDOs and SCOs constitute vertices  $v$  and relationships form edges  $e$ . On a technical level, each SDO, SCO and SRO is serializable as JavaScript Object Notation (JSON)<sup>2</sup> object. As such they are specified to contain a number of mandatory as well as optional attributes (see Appendix A). Some of these attributes are applicable to all STIX2.1 objects whereas others are object specific.

While a standardized and high-quality representation of CTI is a major concern [16], procedural elements also play an important role. The creation of CTI is commonly based on security incidents and thus performed ex post. However, simulations can constitute an additional ex ante approach to create threat intelligence. In general, CTI generation iterates through a number of stages from raw data to actionable CTI.

In organizations CTI is positioned within the wider cybersecurity ecosystem. Security Operations Centers (SOC) commonly integrate threat intelligence and organizational processes [17]. As threat intelligence is about security incidents, it pertains in particular to incident response activities [18]. Tailored and applied to systems as well as decision makers, CTI supports the security posture.

<sup>1</sup><https://www.automationml.org/>

<sup>2</sup><https://www.json.org/>

### III. RELATED WORK

In this section we emphasize on influential works for our research. To the best of our knowledge, digital twins have not yet been used in connection with CTI.

Initial work on security and digital twins by Eckhart and Ekelhart [4] describes the potential of security simulations. On step further, Dietz et al. [8] propose a SOC/SIEM strategy. We also focus on the simulation mode but specifically show how to access the simulation output and transform it into CTI.

The digital twin paradigm is often employed in industrial domains [19], thus works tackling the assessment of ICS security (e.g., [20]) are to be considered as close research areas. Our work enhances ICS security assessment in the following ways: Foremost, we apply a novel industrial technology, the digital twin, to gather security data on ICS. Instead of providing an assessment method, we utilize CTI to structure the data and gather knowledge on ICS security.

Shared CTI must first be produced. CTI generation is assumed to be conducted by security analysts and supported by security tools. The essential activities of threat hunting and discovery have been formally modeled and analyzed with graph-based approaches [21]. Only a few works address the systematic challenge of information structuring with CTI formats. Sadique et al. [22] propose an approach to perform structured threat information generation from network log data using STIX. In comparison, we go beyond low-level CTI and show how to relate objects with relevant higher-level CTI. Besides the transformation, it is necessary to consider process elements for selecting CTI representations.

A CTI generation framework is developed by Landauer et al. [23]. The focus element of their work is anomaly detection within web server log data to derive alarms and attack patterns. Additionally, other works extracted CTI from human-readable threat reports [24], [25]. The specific CTI elements concern IoCs and attack patterns. However, this leaves a gap with regard to threat intelligence from digital twin simulations which can complement existing CTI.

### IV. FRAMEWORK

Our process-based framework demonstrates how digital twins are used to generate threat intelligence. Therefore, we combine a general CTI process with the digital twin and its security simulation capabilities. Sharing the resulting CTI can assist the collaborative security of organizations.

Figure 2 shows our framework in BPMN notation. It starts with a Security Operations Center (SOC) planning the simulation of an incident by using a digital twin. The process ends with generated CTI. Our angle is taken from the data generated by the digital twin – meaning the digital twin as the initial data source is considered in the CTI process steps. Adaption of CTI activities is thus necessary to enable digital twin integration. Sections IV-A to IV-F explain the framework’s activities in detail.

#### A. SOC: Define Incident Parameters

Before security simulation with the digital twin can be performed, the SOC needs to define realistic incident parameters  $\hat{P}_{Inc}$  for the incident simulation. We consider any violation of confidentiality, integrity or availability a security incident. Any method resulting in such violation is considered an attack.

Defining incident parameters may include to determine whether to test system weaknesses, to simulate an attack or else. Moreover, it needs to be specified which system weakness is tested, e.g., by considering CVE<sup>3</sup>-listed vulnerabilities. If the simulation of an attack is chosen, the tactics, techniques and procedures (TTP) of the attacker as well as the entry point to the system need to be set. For instance, the SOC could define the attacker to carry out a DoS-attack. This could also mean to simulate a series of attacks subsequently or at once. Documented attacks (e.g., the WannaCry ransomware-attack [26]) or common sources (e.g., MITRE ATT&CK framework) can provide realistic parameters.

#### B. Digital Twin: Simulate Incident

The security incident simulation is carried out by the digital twin. The simulation mode presents this work’s *security measure*. Typically, different technologies (e.g., Mininet<sup>4</sup> and MiniCPS<sup>5</sup>) are used to simulate systems and their networks [6], [8].

In order to build the simulation, the twin relies on specific *data*. This data  $S_{Sys}$  may consist of specification data as well as historical and state data of the physical counterpart. The simulation can then be build, leading to the general specification of the simulated system  $\hat{S}_{Sys}$ . With the simulated system and the incident parameters defined by the SOC, the digital twin security simulation  $f_{Sim}(Inc)$  can be performed. The output  $O_{Sim}(Inc)$  manifests as incident raw data (e.g., system logs or captured network traffic).

The collaboration of industry experts and simulation developers can be used to safeguard sufficient *fidelity*. Also, the SOC’s knowledge of security operations ensures the determination of high-quality and high-fidelity parameters for the security simulation and its output.

#### C. SOC: Filter State Data

From an operational standpoint, the output of security incident simulation also features less relevant data. To gather the relevant excerpt from the simulation output, filter criteria are applied. In this regard, data processing can be narrowed down to a subset of the original simulation output. SOC personnel can resort to three archetypes of filtering:

- **Time** – filtering data based on time properties (e.g., timestamps)
- **Target** – filtering data based on target devices (e.g., IP addresses)
- **Type** – filtering data based on internal classification (e.g., network packets)

<sup>3</sup><https://cve.mitre.org/>

<sup>4</sup><http://mininet.org/>

<sup>5</sup><https://github.com/scy-phy/minicps>

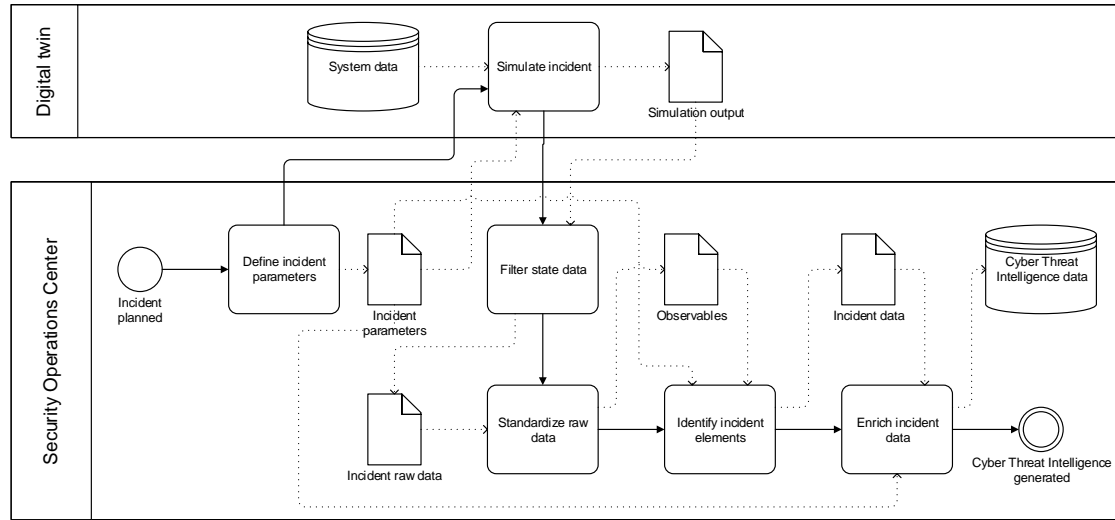


Fig. 2. Process-based framework for structured CTI generation based on digital twin security simulations

The filtered output in form of log files or network traffic captures is then forwarded to perform standardization tasks.

#### D. SOC: Standardize Raw Data

Standardization is a key process step towards structured threat intelligence. Here and in the subsequent process steps, CTI logic and STIX2.1 elements are present.

Within STIX2.1 a total of 18 SCOs address technical, low-level CTI. While in theory all SCOs are used for standardization, preselection of SCOs dependent on two generic categories of simulated incidents – network (e.g., *IPv4 Address* or *Network Traffic*) and host (e.g., *File* or *Process*) – fosters efficient processing within a given context. Considerations prior to standardization must take all SCOs into account but can exclude certain SCOs. We focus mainly on network-centric SCOs as our incident simulation is centered on network-based attacks. Standardized simulation output results in a number of SCOs. These SCOs represent valid observables.

#### E. SOC: Identify Incident Elements

Up to this point only low-level SCOs have been identified and standardized. To reach a complete representation of the simulated incident it is necessary to further identify high-level incident elements expressed as SDOs. For this purpose, the SOC can make use of four data sources that serve as input for Algorithm 1:

- 1) A list of standardized SCOs (i.e., the output of the previous process step)
- 2) A complete STIX2.1 relationship list retrieved from the STIX2.1 specification including both direct and embedded relationships

- 3) Formal incident parameters  $\hat{P}_{Inc}$  defined within the first step of the process framework
- 4) A user defined threshold value specifying the maximum number of degrees of relationship

The identification of incident elements follows a bottom-up approach from SCOs to SDOs and onward to other SDOs. It is the goal of Algorithm 1 to establish relevant relationships between instantiations of STIX2.1 objects. As the algorithm proceeds in multiple rounds, we name relationships according to the round they have been identified. For instance, first degree relationships always refer to relationships with at least one SCO involved while second degree involve at least one SDO and higher relationships always refer to relationships between SDO and SDO. Subsequently, it takes two relationships to reach a SCO from any given second degree SDO.

Algorithm 1 has two parts. The function *Identify* describes the procedure to identify incident elements. A given list of STIX objects (i.e., SCOs or SDOs) is iterated and mapped against possible relationships with SDOs. Direct relationships between STIX2.1 objects are then assessed for applicability based on formal incident parameters. The assessment by SOC personnel determines whether the relationship is relevant for the given simulation. In the case of applicability the relationship's SDO and a new SRO are instantiated and added to the respective lists. For embedded relationships the approach differs in referencing the given STIX2.1 object within the SCO or SDO and then appending only a new SDO.

Besides the function *Identify*, Algorithm 1 highlights the identification of multiple degrees of relationships between STIX2.1 objects. First degree relationships are established

**Algorithm 1:** Identify Incident Elements

---

**Data:** SCO list, STIX2.1 relationship list,  $\hat{P}_{Inc}$ , user threshold

**Result:** SCO list, SDO list, SRO list  
initialization;  
 $degree\ of\ relationship = 1$ ;

**Function**  $Identify(STIX2.1\ object\ list)$

```

foreach element in STIX2.1 object list do
  drel  $\leftarrow$  find direct relationships;
  foreach direct relationship in drel do
    assess applicability based on  $\hat{P}_{Inc}$ ;
    if applicable then
      append SDO to SDO list;
      append SRO to SRO list;
    else
      continue;
  erel  $\leftarrow$  find embedded relationships;
  foreach embedded relationship in erel do
    assess applicability based on  $\hat{P}_{Inc}$ ;
    if applicable then
      reference element in SCO/SDO;
      if SDO then
        append SDO to SDO list;
      else
        continue;
    else
      continue;
while degree of relationship < user threshold do
  if degree of relationship = 1 then
    |  $Identify(SCO\ list)$ 
  else
    |  $Identify(SDO\ list)$ 

```

---

using the SCO list as input. Afterwards, the algorithm proceeds to identify further SDOs based on a given SDO list up until a threshold for the degree of relationships is reached. This threshold value is set by the user and based on the simulation.

The identification of incident elements results in structured incident data. In the next step SDOs are complemented.

**F. SOC: Enrich Incident Data**

Having identified and listed all relevant incident elements and their relationships provides guidance for security experts. However, this information can be further enriched. CTI aims to represent security incidents and attacks in detail. This is achieved by storing information in dedicated object properties.

SDOs deemed relevant by the previous process step are first augmented with information obtained from formal incident parameters defined early on. Here, the specific attributes of the various types of SDOs are filled with contextual values. Exemplary, the use of the STIX2.1 patterning language and the *Indicator* SDO allow to capture attacker behavior.

External references then introduce enumerations and other types of standardized knowledge. For instance, the *Attack Pattern* SDO can refer to information within the *Common Attack Pattern Enumeration and Classification (CAPEC)*. Other

relevant information can be drawn from enumerations (e.g., CVE) and frameworks (e.g., cyber kill chains).

The final output of the digital-twin based Cyber Threat Intelligence generation framework comprises low-level SCOs, high-level SDOs and relationships expressed by SROs and embedded references. These CTI artifacts document the simulated security incident but also build the foundation to derive appropriate incident response measures. This is often followed by sharing CTI with others. Having stored the generated CTI the process ends.

**V. EVALUATION**

To prove our framework’s applicability, we chose technologies in Section V-A. Section V-B defines a use case of digital twin security simulation. Then, Section V-C shows how the simulation output is further processed to generate CTI. Finally, Section V-D presents the results of the evaluation.

**A. Experimental setup**

*Digital twin security simulation.* We selected MiniCPS for ICS simulation. Developed in academia and based on Mininet, it is specially tailored to ICS and incident simulation [27]. MiniCPS is able to simulate ICS network communication using Ethernet/IP (ENIP) or Modbus protocol. Network topologies can include programmable logic controllers (PLCs), human machine interfaces (HMIs), and supervisory control and data acquisition (SCADA) systems. Physical components (e.g., tanks), sensors and actuators can also be modeled. The ICS components’ logic and the underlying physical processes can be implemented abstractly. MiniCPS provides two APIs: a network layer API for mimicking network traffic and a physical layer API to represent physical events.

*CTI generation.* The STIX2.1 specification [15] provides the initial information on how to generate CTI. At various points the specification explicitly details relationships intended to connect observables and other types of CTI. Our experimental setup includes a list of all explicitly mentioned relationships extracted from the STIX2.1 specification. In total 463 relationships have been identified. Most notably, not only direct relationships between STIX2.1 objects but also embedded relationships exist and must be assessed for applicability. We import the list of all STIX2.1 relationships in Python for later processing. For CTI standardization and generation, the STIX2 Python library<sup>6</sup> is used.

**B. Use case**

We implement a use case and digital twin setup to evaluate our framework. The selected use case has been reviewed by industrial experts and considers a digital twin of a conveyor belt. Figure 3 illustrates the system topology.

The conveyor belt is driven by a motor. The belt’s velocity is measured by a sensor. Another sensor measures the temperature of the motor. The PLC and the HMI control the conveyor

<sup>6</sup><https://www.github.com/oasis-open/cti-python-stix2>

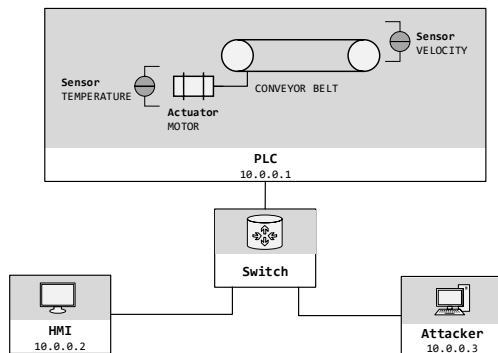


Fig. 3. Use Case – Conveyor belt

belt. They can alter the velocity of the motor, and turn it on and off. The PLC checks whether the temperature and velocity thresholds are adhered to. If so, the belt can operate as usual, otherwise the PLC turns off the motor. The HMI can receive temperature and velocity values from the PLC and send the desired velocity to the PLC.

*Define incident parameters.* A third party with malicious intent enters the network. The attacker aims at over-flooding the PLC with connection requests without finishing them (SYN flood DoS-attack). To reach this goal, the attacker utilizes the `hping3`<sup>7</sup> module hiding the attacker IP address and using the HMI IP address for disguise.

*Simulate incident.* The network topology and physical process, as illustrated in Figure 3, are realized with `MiniCPS`. ENIP is the protocol used for network communication. Further, the logical parts (PLC, attacker) are implemented. Once the attacker logic is started, the PLC gets over-flooded with SYN requests so that the HMI does not receive any responses from the PLC. The simulation produces the corresponding system logs of the PLC and the HMI as well as the captured network traffic. The code for this use case, exemplary system logs and an excerpt from the captured network traffic can be found at GitHub<sup>8</sup>.

### C. Data processing

Data processing for the use case is performed based on our prototypical implementation of general CTI generation utilities in `Python`. These utilities, described in the following, are one attempt to support the process-based framework for structured CTI generation and can be found at GitHub<sup>9</sup>.

*Filter state data.* First, to realize filtering of simulation output both log files and network traffic (pcap frames) are converted. This is necessary for further data handling and thus we implement a `Python` class for each output type. Upon

<sup>7</sup><https://tools.kali.org/information-gathering/hping3>

<sup>8</sup><https://github.com/FrauThes/DigitalTwin-ConveyorBelt>

<sup>9</sup><https://www.github.com/digitaltwinCTI/CTI-DT-utilities>

conversion, implemented functions allow security analysts to get familiar with the simulation output. Through the use of filter functions the simulation output is lastly reduced to a relevant excerpt forming the input data for standardization.

*Standardize raw data.* Standardization of the filtered raw data is based on the security analyst's input. Therefore, we provide the functionality to first show host-based, network-based or all available SCOs of the STIX2.1 standard. This is followed by the option to create a custom list with relevant SCOs which can then be used for standardization and identification of SDOs. Further, we provide functions to extract information from log file entries or pcap frames and storing it in pre-defined SCO `Python` objects.

*Identify incident elements.* We implement a search function to query related objects explicitly mentioned by the STIX2.1 standard. These SCOs or SDOs and associated relationships can then be selected and instantiated by a security analyst with contextual information of the simulated incident. In this regard the Algorithm 1 provides the procedural guidance but relies on external input (i.e.,  $\hat{P}_{Inc}$ ) for applicability assessment and content.

*Enrich incident data.* The final step within the framework is enriching STIX2.1 objects with external CTI. Here, a security analyst can modify existing SCO or SDO `Python` objects and add relevant data. In our use case, suitable descriptions and CAPEC IDs are added to the *attack pattern* SDO.

### D. Experimental results

The experimental results show that the output of digital twin simulations in form of log data or network traffic is extensive and must be reduced to a relevant excerpt. To facilitate the processing, we provide a search functionality.

The multitude of possible CTI objects and relationships in the STIX2.1 specification can be used for description. As we extracted all explicitly mentioned relationships, the workload is reduced. However, the assessment for applicability through a security analyst still requires detailed knowledge of the standard and can be perceived as cumbersome. In particular, for SCOs and lower degree relationships multiple options demand interpretation. Figure 4 shows an aggregated view on possible relationships. We highlight the STIX2.1 objects identified for the use case by a bold frame. With higher transparency we further emphasize higher degrees of relationship between the individual objects.

By analyzing the resulting log data from the DoS-SYN flooding-attack simulation, the SCOs *IPv4 address* and *process* are extracted: `main loop` refers to the process and the IP addresses of HMI and PLC represent the corresponding IP addresses in the network. In a next step, the network traffic is analyzed which yields the SCOs *MAC address* and *network traffic*. This documents the communication between the components of the digital twin. In addition, the attacker MAC address `00:00:00:00:00:03` is present.

In the first search for relationships, we identify embedded relationships between SCOs (e.g., *IPv4 or MAC addresses* and

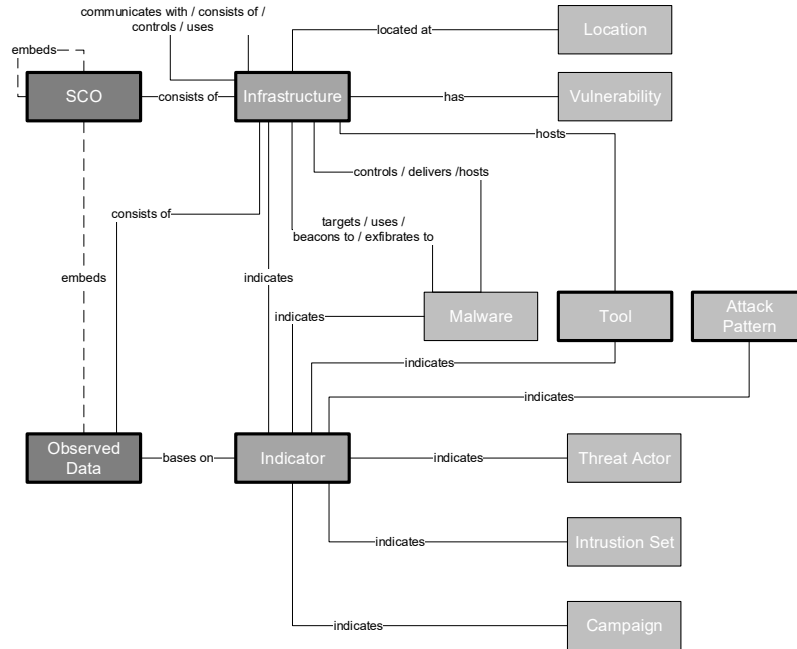


Fig. 4. Aggregated view on relationships (results are framed bold)

network traffic) as well as direct relationships of the previously identified IP addresses and the process to the *infrastructure* SDO. Additionally, network traffic is connected to the *observed data* SDO. Most notably, the `number_observed` attribute indicates a high number of traffic. Also, one existing IP address object is updated as it resolves to two MAC addresses.

Searching for second degree and higher relationships, we identify the *indicator* SDO based on the highly repetitive amount of network traffic originating from multiple ports. Another indicator objects centers on the observed duplicate IP address assignment. Onward, we use the *attack pattern* SDO to describe the context of the simulated attack. Characteristics about `hping3` are consequently connected and stored within a *tool* SDO. At last, the generated STIX2.1 objects are connected to a report object and bundled.

We want to point out that *location* and *threat actor* represent two SDOs not applicable to simulations and CTI generation based on digital twins. This is because a real threat actor, its motivations and its geolocation are not present in simulations. We provide an example visualization for the resulting CTI report of our use case in the Appendix B.

## VI. DISCUSSION

In the following, we reason about the significance of the experimental results, the employment of digital twins and the systematic CTI generation.

We constructed our use case and attack simulation close to possible real-world incidents and ICS. In our use case, the network layer of the conveyor belt scenario provides the same protocols and network packages as used in the real-world. As the simulated DoS-attack affects the network-layer, the results remain reliable. Still, we want to mention that digital twins can never present all aspects of their real-world counterpart. The benefit of conducting digital twin security simulation remains and is documented by our DoS-attack scenario. Attacking the ICS would result in halting the conveyor belt and might cause direct financial losses. Derived insights from the resulting CTI reports can ultimately lead to adequate system protection regardless of prior attack occurrence. Especially, since the traffic between industrial components is often not encrypted the implications for other possible attack scenarios are far reaching.

The intention behind using simulation data should not lie in increasing the amount of CTI but in providing more comprehensive and shareable threat information to foster security. We argue that as long as sufficient fidelity of the digital twin simulations is provided, the use of digital twin simulation data is beneficial. Moreover, this way, countermeasures can be tested.

Currently few works tackle digital twin security simulations [6], [8]. Thus, benchmark comparisons or performance evaluations do not exist. Our use case is based on input from industrial experts. In simulation mode, we neglect traffic



between real-world system and digital twin. Nevertheless, the digital twin has to be sufficiently secured to uphold the advantages of security simulations.

The STIX2.1 format provides a comprehensive approach to generate structured threat information and allows for vendor-independent CTI sharing. Despite its threat reporting roots, CTI goes beyond being purely informative and can also serve as transition to security technologies. Multiple security technologies (e.g., IDS) rely on some form of structured data, which CTI can provide. CTI generation typically follows a bottom-up approach. Starting with large amounts of simulation output it becomes apparent that STIX2.1 is not intended for storing entire outputs as the number of possible relationships rapidly increases. Instead, the STIX2.1 objective is the description of relevant incident and attack elements. Filtering, foreseen by our framework, and other technologies (e.g., log management) assist this objective. Systematic CTI generation is then based on the initial selection of suitable SCOs. Relationships connect these and lead from observed data to indicators and higher level SDOs. Whereas custom STIX2.1 relationships allow for even further CTI structuring options, we expect that for most digital twin simulations explicit relationships provide sufficient semantic expressiveness.

Organizations sharing CTI can harness digital twin security simulations. That way, they have a method to reconstruct cybersecurity attacks closely related to their systems and networks. Yet, they can abstract confidential aspects of security incidents and generate shareable CTI reports. Other organizations receiving these CTI reports have a structured attack description ready for further processing. For instance, visualization techniques can make relevant attack aspects visible and lead to collaborative cybersecurity knowledge.

## VII. CONCLUSION

In this work, we propose a framework to generate CTI from digital twin simulations. Thereto, a digital twin use case is implemented. We show that digital twins can simulate security incidents and demonstrate successive steps using the STIX2.1 standard. Moreover, real-world data is no longer obligatory as output from security simulations provides a novel CTI source. To facilitate the processing of the simulation output, we also created CTI generation utilities. The final result is an exemplary CTI report to be shared and acted upon by other organizations.

Generally, the creation of digital twins is a complex task and organizations might put emphasis on different characteristics. Thus, not every digital twin might be suitable for security simulations as it is strongly dependent on whether the digital twin contains characteristics an attack is based on. The systematic generation of CTI based on the STIX2.1 specification limits ambiguities, yet provides a multitude of semantic options for representation. Due to this, it is beyond the scope of this paper to automate the CTI generation process. Closer consideration and integration of detection capabilities found in security systems might be a feasible addition.

To improve evaluation of digital twins, future research should define quality criteria for digital twins and their usage in security so that benchmark studies can be realized. To offer more and advanced attack scenarios, different high-fidelity domain-models (e.g., physical, logical, network) of a system might be merged into one digital twin simulation. Additionally, a combination of digital twins and deception technologies (e.g., honeypots) might provide input for CTI reports on attacks the SOC is currently not aware of. For systematic CTI generation and the usage of the STIX2.1 format, aspects of recommender systems are worth investigating. Comparing objects and relationship paths could assist the applicability assessment and lead to the most helpful CTI report generation.

## APPENDIX

### A. STIX2.1 JSON representation

The *Observed Data* object shown in Listing 1 is an exemplary STIX2.1 domain object and contains, among others, the attributes `type`, `id` and `created` (common), `number_observed` (specific) as well as `object_refs` (optional). Furthermore, an embedded relationship connecting the *Observed Data* SDO with an *IPv4 Address* SCO is realized by the `object_refs` attribute.

```
{
  "type": "observed-data",
  "id": "observed-data--1",
  "created": "2021-12-25T19:58:16.000Z",
  "modified": "2021-12-25T19:58:16.000Z",
  "first_observed": "2021-11-21T19:00:00Z",
  "last_observed": "2021-11-21T19:10:00Z",
  "number_observed": 21,
  "object_refs": [
    "ipv4-addr--1"
  ]
}, {
  "type": "ipv4-addr",
  "id": "ipv4-addr--1",
  "value": "213.31.107.1"
}
```

Listing 1. Exemplary STIX 2.1 SDO with embedded relationship and SCO

### B. Threat report visualization

The CTI report of the attack simulation conducted in the digital twin use case can be further processed by common STIX visualization tools<sup>10</sup>. Figure 5 displays the different STIX2.1 objects and their relationships. At its core, network traffic objects are linked to originating MAC addresses. Indicators emphasize the network traffic and IP address 10.0.0.2 resolving two MAC addresses.

## ACKNOWLEDGMENT

This work is partly performed under the ZIM SisseC project, which is supported under contract by the German Federal Ministry for Economic Affairs and Energy (16KN085725). Part of this research is further supported by the German Federal Ministry of Education and Research under the BMBF DEVISE project.

<sup>10</sup><https://oasis-open.github.io/cti-stix-visualization/>

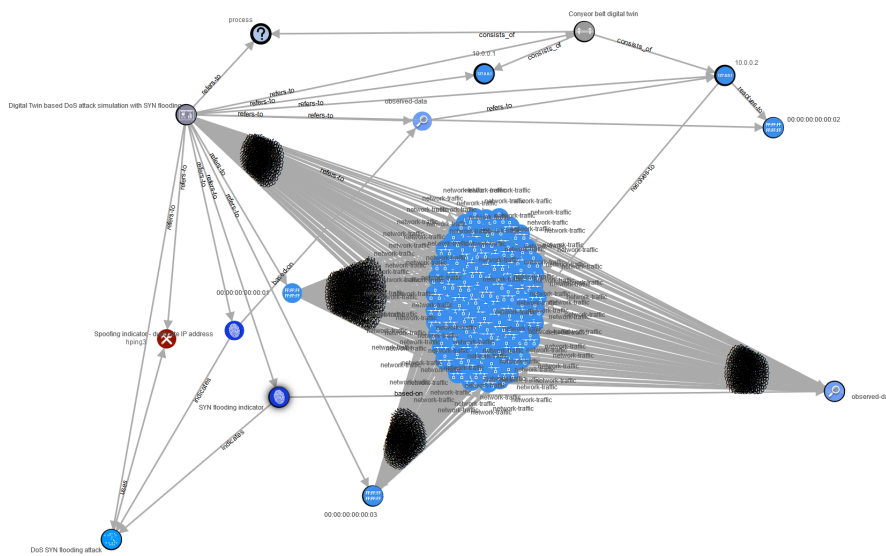


Fig. 5. Visualization of a CTI report for the DoS-attack use case

REFERENCES

[1] W. Tounsi and H. Rais, "A survey on technical threat intelligence in the age of sophisticated cyber attacks," *Computers & Security*, vol. 72, pp. 212–233, 2018.

[2] P. Kieseberg and E. Weippl, "Security challenges in cyber-physical production systems," in *Software Quality: Methods and Tools for Better Software and Systems*, 2018, pp. 3–16.

[3] G. Lumer-Klabbers, J. O. Hausted, J. L. Kvistgaard, H. D. Macedo, M. Frasher, and P. G. Larsen, "Towards a digital twin framework for autonomous robots," in *2021 IEEE 45th Annual Computers, Software, and Applications Conference (COMPSAC)*, 2021, pp. 1254–1259.

[4] M. Eckhart and A. Ekelhart, *Digital Twins for Cyber-Physical Systems Security: State of the Art and Outlook*, 2019, pp. 383–412.

[5] M. Dietz and G. Pernul, "Unleashing the digital twin's potential for ics security," *IEEE Security Privacy*, vol. 18, no. 4, pp. 20–27, 2020.

[6] M. Eckhart and A. Ekelhart, "Towards Security-Aware Virtual Environments for Digital Twins," in *Proceedings of the 4th ACM Workshop on Cyber-Physical System Security (CPSS '18)*, 2018, pp. 61–72.

[7] J. E. Rubio, C. Alcaraz, R. Roman, and J. Lopez, "Current cyber-defense trends in industrial control systems," *Computers & Security*, vol. 87, p. 101561, 2019.

[8] M. Dietz, M. Vielberth, and G. Pernul, "Integrating digital twin security simulations in the security operations center," in *Proc. of the 15th Int. Conference on Availability, Reliability and Security*, 2020, pp. 1–9.

[9] S. Boschert, C. Heinrich, and R. Rosen, "Next Generation Digital Twin," in *Proceedings of the 12th International Symposium on Tools and Methods of Competitive Engineering*, 2018, pp. 209–217.

[10] M. Dietz and G. Pernul, "Digital Twin: Empowering Enterprises Towards a System-of-Systems Approach," *Business & Information Systems Engineering*, vol. 62, no. 2, p. 179–184, 2020.

[11] G. N. Schroeder, C. Steinmetz, C. E. Pereira, and D. B. Espindola, "Digital twin data modeling with automationml and a communication methodology for data exchange," *IFAC*, vol. 49, no. 30, pp. 12 – 17, 2016, 4th IFAC Symposium on Telematics Applications.

[12] M. Grieves and J. Vickers, *Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems*, 2017, pp. 85–113.

[13] R. Brown and R. M. Lee, "The evolution of cyber threat intelligence (cti): 2019 sans cti survey," *SANS Institute*, 2019.

[14] F. Menges and G. Pernul, "A comparative analysis of incident reporting formats," *Computers & Security*, vol. 73, pp. 87–101, 2018.

[15] OASIS Cyber Threat Intelligence (CTI) TC, "Stix™ version 2.1: Committee specification 01," 2020.

[16] D. Schlette, F. Böhm, M. Caselli, and G. Pernul, "Measuring and visualizing cyber threat intelligence quality," *International Journal of Information Security*, vol. 20, no. 1, pp. 21–38, 2021.

[17] M. Vielberth, F. Böhm, I. Fichtinger, and G. Pernul, "Security operations center: A systematic study and open challenges," *IEEE Access*, vol. 8, pp. 227 756–227 779, 2020.

[18] D. Schlette, M. Caselli, and G. Pernul, "A comparative study on cyber threat intelligence: The security incident response perspective," *IEEE Communications Surveys & Tutorials*, 2021.

[19] E. Negri, L. Fumagalli, and M. Macchi, "A Review of the Roles of Digital Twin in CPS-based Production Systems," *Procedia Manufacturing*, vol. 11, pp. 939–948, 2017.

[20] Y. Cherdantseva, P. Burnap, A. Blyth, P. Eden, K. Jones, H. Soulsby, and K. Stoddart, "A review of cyber security risk assessment methods for scada systems," *Computers & Security*, vol. 56, pp. 1 – 27, 2016.

[21] S. M. Milajerdi, B. Eshete, R. Gjomemo, and V. Venkatakrishnan, "Poirot: Aligning attack behavior with kernel audit records for cyber threat hunting," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019, p. 1795–1812.

[22] F. Sadique, S. Cheung, I. Vakilinia, S. Badsha, and S. Sengupta, "Automated structured threat information expression (STIX) document generation with privacy preservation," in *9th IEEE Annual Ubiquitous Computing, Electronics & Mobile Comm. Conf.*, 2018, pp. 847–853.

[23] M. Landauer, F. Skopik, M. Wurzenberger, W. Hotwagner, and A. Rauber, "A framework for cyber threat intelligence extraction from raw log data," in *IEEE Int. Conf. on Big Data*, 2019, pp. 3200–3209.

[24] X. Liao, K. Yuan, X. Wang, Z. Li, L. Xing, and R. Beyah, "Acing the ioc game: Toward automatic discovery and analysis of open-source cyber threat intelligence," in *Proc. of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 755–766.

[25] V. Legoy, M. Caselli, C. Seifert, and A. Peter, "Automated retrieval of ATT&CK tactics and techniques for cyber threat reports," *arXiv*, 2020.

[26] D. Kao and S. Hsiao, "The dynamic analysis of wannacry ransomware," in *20th Int. Conf. on Adv. Comm. Tech. (ICACT)*, 2018, pp. 159–166.

[27] D. Antonioli and N. O. Tippenhauer, "Minicps: A toolkit for security research on cps networks," in *Proc. of the 1st ACM Workshop on Cyber-Physical Systems-Security and/or Privacy*, 2015, pp. 91–100.

## 5 Enhancing Industrial Control System Forensics Using Replication-based Digital Twins

---

Current status:	Published
Conference:	Digital Forensics - 17th IFIP WG 11.9 International Conference, Virtual Event, USA, February 01-02, 2021
Date of acceptance:	November 24, 2020
Full citation:	DIETZ, M., ENGLBRECHT, L. AND PERNUL, G. Enhancing Industrial Control System Forensics Using Replication-based Digital Twins. In <i>Advances in Digital Forensics XVII, IFIP Advances in Information and Communication Technology</i> , vol. 612. Springer, Cham (2021), pp. 21-38.
Authors contributions:	Marietheres Dietz 45% Ludwig Englbrecht 45% Günther Pernul 10%

---

**Conference Description:** The annual IFIP WG 11.9 International Conference on Digital Forensics provides a forum for presenting original research results and innovative ideas related to the extraction, analysis and preservation of all forms of electronic evidence. The conference proceedings are published as an edited volume in the well-known Research Advances in Digital Forensics book in the IFIP Advances in Information and Communication Technology (AICT) series by Springer.



## Chapter 2

# ENHANCING INDUSTRIAL CONTROL SYSTEM FORENSICS USING REPLICATION-BASED DIGITAL TWINS

Marietheres Dietz, Ludwig Englbrecht and Günther Pernul

**Abstract** Industrial control systems are increasingly targeted by cyber attacks. However, it is difficult to conduct forensic investigations of industrial control systems because taking them offline is often infeasible or expensive. An attractive option is to conduct a forensic investigation of a digital twin of an industrial control system. This chapter demonstrates how a forensic investigation can be performed using a replication-based digital twin. A digital twin also makes it possible to select the appropriate tools for evidence acquisition and analysis before interacting with the real system. The approach is evaluated using a prototype implementation.

**Keywords:** Digital forensics, industrial control systems, digital twins

## 1. Introduction

Industrial control systems have long life-spans. Since maintenance is performed only a few times a year, industrial control system firmware and software are updated very infrequently [19]. While safe operations are a priority for industrial control systems, adequate levels of security are generally lacking. As a result, industrial control systems are exposed to numerous threats.

When security is breached, an incident response is initiated to understand the situation, mitigate the effects, perform corrective actions and ensure safe operations. A digital forensic investigation provides the best insights into an incident and also assists in the prosecution of the perpetrators. Digital forensic readiness is essential to maximize the ability

to acquire useful evidence and minimize the costs of investigations [30]. An appropriate enterprise-wide maturity level is needed to implement digital forensic readiness for information technology assets [13]. However, an appropriate maturity level is even more difficult to attain by enterprises with operational technology assets such as industrial control systems.

Conducting digital forensic investigations of industrial control systems is challenging because the systems are required to operate continuously for safety and financial reasons. Since the systems cannot be stopped to acquire evidence and conduct forensic analyses, the only alternative is to reduce the shutdown time. This can be accomplished using digital twins of the real systems to identify where evidence resides in the real systems and to select the right tools for extracting evidence before the real systems are stopped. Digital twins replicate the dynamic behavior of their real counterparts. Unlike other state-of-the-art solutions, using digital twins enable industrial control systems to continue to operate while potential attacks are being investigated. Furthermore, unlike cyber ranges and testbeds, digital twins are well suited to digital forensics due to their fidelity, flexibility and two-way communications between the real systems and their digital twins. This chapter demonstrates how forensic investigations of industrial control systems can be performed using replication-based digital twins.

## 2. Background

This section provides an overview of digital twins, digital twin security and digital forensics.

### 2.1 Digital Twin

A digital twin is a controversial term with different meanings in different domains [23]. Nevertheless, it can be regarded as a virtual representation of a real object over its lifecycle. Although digital twins have been employed in several domains, including smart cities [14], health-care [21] and product management [31], they are commonly deployed in the Industry 4.0 paradigm [23], which is the focus of this work.

According to Kritzinger et al. [20], a digital twin is distinguished from other virtual representations (e.g., digital models) by its data flow. A digital model is a manual flow from a real object to a digital object. In contrast, a digital twin has bidirectional automated data flows between the real and virtual worlds [4, 20]. Thus, the digital twin is able to gather state data from its physical counterpart. However, the twin usually contains other asset-relevant data such as specification data [4]. When

enhanced with semantics [26], this data can support various analyses, optimizations and simulations performed by the digital twin [1, 16].

## 2.2 Digital Twin Security

Several researchers have emphasized that digital twins must have adequate security [16, 26]. However, digital twins can also support industrial control system security [5, 8, 25]. Thus, digital twin security has two perspectives, securing digital twins and using digital twins to implement security. This work focuses on the second perspective and assumes that a digital twin has adequate security. The following security-centered definition of a digital twin is employed in this work [9].

*Definition 1. A digital twin is a virtual double of a system during its lifecycle, providing sufficient fidelity for security measures via the consumption of real-time data when required.*

Various digital twin modes exist to enable secure operations [5]. While a digital twin provides analytical and simulation capabilities, the replication mode, which supports the exact mirroring of a real system and its states, is relevant to this research [5]. The record and play functionality [10] is unique to the replication mode and is the essence of this work (Definition 1). While research has focused on digital twin security, little, if any, work has focused on using digital twins to support digital forensic investigations despite promising characteristics such as system state mirroring.

## 2.3 Digital Forensics

Digital forensics involves the identification, collection, preservation, validation, analysis, interpretation and presentation of digital evidence associated with an incident in a computer system [24]. The collection and analysis of digital evidence should be based on a comprehensive process model (see, e.g., [18]).

Internet of Things devices generate many traces during their operation that can be vital to digital forensic investigations. Clearly, there is a need for tools that can support digital forensic investigations of Internet of Things devices. Servida and Casey [27] discuss the challenges involved in examining Internet of Things devices. Although their work does not explicitly deal with industrial control systems, the three main challenges they present are relevant to this work. First, the computing power of the devices is very low and not suitable for performing complex tasks.

Second, most devices have limited embedded memory or an external SD card. Third, it is difficult to extract evidence due to device heterogeneity.

Digital forensic practitioners require considerable expertise, tools and time to completely and correctly reconstruct evidence given the large amounts of data to be processed. A promising approach is to use digital twins. A digital twin can be used to detect an attack. Additionally, it can provide crucial information and insights during digital forensic analysis.

Another challenge to conducting a digital forensic analysis is that the integrity of the data can be compromised during its recovery and analysis. A digital twin can assist in ensuring data integrity. The digital twin of an industrial control system can be examined and the digital forensic process and results verified before performing any actions on the real system.

### 3. Related Work

The application of digital twins to security and especially forensics has only recently drawn the attention of researchers. The concepts proposed by Eckhart et al. [7, 10] and Gehrman and Gunnarsson [15] are closely related to this research.

In general, a digital twin is a high-fidelity representation of its real counterpart. Eckhart and Ekelhart [7] were the first to study industrial control system state replication using digital twins; their focus was on reflecting the states of the real system virtually. To avoid large bandwidth overhead, Eckhart and Ekelhart [7] proposed a passive approach that identifies stimuli that alter real-world system states and reproduces them in the digital world.

Gehrman and Gunnarsson [15] demonstrated that an active approach is suitable for less complex digital twins with moderate synchronization frequencies that do not create overhead; examples are replications of a single industrial control system or an industrial plant with a few industrial control systems. They showed that synchronizing the states of a real system with a digital twin supports active replication. Gehrman and Gunnarsson also specified security requirements, established a security architecture and implemented secure synchronization between the real object and its digital twin.

According to Eckhart et al. [10], the record and play (replay) mode is a special manifestation of replication using a digital twin. Generally, a digital twin would reflect the real system states at all times. Additionally, restoring the preceding state is enabled, instead of merely replicating the current state that would be lost as soon as the subsequent

state is replicated. The replay mode supports incident management, explicitly tracks infection histories [9] and provides novel functionality with regard to forensics [5]. Therefore, replay is a vital functionality provided by replication using digital twins, especially when applied to digital forensics.

Modern industrial control systems are exposed to security threats because they incorporate commodity hardware and software and operate in highly-interconnected environments. Researchers have attempted to enhance digital forensic capabilities by providing monitoring and logging functionality [3, 33].

The proposed research differs from the research described above in a key way. If an industrial control system is compromised, then its digital twin would exhibit the same malicious behavior as its real counterpart. While this is often considered a downside of replication [7, 15], the proposed research attempts to transform it to an advantage. Specifically, after replicating the exact states of the real system in its digital twin, forensic tools can be applied to conduct deep inspections of a security incident.

Note that this approach differs from traditional intrusion detection and security incident and event management (SIEM) research by focusing on deep inspection and resolution of incidents instead of mere detection. Indeed, the objective is to create a forensically-sound and replicable baseline for forensic analyses of industrial control systems.

The proposed approach advances the state-of-the-art in several respects. Artificial environments such as Mininet can be used to reproduce stimuli (or events) that change the states of industrial control systems or are responsible for the identified activities [7]. Such environments are well-suited to simulating systems and their events [8], but the fidelity of their replication is reduced [7]. Therefore, the proposed approach includes stimuli and events directly from a real system. It is also relevant to note that, while the digital-twin-based security framework of Gehrman and Gunnarsson [15] considers security abstractly by suggesting a single security analysis component for multiple digital twins in a system, the proposed approach incorporates a separate security analysis module for each control system.

The proposed approach also advances previous work by providing a concrete definition and a proof-of-concept implementation of a digital twin environment for forensic (and/or security) analyses. Instead of investigating how digital twins can protect industrial control systems from external attacks [15], the focus is on digital forensics in the factory domain. The replication-based approach incorporates real communications between multiple control systems. The influences on the logical compo-



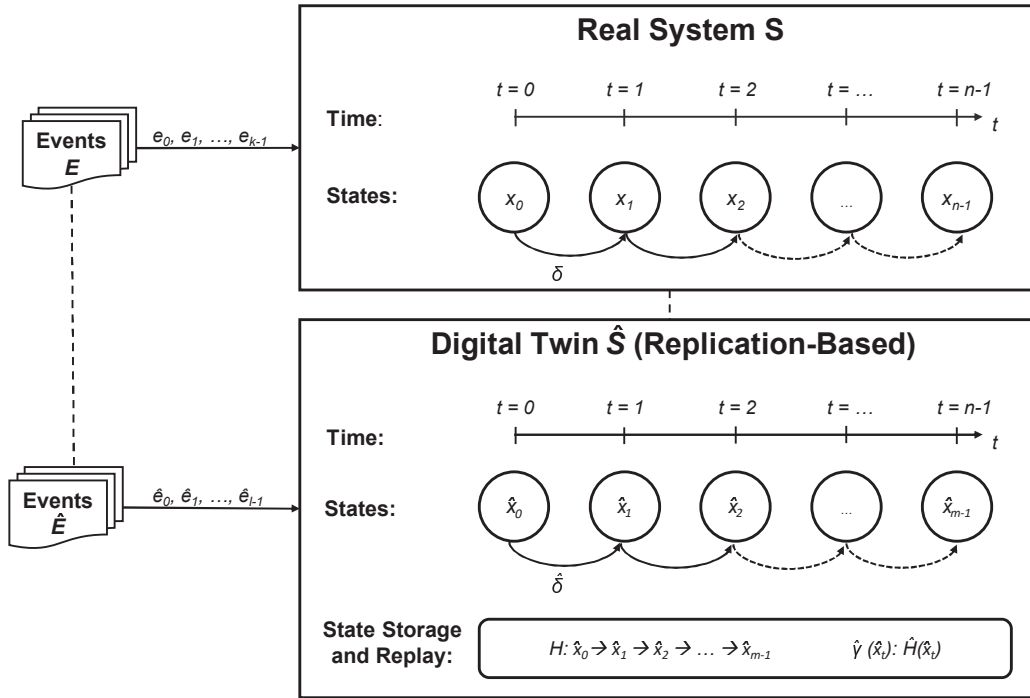


Figure 1. Replication-based digital twin with state storage and replay.

nents and filesystems of industrial control systems are also mimicked. Whereas other work relies on logging and monitoring mechanisms [3, 33], the proposed approach focuses on filesystem changes by making periodic recordings of content in digital twins. This enables the replaying of all the evidence generated by the replicated digital twins.

#### 4. Replication Using Digital Twins

This section provides the theoretical foundations for the proposed approach using digital twins. Four theorems formalize the requirements for replication-based digital twins with replay capabilities.

##### 4.1 Replication and Replay Theorems

Definition 1 above provides the security-centered definition of a digital twin [9]. Formal notions pertaining to state replication with digital twins are provided in [7, 15]. These notions and additional definitions create the basis for digital forensic analyses. Figure 1 presents the proposed replication-based digital twin with replay functionality that can be used for digital forensic analyses.

Sufficient fidelity of digital twins in the replication mode is required to support analyses. This concept is formalized by Theorems 1, 2 and 3.

*Theorem 1 (Representation of States).* The finite set of states  $X = \{x_0, x_1, \dots, x_{n-1}\}$  of a real system is represented in its replication-based digital twin as  $\hat{X} = \{\hat{x}_0, \hat{x}_1, \dots, \hat{x}_{m-1}\}$ . A high-fidelity digital twin is replicated as a subset of the real system corresponding to  $\hat{X} \subseteq X$  where  $m \leq n$ . In an ideal digital twin,  $\hat{X} = X$ .

*Theorem 2 (Timely Orderliness).* To replicate a real system accurately, the concept of time has to be considered. Let  $x_t \in X_t$  represent the real system at time  $t$  where the initial state is  $x_0$ . The digital twin replicates each state  $\hat{x}_t$  in chronological order so that  $x_0 < x_1 < x_2 < \dots < x_{n-1}$ . Time delays may occur between the real system states and digital twin states, but they do not affect digital forensics much because forensic investigations are typically conducted post mortem.

In addition to having sufficient fidelity, a digital twin must be able to consume real-time data and replay it. This motivates Theorem 3.

*Theorem 3 (Integration of Events.)* If a system changes from one state to another, certain input data is required, which is referred to as events. Events might occur due to the inner workings of the system or due to its external environment that may not be covered by its digital twin. For example, commands from the system's program are internal events whereas network traffic from other systems are external events. Real events  $E = \{e_0, e_1, \dots, e_{k-1}\}$  and the events replicated in the digital twin  $\hat{E} = \{\hat{e}_0, \hat{e}_1, \dots, \hat{e}_{l-1}\}$  express these inputs, where  $\hat{E} \subseteq E$  and  $l \leq k$ . Furthermore, the transition function  $\delta$  expresses the changes of states in the real system:  $\delta : X \cdot E \rightarrow X$ , i.e.,  $x_{t+1} = \delta(x_t, e_t)$ . Likewise,  $\hat{\delta}$  expresses the changes of states in its digital twin. Events lead to state changes that in turn may leave traces such as new files or updates of internal values in the real system. As a result of replication, the same traces will be found in the digital twin.

The highly-desirable replication-based replay functionality imposes additional requirements. This motivates a fourth theorem.

*Theorem 4 (Accuracy in Replay).* The replay function resets a digital twin to a starting state. Deviations from the previously-observed states of the digital twin should not occur when retrieving its historical events [9]. First, the transition function leading to a subsequent state  $x'$  is repli-

cated to achieve similar states in the digital twin:  $\delta(x, e) = \hat{\delta}(\hat{x}, \hat{e}) \iff x' = \hat{x}'$ . Thus, starting with the initial state, a chain of historic states  $\hat{H} : \hat{x}_0 \mapsto \hat{x}_1 \mapsto \hat{x}_2 \mapsto \dots \mapsto \hat{x}_n$  can be constructed. This chain can be used to reset states and replay the subsequent states in chronological order. The replay function  $\hat{\gamma}(\hat{x}_t) : \hat{H}(\hat{x}_t)$  expresses a reset to state  $\hat{x}_t$  and the traversing of the states in chronological order.

Finally, the real system is a deterministic system defined by tuples  $S := (X, x_0, E, \delta)$ . The digital twin is represented similarly by incorporating state storage and replay functionality  $\hat{S} := (\hat{X}, \hat{x}_0, \hat{E}, \hat{\delta}, \hat{H}, \hat{\gamma})$ .

## 4.2 Conceptual Framework

The framework collects data (e.g., network traffic) from a real system that is imported by a high-fidelity digital twin of the real system. The replication ensures that the states of the real system are mirrored by the digital twin. Thus, the digital twin would have digital evidence traces that mirror those in the real system, enabling the digital twin to be analyzed in a forensic investigation.

Forensic investigations, however, require the recording of the current replicated states as well previous states and their associated traces. Furthermore, the states should be accountable and in chronological order (specified in Theorem 2). Therefore, the replication-based digital twin framework also incorporates state storage and replay functionality.

Figure 2 shows the replication-based digital twin framework with state storage and replay functionality. Note that the events  $\hat{E}$  in the framework may be external as well as internal.

The framework has four key building blocks: (i) data collection, (ii) digital twin replication, (iii) digital twin state storage and replay, and (iv) digital forensic analysis:

- **Data Collection:** Input data is required to replicate a real system with sufficient fidelity. The data collection component gathers the inputs (events  $\hat{E}$ ) as specified in Theorem 3. The input data may be internal (static) or external (dynamic). Internal data typically can be obtained directly from the real system, such as program code that may alter the system state or commands that are sent in response to external events. They are static because they do not change often and do not exhibit streaming characteristics. In contrast, external data typically corresponds to events that affect the real system. They often occur outside the real system, but within its environment. External data can be characterized as mainly dynamic because it can emerge at any time and in a constant manner

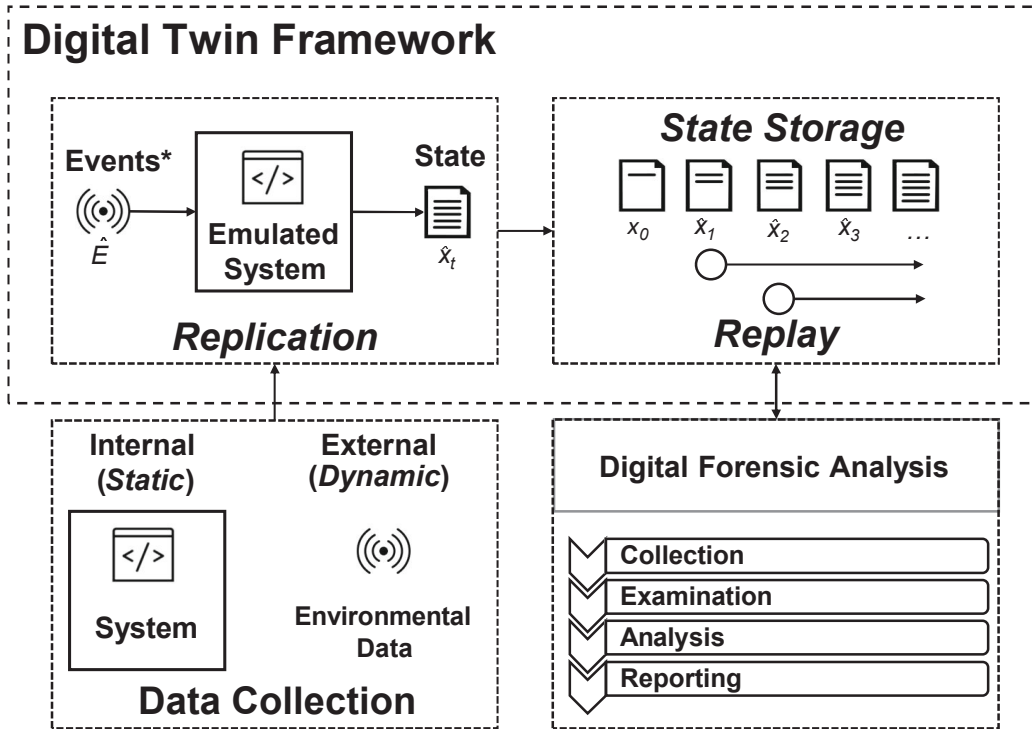


Figure 2. Digital twin framework.

(streaming characteristics). Examples of external data are network traffic from the real system’s environment and sensor values.

- Digital Twin Replication:** The collected data is used to replicate the real system as a digital twin. Internal data such as program code and system configurations are used to emulate the real system. It is important to choose the right technology for replication along with the system itself, desired degree of detail and levels of representation (e.g., network, software and operating system).

While internal events automatically occur by integrating program code and software, external events have to be input to the system. This enables the replication of system behavior based on internal events and on external stimuli. The greater the amount of data integrated, the more accurate the replication, but a trade-off should be performed between replication fidelity and cost. According to Theorems 1, 2 and 3, the states of the real system are replicated in the desired manner. In each state, different traces are created based on the transition functions. For instance, a file could be created in state *A* while its content is changed in state *B*.

- **Digital Twin State Storage and Replay:** This component is vital to digital forensic investigations and elucidation. Stage storage keeps the various system states in chronological order ( $\hat{H}$  in Theorem 4), ensuring that they remain accountable. The lateral movement of an attack can be elucidated in a step-by-step manner using system traces. It is critical that the actual content of the generated or modified system data are examined, as initiated by the recorded and replayed network traffic of the real system.

Replay relies on state storage. It enables the replication of previous states and all their subsequent states (the transition function can be deduced by considering consecutive states). The replay functionality enables the replication to be reset to a desired state and to play all the succeeding states ( $\hat{\gamma}(\hat{x}_t)$  in Theorem 4). A digital forensic practitioner can stop the replication at a state of interest, conduct a forensic analysis and continue.

- **Digital Forensic Analysis:** With state storage and replay, a forensic practitioner can gain valuable insights into previous states and processes. There are many ways to use these new opportunities in digital forensics. For example, a practitioner could reset the replication to a certain state using the replay functionality and use analysis tools in turn until a suitable tool is found. The exact replication of events and their storage in chronological order enhances the understanding of an attacker’s tactics, techniques and procedures. With each consecutive state, the pattern of the attack might become clearer to the practitioner. An attack that deletes its traces can be investigated by leveraging the replication-based approach with state storage and replay according to the digital forensic process defined by Kent et al. [18].

## 5. Implementation and Evaluation

The implementation involved a real system comprising a windmill with a programmable logic controller (PLC). The real system and its digital twin had identical Unix operating systems, ran the OpenPLC programmable logic controller software and communicated using TCP/IP.

### 5.1 Implementation and Experimental Setup

The real system was replicated by passively capturing its network traffic using `tshark` and re-transmitting the traffic to the emulated system. The configuration supports forensic analyses of the real system during execution. This section provides details of the implementation.

The implementation incorporated four main components: (i) data collection, (ii) digital twin replication, (iii) digital twin state storage and replay, and (iv) digital forensic analysis:

- **Data Collection:** The real system and its digital twin employed Unix operating systems running OpenPLC software. OpenPLC supports the IEC 61131-3 standard [17], which defines the software architecture and programming languages for programmable logic controllers. Network traffic created during the operation of the real system was recorded as a PCAP file and re-transmitted to the digital twin.
- **Digital Twin Replication:** The real system was replicated by the digital twin, which executed in a virtual environment running the same Unix operating system and OpenPLC software. All the OpenPLC variables in the real system were refactored as “memory storage” in the digital twin. This facilitated the persistent storage feature of OpenPLC whereby values at various programmable logic controller addresses were saved to disk to provide insights into their changes during programmable logic controller operation.

The open-source software Polymorph [29] was used to re-transmit network traffic. Polymorph translated the PCAP file data to templates to enable situation-aware interactions. It also facilitated dynamic integration of the components instead of the pure transmission of data. The templates were modifiable for subsequent re-transmission, which was vital because the digital twin had to respond correctly to commands issued in the real system. The digital twin operated (according to the transition function in Theorem 3) as closely as possible to the real system.

- **Digital Twin State Storage and Replay:** The storage and replay component was hosted on a virtual machine (VM), which created a storage snapshot of system state every minute. The snapshots enabled the entire virtual system to be re-created at any point in time. Since data was written and deleted during the execution of OpenPLC, the freed data area in the digital twin could be overwritten, which was problematic. Therefore, a modified version of stateless continuous data protection (CDP) software was employed. Continuous data protection is a technology that continuously captures and stores data changes, enabling data from any point in the past to be recovered [28, 34]. The software also enables the monitoring and restoration of all files generated during system execution.

The `sauegardeEx` tool [12] was created to gain insights into the generation of files during OpenPLC operation. It is based on `sauegarde`, an open-source, stateless implementation of continuous data protection [6].

Compared with traditional backup technologies, continuous data protection mechanisms improve the recovery point objective (RPO) metric [22]. The RPO metric defines the time between two successful backups and, thus, the maximum amount of data loss during a successful recovery. The RPO is zero for a system with fully synchronized protection. The RPO metric of zero provided by continuous data protection theoretically allows unlimited recovery points [28]. The `sauegardeEX` tool implements continuous data protection at the block level (virtual machine snapshots) and at the file level. The virtual machine snapshots were also taken every minute to recover the running system, including volatile RAM memory. With more resources, snapshots could be taken more frequently.

The digital twin framework, which was equipped with the client version of `sauegardeEx`, sent every file alteration along with the file content to the server. This enabled a specific file to be restored at any point in time and also addressed the problem of overwriting a freed storage area (due to file deletion or update). Instituting this mechanism during OpenPLC execution and replicating the environment via Polymorph ensured that all possible traces were recorded.

- **Forensic Analysis:** Continuous data protection was also exploited to generate data for digital forensic analysis. Continuous data protection technology has not been considered in the digital forensic context, but it is certainly important. In fact, the state storage and replay functionality supported by the modified continuous data protection mechanism enables different forensic tools to be used without the risk of compromising data in the real system or even rendering the data unusable. Indeed, the replication-based approach with state storage and replay completely supports the digital forensic process specified by Kent et al. [18].

## 5.2 Results and Evaluation

To evaluate the framework, network traffic between OpenPLC (master) and the control unit of the windmill (slave) was captured. The standard Modbus TCP protocol was used for communications. A Python-

based Modbus simulation tool `pyModbusTCP` was used to generate realistic sensor data. A `sensordata.py` script simulated the sensor data for the wind speed around the windmill. Eight registers in the Modbus slave device were relevant. The first four registers contained the current sensor data and the other four indicated the corresponding system status. The sensor data values ranged from one to ten. The sensor values were grouped into three system state categories, green, yellow and red:

- **Values 1-5:** System state is green (Statuscode 200).
- **Values 6-8:** System state is yellow (Statuscode 300).
- **Values 9-10:** System state is red (Statuscode 400).

Pseudorandom sensor values were written to the registers every ten seconds by the Python script. Pseudorandom numbers were used so that changes to the wind speed corresponded to ascending or descending patterns and the values did not vary too much. OpenPLC updated the system status every five seconds based on the sensor values. A slight delay always occurred before the new system status was stored in the registers.

Traffic between OpenPLC and the Modbus slave was transformed to the network templates by Polymorph. This mimicked the external behavior based on traffic content.

By applying Polymorph and `sauvegardeEx` to the replicated system running OpenPLC with persistent storage, all the states of the OpenPLC addresses and related file changes during execution were recorded. All the system artifacts were simultaneously recorded at the digital twin. This enabled the determination of the relationships between the changed states and file content at any point during execution. Since the framework was designed to acquire the actual file content of written files on the hard drive as well as volatile memory content, VirtualBox and its live snapshotting functionality were leveraged.

During the evaluation, 1,432 state changes were observed on the persistent data storage during a five-minute period. The periodic virtual machine snapshots enabled a retrospective analysis of the running system.

Several forensic tools were applied at three points in time during execution. Data used for the evaluation is available at [11]. Table 1 shows an excerpt of the analysis and suitable digital forensic tools. The tools presented in [32] were considered in the evaluation.

An important component of the evaluation was to analyze discrepancies between the real system and its digital twin. This was accomplished by comparing file-level evidence between the real system and its digital



Table 1. Excerpt of the recorded state changes and suitable digital forensic tools.

Timestamp	State Change	File Name	Suitable Tools
2020-09-11 09:16:59.123	File modification	<code>persistent.file</code>	CPLCD
2020-09-11 09:17:01.102	File modification	<code>openplc.db</code>	Bring2lite
2020-09-11 09:17:23.322	File modification	<code>persistent.file</code>	CPLCD
2020-09-11 09:18:00.000	VM snapshot	<code>%/disk.vmdk</code>	Autopsy, CPLCD
2020-09-11 09:18:01.202	File modification	<code>persistent.file</code>	CPLCD
2020-09-11 09:18:01.302	File modification	<code>openplc.db</code>	Bring2lite

twin using the approximate hashing function of Breitingner and Baier [2]. This hashing function was used instead of the SHA-256 hashing function because it provides measures of file similarity. Frequent comparisons of the recorded files helped identify and verify time-event correlations. Comparisons of the 1,432 recorded state changes yielded an average similarity of 98%.

## 6. Discussion

The proposed approach is easily implemented on architectures with open-source programmable logic controller software. All that is needed is knowledge about the real system and adequate recordings of network traffic.

Although a digital twin adequately replicates a real system, this research has omitted formal measurements of the similarity between them. In order for evidence from a digital twin to be admissible, it is vital their similarity be measured and documented. One approach is to use the synchronization function proposed by Gehrman and Gunnarsson [15]. However, this mechanism can introduce time differences between the digital twin and its real counterpart. Specifically, system states caused by an attacker in the real system would manifest themselves earlier than in the digital twin.

An interesting possibility is to incorporate control theory in a digital twin. This would make the digital twin a better replication of the real system that would, in turn, contribute to the admissibility of the extracted evidence.

The proposed approach provides recordings of file content at various points in time (via `sauvegardeEx`) and system-wide snapshotting of the running programmable logic controller software (via `VirtualBox`). These features make it possible to detect and analyze RAM-based malware. However, a limitation is that the implementation employed Unix and open-source OpenPLC software instead of industrial control system

firmware. Although the underlying theory is sound, the open-source implementation would hinder its application in industrial environments.

The implementation of a digital twin for forensic investigations can be expensive. In addition to creating a digital twin and verifying its fidelity, it would be necessary to constantly modify the digital twin and verify that it keeps up with any and all changes made to the real system. This would require digital forensic professionals to have considerable industrial control system expertise, which would be an expensive proposition.

## 7. Conclusions

As attacks on critical infrastructure assets increase, it is imperative to develop digital forensic techniques targeted for industrial control systems. However, taking an industrial control system offline to conduct a digital forensic investigation is infeasible and expensive. An attractive alternative is to conduct a forensic investigation of a digital twin of an industrial control system. Implementing a digital twin with replication-based state storage and replay enables the acquisition and analysis of file-level evidence. Additionally, the digital twin could be used to select the appropriate forensic tools for evidence acquisition and analysis before interacting with the real system, thereby reducing system downtime when conducting the investigation.

## Acknowledgement

This research was partly conducted for the ZIM SSSeC Project under Contract no. 16KN085725 from the German Federal Ministry of Economic Affairs and Energy.

## References

- [1] S. Boschert, C. Heinrich and R. Rosen, Next generation digital twin, *Proceedings of the Twelfth International Symposium on Tools and Methods of Competitive Engineering*, pp. 209–217, 2018.
- [2] F. Breitinger and H. Baier, Similarity preserving hashing: Eligible properties and a new algorithm MRSH-v2, *Proceedings of the Fourth International Conference on Digital Forensics and Cyber Crime*, pp. 167–182, 2012.
- [3] C. Chan, K. Chow, S. Yiu and K. Yau, Enhancing the security and forensic capabilities of programmable logic controllers, in *Advances in Digital Forensics XIV*, G. Peterson and S. Sheno (Eds.), Springer, Cham, Switzerland, pp. 351–367, 2018.

- [4] M. Dietz and G. Pernul, Digital twins: Empowering enterprises towards a system-of-systems approach, *Business and Information Systems Engineering*, vol. 62(2), pp. 179–184, 2020.
- [5] M. Dietz and G. Pernul, Unleashing the digital twin’s potential for ICS security, *IEEE Security and Privacy*, vol. 18(4), pp. 20–27, 2020.
- [6] `dupgit`, `cdpfgl`: Continuous Data Protection for GNU/Linux, GitHub ([github.com/dupgit/sauvegarde](https://github.com/dupgit/sauvegarde)), 2021.
- [7] M. Eckhart and A. Ekelhart, A specification-based state replication approach for digital twins, *Proceedings of the Workshop on Cyber-Physical Systems Security and Privacy*, pp. 36–47, 2018.
- [8] M. Eckhart and A. Ekelhart, Towards security-aware virtual environments for digital twins, *Proceedings of the Fourth ACM Workshop on Cyber-Physical System Security*, pp. 61–72, 2018.
- [9] M. Eckhart and A. Ekelhart, Digital twins for cyber-physical systems security: State of the art and outlook, in *Security and Quality in Cyber-Physical Systems Engineering*, S. Biffi, M. Eckhart, A. Lüder and E. Weippl (Eds.), Springer, Cham, Switzerland, pp. 383–412, 2019.
- [10] M. Eckhart, A. Ekelhart and E. Weippl, Enhancing cyber situational awareness for cyber-physical systems through digital twins, *Proceedings of the Twenty-Fourth IEEE International Conference on Emerging Technologies and Factory Automation*, pp. 1222–1225, 2019.
- [11] L. Englbrecht, `DTDFEvaluation`, GitHub ([github.com/LudwigEnglbrecht/DTDFEvaluation](https://github.com/LudwigEnglbrecht/DTDFEvaluation)), 2021.
- [12] L. Englbrecht, `sauvegardeEX`, GitHub ([github.com/LudwigEnglbrecht/sauvegardeEX](https://github.com/LudwigEnglbrecht/sauvegardeEX)), 2021.
- [13] L. Englbrecht, S. Meier and G. Pernul, Towards a capability maturity model for digital forensic readiness, *Wireless Networks*, vol. 26(7), pp. 4895–4907, 2020.
- [14] M. Farsi, A. Daneshkhah, A. Hosseinian-Far and H. Jahankhani (Eds.), *Digital Twin Technologies and Smart Cities*, Springer, Cham, Switzerland, 2020.
- [15] C. Gehrman and M. Gunnarsson, A digital twin based industrial automation and control system security architecture, *IEEE Transactions on Industrial Informatics*, vol. 16(1), pp. 669–680, 2020.

- [16] M. Grieves and J. Vickers, Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems, in *Transdisciplinary Perspectives on Complex Systems*, F. Kahlen, S. Flumerfelt and A. Alves (Eds.), Springer, Cham, Switzerland, pp. 85–113, 2017.
- [17] International Electrotechnical Commission, IEC 61131-3:2013 Programmable Controllers – Part 3: Programming Languages, Geneva, Switzerland, 2013.
- [18] K. Kent, S. Chevalier, T. Grance and H. Dang, Guide to Integrating Forensic Techniques into Incident Response, NIST Special Publication 800-86, National Institute of Standards and Technology, Gaithersburg, Maryland, 2006.
- [19] P. Kieseberg and E. Weippl, Security challenges in cyber-physical production systems, in *Software Quality: Methods and Tools for Better Software and Systems*, D. Winkler, S. Biffi and J. Bergsmann (Eds.), Springer, Cham, Switzerland, pp. 3–16, 2018.
- [20] W. Kritzinger, M. Karner, G. Traar, J. Henjes and W. Sihn, Digital twins in manufacturing: A categorical literature review and classification, *IFAC-PapersOnLine*, vol. 51(11), pp. 1016–1022, 2018.
- [21] Y. Liu, L. Zhang, Y. Yang, L. Zhou, L. Ren, F. Wang, R. Liu, Z. Pang and M. Deen, A novel cloud-based framework for elderly healthcare services using digital twins, *IEEE Access*, vol. 7, pp. 49088–49101, 2019.
- [22] M. Lu and T. Chiueh, File versioning for block-level continuous data protection, *Proceedings of the Twenty-Ninth IEEE International Conference on Distributed Computing Systems*, pp. 327–334, 2009.
- [23] E. Negri, L. Fumagalli and M. Macchi, A review of the roles of digital twins in CPS-based production systems, in *Value Based and Intelligent Asset Management: Mastering the Asset Management Transformation in Industrial Plants and Infrastructures*, A. Crespo Marquez, M. Macchi and A. Parlikad (Eds.), Springer, Cham, Switzerland, pp. 291–307, 2020.
- [24] G. Palmer, A Road Map for Digital Forensic Research, DFRWS Technical Report, DTR-T001-01 Final, Air Force Research Laboratory, Rome, New York, 2001.
- [25] J. Rubio, R. Roman and J. Lopez, Analysis of cybersecurity threats in Industry 4.0: The case of intrusion detection, *Proceedings of the International Conference on Critical Information Infrastructures Security*, pp. 119–130, 2017.

- [26] G. Schroeder, C. Steinmetz, C. Pereira and D. Espindola, Digital twin data modeling with automationML and a communication methodology for data exchange, *IFAC-PapersOnLine*, vol. 49(30), pp. 12–17, 2016.
- [27] F. Servida and E. Casey, IoT forensic challenges and opportunities for digital traces, *Digital Investigation*, vol. 28(S), pp. S22–S29, 2019.
- [28] Y. Sheng, D. Wang, J. He and D. Ju, TH-CDP: An efficient block level continuous data protection system, *Proceedings of the International Conference on Networking, Architecture and Storage*, pp. 395–404, 2009.
- [29] shramos, Polymorph (v2.0.5), GitHub ([github.com/shramos/polymorph](https://github.com/shramos/polymorph)), 2020.
- [30] J. Tan, Forensic readiness: Strategic thinking on incident response, presented at the *Second Annual CanSecWest Conference*, 2001.
- [31] F. Tao, J. Cheng, Q. Qi, M. Zhang, H. Zhang and F. Sui, Digital twin driven product design, manufacturing and service with big data, *International Journal of Advanced Manufacturing Technology*, vol. 94(9), pp. 3563–3576, 2018.
- [32] T. Wu, F. Breitingner and S. O’Shaughnessy, Digital forensic tools: Recent advances and enhancing the status quo, *Digital Investigation*, vol. 34, article no. 300999, 2020.
- [33] K. Yau, K. Chow and S. Yiu, A forensic logging system for Siemens programmable logic controllers, in *Advances in Digital Forensics XIV*, G. Peterson and S. Shenoi (Eds.), Springer, Cham, Switzerland, pp. 331–349, 2018.
- [34] X. Yu, Y. Tan, Z. Sun, J. Liu, C. Liang and Q. Zhang, A fault-tolerant and energy-efficient continuous data protection system, *Journal of Ambient Intelligence and Humanized Computing*, vol. 10(8), pp. 2945–2954, 2019.

## 6 A Distributed Ledger Approach to Digital Twin Secure Data Sharing

---

Current status:	Published
Conference:	Data and Applications Security and Privacy (DBSec) - 33rd IFIP WG 11.3 Annual Conference, Charleston, SC, USA, July 15-17, 2019
Date of acceptance:	April 14, 2019
Full citation:	DIETZ, M., PUTZ, B. AND PERNUL, G. A Distributed Ledger Approach to Digital Twin Secure Data Sharing. In <i>Data and Applications Security and Privacy XXXIII. Lecture Notes in Computer Science</i> , vol. 11559. Springer, Cham (2019), pp. 281-300.
Authors contributions:	Marietheres Dietz      45% Benedikt Putz      45% Günther Pernul      10%

---

**Conference Description:** The annual IFIP WG 11.3 Conference on Data and Applications Security and Privacy (DBSec) tackles novel research on theoretical and practical aspects concerning data protection, privacy, and applications security. The conference proceedings are published by Springer under the Lecture Notes in Computer Science (LNCS) series.



# A Distributed Ledger Approach to Digital Twin Secure Data Sharing

Marietheres Dietz<sup>(✉)</sup>, Benedikt Putz, and Günther Pernul

University of Regensburg, Regensburg, Germany  
{marietheres.dietz,benedikt.putz,guenther.pernul}@ur.de

**Abstract.** The Digital Twin refers to a digital representation of any real-world counterpart allowing its management (from simple monitoring to autonomy). At the core of the concept lies the inclusion of the entire asset lifecycle. To enable all lifecycle parties to partake, the Digital Twin should provide a sharable data base. Thereby, integrity and confidentiality issues are pressing, turning security into a major requirement. However, given that the Digital Twin paradigm is still at an early stage, most works do not consider security yet. Distributed ledgers provide a novel technology for multi-party data sharing that emphasizes security features such as integrity. For this reason, we examine the applicability of distributed ledgers to secure Digital Twin data sharing. We contribute to current literature by identifying requirements for Digital Twin data sharing in order to overcome current infrastructural challenges. We furthermore propose a framework for secure Digital Twin data sharing based on Distributed Ledger Technology. A conclusive use case demonstrates requirements fulfillment and is followed by a critical discussion proposing avenues for future work.

**Keywords:** Trust frameworks · Distributed systems security · Distributed ledger technology · Digital twin

## 1 Introduction

Hardly anything has revolutionized society as much as digitization. At its beginning, data from everyday life was captured and stored digitally. After reaching significant amounts of digital data, recent years have been devoted to gaining relevant insights into data by leveraging Big Data Analytics, Artificial Intelligence and so on. A next step in digitization is now emerging in the form of the Digital Twin (DT) paradigm.

The Digital Twin refers to a digital representation of any real-world counterpart, at most times an enterprise asset. Its core building blocks are asset-specific data items, often enhanced with semantic technologies and analysis/simulation

---

The first two authors have contributed equally to this manuscript.

© IFIP International Federation for Information Processing 2019  
Published by Springer Nature Switzerland AG 2019  
S. N. Foley (Ed.): DBSec 2019, LNCS 11559, pp. 281–300, 2019.  
[https://doi.org/10.1007/978-3-030-22479-0\\_15](https://doi.org/10.1007/978-3-030-22479-0_15)

environments to explore the real-world asset digitally. The DT thus allows management of such an asset ranging from simple monitoring to autonomy. An essential part of the concept is the inclusion of the whole asset lifecycle. To integrate all lifecycle participants, the DT should provide comprehensive networking for its data, allowing it to be shared and exchanged [4].

Although the DT concept certainly advances digitization, it nevertheless poses new challenges in terms of IT security, especially in industrial ecosystems [10, 18]. Most notably, security must be maintained during the exchange of DT data between different, non-trusting parties. For instance, consider the DT of a power plant. Synchronizing tasks between twins should uphold integrity to avoid manipulated operations on the power plant. Also, involved parties should not be able to read every shared data element (e.g. the manufacturer of the power plant need not know the plant's current status), resulting in confidentiality requirements. To the best of our knowledge, current DT frameworks do not permit secure data sharing. Bridging this gap, our work provides a framework introducing security-by-design in DT data sharing.

To achieve this goal, we consider Distributed Ledger Technology (DLT). DLT is the umbrella term for distributed transaction-based systems, shared among several independent parties in a network. Distributed Ledgers have built-in mechanisms for access control and asset management, including authentication and authorization mechanisms. We focus on permissioned distributed ledgers, which target enterprise usage by restricting access to fixed set of independent and semi-trusted participants. One of the main reasons for using a Distributed Ledger is disintermediation, replacing the need for trust in a third party or central operator through a replicated and integrity-preserving database. Inherent transparency and auditability are additional advantages over centralized solutions. Due to these properties, DLT is uniquely suited to solve the challenges of DT secure data sharing.

Accordingly, this work proposes a framework for secure DT data sharing across an asset's lifecycle and collaborating parties based on DLT. We contribute to the body of knowledge by offering a solution without a trusted third party (TTP) based on security-by-design. The remainder of this paper is organized as follows: Sect. 2 introduces the background of our work. Afterwards, we proceed to the description of the current problems in DT data sharing and name the resulting requirements for secure DT data sharing (Sect. 3). In Sect. 4, we provide a framework for secure DT data sharing for multiple parties based on DLT. To show practical relevance and the functionality of our framework, a use case is provided in Sect. 5. In Sect. 6, we evaluate our approach in terms of fulfillment of the stated requirements. To conclude, Sect. 7 sums up the main contributions and gives an outlook for future work.

## 2 Background

At present, the *Digital Twin* phenomenon is still in its infancy. Nevertheless, implementation and design of this concept are addressed to date, especially in



the area of Industry 4.0. With strong focus on the industrial domain, the major part of research suggests DT implementation through AutomationML-formatted descriptive data of the real-world counterpart, e.g. [2,6,20]. The XML-based AutomationML (AML) format describes industrial assets and provides object-orientation for modeling the asset's physical and logical components [20]. Eckhart and Ekelhart [6] propose a framework for using a DT's simulation mode for security purposes such as pen testing. While these works focus on an initial development of a DT, the consideration of data sharing functions are still missing. However, exchanging data is vital for enabling the lifecycle integration and collaboration [4]. Our work builds on existing DT propositions, resulting in a concept that can be applied in a complementary way to enable secure DT data sharing.

Regarding DT data sharing, both the communication between lifecycle parties and the bidirectional communication between the DT and its real-world asset counterpart need to be considered. Bidirectional communication consists of the DT's instructions for the asset and the asset's status update for the DT. To uphold integrity among multi-domain DT models, Talkhestani et al. [21] offer a solution. They detect model changes by applying anchor points, and upon detection synchronize the DT while keeping model dependencies consistent. However, this includes drawbacks such as the manual creation of anchor points and reliance on a Product Lifecycle Management (PLM) system, while our solution offers platform-independence. Security aspects, such as the guarantee for all lifecycle partners to access the data while upholding confidentiality, are not considered to date, but integrated in our solution.

DT management is a form of enterprise asset management, which is one of the prime use cases of Distributed Ledgers [1]. Distributed Ledgers are able to track events and provenance information along an asset's lifecycle and increase transparency for all participants. For example, Litke et al. [12] studied the benefits of Distributed Ledgers for different actors in supply chain asset management, a research area closely related to DT asset management. In another study, Meroni and Plebani [14] investigate how the blockchain technology can be used for process coordination among smart objects. Smart objects are similar to DTs in that they are applied for monitoring physical artifacts. An issue with their proposed approach is that sensor data is also stored on the blockchain, which can be detrimental to performance and scalability. We consider this issue and provide a solution to overcome this obstacle.

### 3 Problem Statement

On the one hand, DTs should facilitate the access to asset information for different stakeholders along its lifecycle [17]. It is a task which enables feedback loops, while stepping towards a circular economy [3]. On the other hand, the involved parties do not necessarily trust each other, resulting in a confidentiality dilemma. A useful example is given in [13]: Two separate standalone DTs exist for a single device instance, one for the manufacturer and the other at the


284 M. Dietz et al.

customer site – due to information security reasons. Additionally, current works state that enterprise infrastructures need to overcome the following obstacles to provide secure DT data sharing:

- application of different tools [13,24]
- usage of various data formats [13]
- missing standards [4]
- broken information flow across lifecycle phases [13,24]
- clarification of the ownership of information [13].

This calls for a holistic approach that provides confidentiality and integrity, two central security dimensions in networks [26].

### 3.1 Digital Twin Model



**Fig. 1.** Overview of the asset lifecycle participants interacting with the DT.

Figure 1 illustrates DT data sharing and an exemplary set of lifecycle stakeholders. The depicted DT model comprises different *capabilities* and two types of asset-specific data. *Descriptive data* refers to static properties of the device and infrequently changing state information. This data is mainly produced by users. *Sensor data* occurs frequently and should be available in near real-time. It is generated by sensors of the physical asset or in its proximity, which provide valuable information on the asset’s environmental conditions. Moreover, data of both types needs to be synchronized with the physical counterpart. Therefore, the *sync* capability compares the state of the DT to its real-world counterpart and resolves possible discrepancies.

The *access control* capability provides authentication and authorization modules to enable data sharing of involved parties without hampering confidentiality. The *monitoring*, *simulation* and *analysis* capabilities represent advanced operations of the DT. Depending on the extent of the operations present in a DT, DT status data can be returned to the participant or the real-world counterpart’s state can be modified.

The depicted information flows show how information about the physical device is gathered from and sent to the lifecycle parties. Generally, the type of data accessed and shared by the different lifecycle parties depends on the real-world twin, the parties' roles in its lifecycle and thus, the specified access control mechanisms in the DT. The system flows represent necessary bidirectional synchronization between the DT and its real-world counterpart as stated in Sect. 2. Both flows contribute to making the data sharing activities of the involved parties traceable. This enables feedback from the latest stages of the asset lifecycle to the earliest ones [17].

### 3.2 A Formal Basis for Secure Digital Twin Data Sharing

Although a methodological literature analysis to establish requirements is the state-of-the-art approach, it is currently not sensible to carry out with regard to our research focus. On the one hand, this is due to the fact that only a small number of publications exist. In addition, data sharing has not yet been a focus in DT literature to date. Moreover, security-by-design concepts have not been considered yet. Therefore, we establish a formally valid basis in order to create a uniform understanding of DT data sharing. To derive the requirements, the

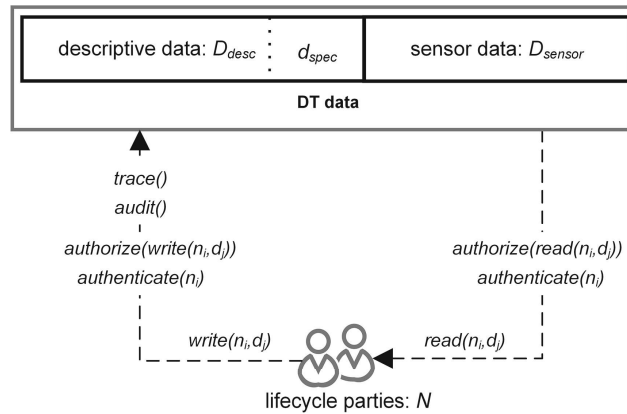


Fig. 2. Control flows for a single DT.

mechanisms to achieve the central goal of **secure DT data sharing** have to be examined in detail. Figure 2 illustrates the formal functions required to achieve this goal, which are also described hereafter.

**DT Data:** We see DT data twofold: At first, there is a set of descriptive data elements  $D_{desc} := \{d_1, \dots, d_m\}$  varying from documents to models or analytic outcomes. Its essential data element is the specification of the DT  $d_{spec} \in D_{desc}$ . The second set contains environmental, device-produced data, namely sensor data  $D_{sensor} := \{d_1, \dots, d_n\}$ , whereby  $D_{desc} \setminus D_{sensor}$ .

286 M. Dietz et al.

**Sharing:** A finite set of lifecycle parties  $N := \{n_1, \dots, n_k | k \geq 2\}$  can share the respective data elements of  $D_{desc}$  (write operation) or access the data elements of  $D_{desc}, D_{sensor}$  (read operation). This results in the following necessary functions:

$write(n_i, d_j | d_j \in D_{desc})$  and  $read(n_i, d_j | d_j \in D_{desc} \vee d_j \in D_{sensor})$ .

Note that  $1 \leq i \leq k$  as well as  $j \begin{cases} 1 \leq j \leq m & \text{if } j \in D_{desc} \\ 1 \leq j \leq n & \text{if } j \in D_{sensor} \end{cases}$ .

**Security-by-design:** Security-by-design infers introducing security mechanisms at the very beginning of a system's design [22]. In terms of DT data and sharing security, data integrity and confidentiality mechanisms are of special interest. Confidentiality in terms of securing data from view of non-trusted third parties can be reached by access control mechanisms [19]:

authentication:

$$authenticate(n_i)$$

authorization:

$$\begin{aligned} &authorize(read(n_i, d_j)) \\ &authorize(write(n_i, d_j)) \end{aligned}$$

Integrity of data can be achieved by auditability and traceability of write operations. Given  $D_{desc}$  as the origin set of data,  $D'_{desc}$  is the set of data after a data element  $d_j$  is added to the origin set. The following functions can cover integrity aspects:

auditability:

$$\begin{aligned} audit() : D_{desc} \rightarrow D'_{desc} &\iff write(n_i, d_j) \wedge \\ D_{desc} \not\rightarrow D'_{desc} &\iff \neg write(n_i, d_j) \end{aligned}$$

traceability:

$$trace() : D_{desc} \rightarrow D'_{desc} \implies D_{desc} \circ D'_{desc}$$

Thereby, auditability guarantees that  $D_{desc}$  is transformed to  $D'_{desc}$  in case of an authorized write operation whereas other operations are not able to transform the data in any way. Traceability ensures that authorized writes of data elements and thus, transformations of  $D_{desc}$  to  $D'_{desc}$ , are chained up. In conclusion, data integrity is ensured as the data cannot be manipulated or tampered with in retrospect.

### 3.3 Requirements for Secure DT Data Sharing

To provide a sound solution for secure DT data sharing, the following requirements were derived from the formal basis and the aforementioned challenges identified in the literature analysis.

**R1. Multi-party Sharing.** To enable lifecycle inclusion, a vital characteristic of the DT paradigm [4], the multiple stakeholders  $N$  involved in the lifecycle

have to be considered. As described in Fig. 1, parties can vary from manufacturer to maintainer. However, all involved parties are pre-registered and therefore determinable.

**R2. Data Variety Support.** At the heart of the DT lie the relevant digital artifacts  $D_{desc}$ ,  $D_{sensor}$ , which vary from design and engineering data to operational data to behavioral descriptions [4]. Thus, different data types and data formats [13] need to be supported during data sharing. For instance, Schroeder et al. claim that using the semi-structured AutomationML format to model attributes related to the DT ( $d_{spec}$ ) is very useful for DT data exchange [20]. In addition to semi-structured data, structured data (e.g. sensor tuples in  $D_{sensor}$ , database entries) and unstructured data such as human-readable documents can be asset-relevant and shared via the DT.

**R3. Data Velocity Support.** Often, DT data is distinguished between descriptive, rather static data, and behavioral, more dynamic data (see Fig. 1). The latter changes with time along the lifecycle of the real-world counterpart [20]: With each lifecycle stage the asset-related information evolves, resulting in different versions and a dynamic information structure [17]. Naturally, dynamic data includes sensor data  $D_{sensor}$  – which mostly refers to the actual state of the real-world counterpart [8]. While the infrequently changing data  $D_{desc}$  might not require high throughput, sensor and dynamic data  $D_{sensor}$  accrues in intervals ranging from minutes to milliseconds. Therefore, the solution must support high throughput and low sharing latency for efficient sharing of dynamic data – thus supporting data velocity.

**R4. Data Integrity and Confidentiality Mechanisms.** An important requirement is taking into account data security features, especially integrity and confidentiality. At first, this requirement aims at safeguarding data integrity to avoid wrong analytic decisions based on manipulated data. It can be ensured by *audit()* and *trace()* mechanisms. The second main security objective is to avoid confidentiality and trust problems while enabling multi-party participation. This calls for restricted data access dependent on the party through *authenticate()* and *authorize()* functions, while ideally keeping the effort for user registration low. Different levels of confidentiality should be possible for different data elements. For instance,  $D_{sensor}$  might need a lower level of protection than  $D_{desc}$ , as the latter might include sensitive corporate information such as blueprints. Detailed *authorize()* functions, providing access-restrictions for each data element, can cover this aspect.

**R5. Read and Write Operations.** To interact with DT data, a DT data sharing solution must provide *read()* and *write()* data operations for the sharing parties. The allowance of operation modes for the data elements should be chosen carefully for each party to ensure R4 (cf. Fig. 2).

Overall, we do not claim that these requirements are complete. There may be other requirements of importance, but regard these as essential for the following reasons. On the one hand, these requirements were found to be mentioned most often in the reviewed literature, while others were less frequently mentioned and

are therefore considered of lower importance (see Sect. 6.2 for further explanation). On the other hand, the stated requirements were also the main focus in various practitioners reports (e.g. [9, 16, 25]) and during discussions with experts.

## 4 Solution Architecture

In order to develop a framework for secure DT data sharing, we first evaluate the suitability of DLT in Sect. 4.1. Afterwards, Sect. 4.2 explains the system architecture and Sect. 4.3 explains how the various data types are stored. Section 4.4 details the inclusion of the DT capabilities as part of the DLT solution. Finally, Sect. 4.5 explains the initial setup procedure for our framework.

### 4.1 Technology Selection

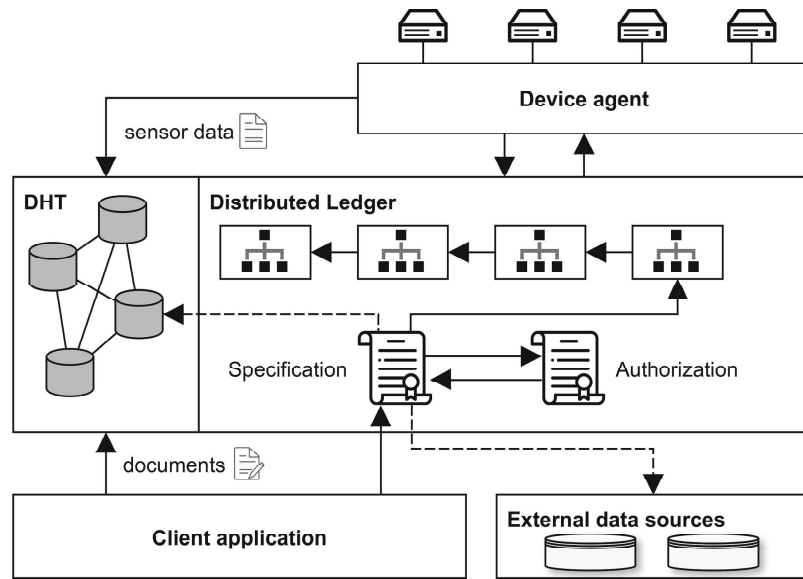
To develop a solution architecture, we first evaluate different data storage solutions' properties to select the technology best suited to fulfill the requirements.

A centralized solution could be created in the form of a portal, operated by a third party or the operator of the twinned device. This requires trust of the participating parties towards the portal maintainer, as the maintainer could manipulate data or revoke access to the DT for other parties. A distributed approach jointly operated by all participants could solve this trust issue. Distributed Ledgers represent such a distributed solution. They permit verifiable decentralized execution of business logic via *smart contracts*, ensuring that rules and processes agreed upon by the lifecycle participants are followed.

We evaluate the applicability of Distributed Ledgers to our DT data sharing requirements based on the blockchain applicability evaluation framework by Wüst and Gervais [28]. As illustrated in Fig. 1, there is a need to store various types of data as part of the DT state. Multiple parties interact with the twin during its lifecycle who do not fully trust each other. These writers are usually known in advance or change infrequently (i.e. the maintenance service provider changes). These characteristics lead to the choice of a public or private permissioned blockchain in the framework [28]. In our case, this choice depends on whether public auditability is required or not. While use-case dependent, we focus on private permissioned blockchains for the rest of this paper. If needed, public read-only access to blockchain data can be enabled during implementation for most permissioned blockchain frameworks (i.e. through a REST API).

### 4.2 System Architecture

The proposed DLT-based architecture for secure DT data sharing is shown in Fig. 3. Every participant runs three components: a node of a **Distributed Hash Table (DHT)**, a node of the **Distributed Ledger** and a **client application**. The DHT and Distributed Ledger make up the shared data storage, while the client application is responsible for the user interface and backend logic for retrieving and processing the data stored on the ledger and DHT. For owners



**Fig. 3.** DLT-based architecture for DT data sharing.

of twinned physical devices, a **Device agent** manages the physical devices and coordinates their interactions with the system. As part of operational technology, the Device agent functions as a bridge between the cross-organizational asset management system and the physical devices controlled by a single organization.

Data storage systems based on distributed ledgers have two ways of storing data: on-chain and off-chain [29]. On-chain storage is restricted to transactions and the internal state storage of smart contracts. Due to full replication of on-chain data, items larger than a few kilobytes in size need to be stored in a different, off-chain location. Using a traditional database would however result in a single point of failure or reintroduce a trusted party.

For this reason, we resort to a structured DHT for large data items. DHTs are distributed key-value stores, where all key-value pairs are mapped to one or more nodes. The DHT entries can be linked to the corresponding on-chain asset based on the DHT key hash. By storing the hash on the blockchain, integrity of the off-chain data can be verified after retrieving it from the DHT. To maintain confidentiality and availability, data stored on the DHT is encrypted, sharded and replicated. Correspondingly, an access control mechanism is needed to allow authorized parties to access the data. The k-rAC scheme illustrates how a DHT can implement the required functionality [11]. In k-rAC, access control is implemented using access control lists (ACL) stored along with each key-value pair on the DHT. We propose reusing the Distributed Ledger’s public key identities for DHT authentication. A symmetric key is used for encryption, which is then available to authorized parties by encrypting it with their public key. The encrypted access keys are distributed with each data item’s ACL. Manipulation of the ACL is prevented by requiring a quorum of  $2k + 1$  nodes for write operations, where  $k$  is the number of tolerated malicious nodes.

### 4.3 Data Storage

There are two types of **descriptive data** that need to be stored by the system: a machine-readable specification and device-related unstructured data (i.e. human-readable documents). The **specification** includes a description of the device's hardware components as well as their functions. The DT's physical properties are derived from this specification. For our work we assume that AML is used to describe the physical asset. The AML specification is stored on the ledger in a modifiable way. This approach guarantees that updates to the device specification are observed by all parties. Distributed Ledgers can store complex modifiable state by using *smart contracts*. We thus refer to the resulting contract as the *specification contract*.

**Unstructured data** can be uploaded to the system and may subsequently be annotated or modified by other parties. Due to its size it cannot easily be parsed and stored in contracts. For this reason, it is stored off-chain and registered in the smart-contract with a document title and a hash of the contents. To update a document, a new version must be uploaded to the DHT and the smart contract reference updated. This ensures that changes to the documents are traceable.

**Sensor data** needs to be stored off-chain due to its frequent updates and the considerable amount of generated data. A history of the sensor data is kept to allow for further analysis, e.g. predictive maintenance or troubleshooting. The link to the on-chain data is established via a pointer to the off-chain storage location, stored on-chain in the specification contract. To avoid having to update the storage location hash every time new sensor data is uploaded to the DHT, we take advantage of *DHT feeds*. This concept is inspired by the Ethereum network's DHT Swarm [7]. In Swarm, a feed is defined by a feed manifest with a unique address on the network. The feed manifest's owner (i.e. the physical device) is the only user permitted to upload signed and timestamped data to this address. Any data format can be used and a history of uploaded data is kept. The DHT feed enables frequent sensor data sharing without having to update an on-chain reference. Based on the feed, the client application may compare sensor updates with expected values derived from the specification contract to detect anomalies. Additionally, there is no need for directly accessing the physical device, which may reside in a protected network. Instead, data updates are readily available on the DHT for authorized participants.

Many organizations also have additional internal data sources or microservices that provide structured data relevant to the Digital Twin. These data sources can be included in the twin by adding references (i.e. an URI) to the DT specification contract. This allows inclusion of legacy data sources and complex data which cannot easily be stored on a DHT (i.e. relational data). If the external data source requires authentication, it is the responsibility of the data source provider to ensure access rights for the DT ledger's identities.

Listing 1.1 shows a pseudocode representation of the data types stored in the specification contract. The syntax is inspired by Ethereum's Solidity smart contract programming language. All data stored on the contract is readable



by all lifecycle participants. Besides general device metadata, the contract also includes a program call queue for interaction with the physical device’s program interfaces (see also Sect. 4.4). Since smart contracts must be deterministic and thus cannot interact with files, the AML specification is stored in a string variable. This variable can later be parsed and modified, as illustrated in Sect. 4.4. Hash references to new original documents on the DHT are kept track of in the `documents` mapping. The hash serves as an identifier, while the `document` struct provides metadata. Updated versions of each document are stored in the `documentVersions` mapping. The `componentID` and corresponding feed reference of the sensor data stream on the DHT are stored in the `sensorFeeds` mapping.

```

/* metadata and specification*/
string deviceName
string deviceID
string deviceAML
string[] callProgramQueue

/* additional descriptive data */
struct Document {
    uint timestamp
    string description
    address owner
}

struct ExternalSource{
    string URI
    address owner
}

mapping(string=>Document) documents
mapping(string=>string[]) documentVersions
ExternalSource[] externalSources

/* sensor data */
mapping(string=>string) sensorFeeds

```

**Listing 1.1.** Data structures of the specification contract

```

/* descriptive data interfaces */
function addDocument(document)
function addDocumentVersion(string hash)
function removeDocument(string hash)

function addExternalSource(string URI)
function removeExternalSource(string URI)

/* sensor data interfaces */
function addSensorFeed(string componentID,
    string reference)

function removeSensorFeed(string componentID)

/* interaction with the specification */
function insertAML(string amlCode, string
    parentID, string afterID)
function removeAML(string ID)
function callProgram(string programName,
    string parameters[])

```

**Listing 1.2.** Function interfaces of the specification contract

#### 4.4 Capabilities

We focus on the three capabilities required for accessing and publishing DT data: *DT interaction*, *access control* and *sync*.

*DT interaction* refers to the information flows in Fig. 1, which allow users to interact with the twin’s data. The specification contract implements this functionality. It allows users to read and potentially modify the DT instance. The relevant interfaces that can be called with transactions are shown in Listing 1.2. New or updated references to documents may be appended by any authorized user. The same applies to external data sources and sensor feed references to the DHT. The specification can be manipulated by inserting or removing specific AML segments, which are identified by their ID. To determine the position of a new AML code segment in the AML document, the parent ID and the ID of the preceding element need to be passed as parameters. The twin’s program

interfaces for setting device parameters can be accessed via `callProgram`. This function checks authorization, finds the requested program in the AML specification and places it in a queue for the Device agent to retrieve. The agent then forwards the program call to the device for execution.

The *access control* capability is responsible for authentication and authorization of user interactions with the DT data. For user authentication, accounts are created on the blockchain and represented by their public key. An initial solution could be provided by the framework's built-in identity management, for example Hyperledger Fabric's Membership Service Provider (MSP) [1]. The MSP lists the certificate authorities who may issue digital identities for the Distributed Ledger. The same identity can then be reused for authentication in the DHT. Authorization is realized in a separate access control smart contract. Any protected interaction with the Digital Twin is first authorized through that contract. Such interactions are for example modifications of the twin's properties, like changing parameters or modifying its specification. A query from the client application provides an identity to the specification contract, which then interacts with the authorization contract to determine if the user is allowed to perform the action. Authorization is then granted or denied based upon a stored role-permission mapping. Accordingly, the contract's interfaces are based upon a Role-based Access Control (RBAC) scheme. We do not describe the access control contract in detail here, as there are other works describing blockchain-based access control schemes [5].

The *sync* capability requires regular interaction between the Device agent and the Distributed Ledger. For synchronization, the Device agent pulls updates from the real-world asset and uploads them to the off-chain DHT sensor data feed. The Device agent monitors the ledger and pushes any modifications instructed by committed on-chain transactions to the asset. The synchronization interval depends on the use case.

Other DT capabilities like *monitoring*, *simulation* and *analysis* can be executed off-chain by interacting with the local copy of the ledger. Simulation or analysis instructions and results can be shared on the ledger as documents. This would allow other parties to verify the results, should they desire to do so.

#### 4.5 Setup Process

Initially, each lifecycle participant sets up one network node running both a DHT and a Distributed Ledger node. These serve as local replicas of ledger data and access points for off-chain data. They may also be used for transaction-based interaction with the smart contracts. Additionally, an identity provider must be set up to allow federated identities from all participating organizations based on public key certificates.

Once the network is set up, a Digital Twin instance can be created on the ledger by the device owner. The manufacturer should first provide the AML file to the owner, who then proceeds to set up a Digital Twin sharing instance on the ledger. The client application provides the interface to upload the file and create a smart contract based on it. Before uploading, the owner also needs to

specify the access rights associated with the various parts of the specification. Although use case dependent, sensible default values could be *write* access by owner and maintainer and *read* access by everyone else.

In this way, any number of Digital Twin instances can be created by the various parties on the network. Each instance is represented by a specification contract. Subsequent modifications take place via authorized on-chain transactions and are stored as part of the contract's internal state. As a result, auditing the twin is possible by (actively or retroactively) monitoring smart contract transactions for anomalies.

## 5 Use Case

This chapter intends to show how the theoretical framework developed in Sect. 4 is traversed in a use case. To begin with, the overall setting of the use case is described in Sect. 5.1, while the subsequent Sect. 5.2 iterates the use case through the solution architecture. At last, a summary is given, focusing on the automation degree in data sharing and the reading operation (Sect. 5.3).

### 5.1 Setting

The setting is chosen close to reality. The asset, the real-world counterpart to the DT, is a bottling plant, where bottles are filled with beverages. The parties involved in the asset lifecycle are a manufacturer, an owner, a maintainer of the bottling plant and an external auditor that audits the safety of our bottling plant. For our use case, we consider the following scenario: The bottles are flooding due to a broken sensor in the bottling plant. Consequently, the maintainer detects the damage and changes the broken sensor in the bottling plant.

This entails the following shared data interactions. At first, the specification of the plant needs to be updated by replacing the broken sensor's specification entry with the newly added sensor. Additionally, the new sensor's data stream has to be integrated in place of the old sensor stream. Other documents concerning the maintenance task might also be shared, such as a maintenance report.

While the maintainer is the only party sharing data in this scenario, the owner should also be updated on the state of the bottling plant. Furthermore, the manufacturer needs to be informed that the sensor is broken, so that an analysis of the time and circumstances can be conducted. This way relevant insights for future plant manufacturing can be gained. Additionally, the external auditor needs to access the information about the maintenance task to review the procedure in terms of safety compliance.

### 5.2 Framework Iteration

This use case triggers a specific logical order of events in the framework, which are highlighted in Fig. 4 and described hereafter. The framework first comes into play when the maintainer replaces the broken sensor.

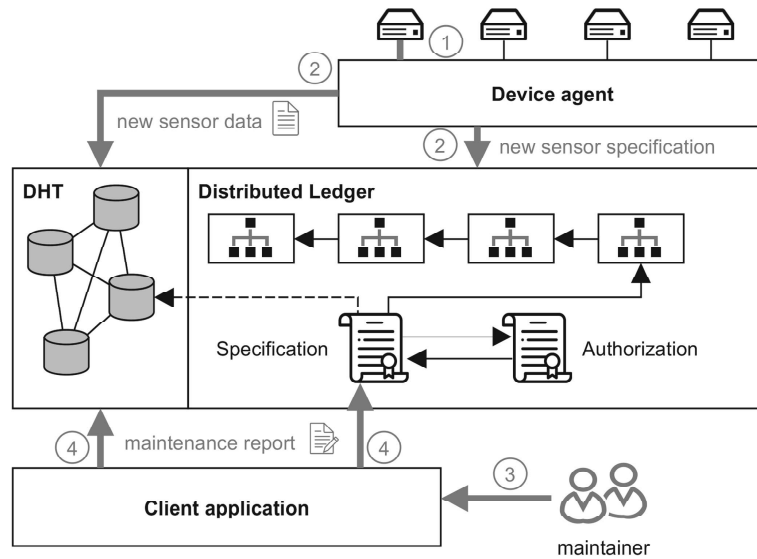


Fig. 4. Use case tailored architecture for DT data sharing.

1. All devices are connected with the **Device agent**, which registers the exchange of the broken sensor. Additionally, it gathers information about the new sensor.
2. Following the new sensor connection, the **Device agent** forwards the new incoming data stream of the sensor into the **DHT**. The location of the stored sensor stream in the **DHT** is registered by the **Device agent**.

The **Device agent** then sends a transaction containing the new sensor specification to the **Distributed Ledger**. This transaction invokes the specification contract, resulting in several updates. First, the old sensor entry is removed and the new sensor specification given by the **Device agent** is added. Secondly, the storage location of the sensor stream on the **DHT** is added by a reference to the location. These three transactions concerning the specification are stored on the **Distributed Ledger**.

3. Having performed the maintenance task, the maintainer writes a maintenance report and pushes it onto the **Client application**.
4. The **Client application** adds the maintenance report by performing two actions. Firstly, it adds the report to the off-chain **DHT**. Secondly, it stores the reference to the **DHT** location of the report on the specification contract. Thereby, the location is added to the entry of the sensor specification.

### 5.3 Results

In a nutshell, the recognition of new sensor and the AML update with the new component is already accomplished by the **Device agent** without requiring human interaction. The new data stream is automatically forwarded to the **DHT** and the reference to the new storage location of the component's data stream is added to the specification contract. Additional unstructured non-specification

data (e.g. the maintenance report) can be added manually. The **Client application** takes care of the necessary background work by inserting the file into the **DHT** and adding the respective storage reference into the specification contract.

All participating parties can view the latest transactions on the ledger – presented in a comprehensive way in the **Client application**. Advanced **Client applications** could also notify the user whenever an ledger update takes place.

Considering security, the advantages of this framework shine when compared to the alternative solution: A TTP could deliberately transfer shared information and know-how to rival enterprises. For instance, confidential sensor data or blueprints could be leaked to competitors, which may then deduce quality issues of the rival product. The service of the TTP could also be compromised by attackers, resulting e.g. in a violation of integrity so that the sharing parties receive inconsistent asset versions.

## 6 Evaluation

To evaluate our framework, Sect. 6.1 discusses the suitability of the framework in reference to the requirements. Finally, the results are discussed in Sect. 6.2.

### 6.1 Requirements Fulfillment

To sum up, our approach fulfills the requirements **R1–R5**. The following paragraphs explain how each requirement was addressed in our solution architecture.

**R1. Multi-party Sharing.** The main argument for using Distributed Ledgers is the involvement of multiple parties  $N$  who produce and consume data. Next to the ledger, our approach provides a client application for all parties that accesses the data on the ledger and the DHT. Therefore, our approach clearly fulfills **R1**.

**R2. Data Variety Support.** To enable the sharing of different data in various formats, our approach provides a central documentation and two storage options. The standardized asset description  $d_{spec}$  is included in the Distributed Ledger and serves as the basis of the DT within the specification contract. All other data of  $D_{desc}$  as well as the sensor data  $D_{sensor}$  are stored off-chain in the DHT. Moreover, each stored data element in DHT is registered in the central specification contract as a reference to the storage location of the data element. For instance, a sensor in the specification contract contains a reference to the storage position of its data stream in the DHT. Hence, **R2** is met.

**R3. Data Velocity Support.** Modern sensor data streams' frequency and volume exceed the performance characteristics of current Distributed Ledger frameworks. Since the data streams  $D_{sensor}$  do not describe main features of the DT ( $d_{spec}$ ), they are stored off-chain in the DHT. This way, high throughput of  $D_{sensor}$  is supported, while the sharing latency is also kept low (seconds). The Distributed Ledger maintains verifiability by storing the hash reference to the data stream on the DHT in the specification contract. This ensures no loss in performance and data access through the DHT, supporting **R3**.

**R4. Data Integrity and Confidentiality Mechanisms.** With respect to data integrity, the Distributed Ledger attaches every new data element (*trace()*) and prevents manipulation of the data by replicating it among all involved parties. A manipulation would result in a version mismatch or loss of consensus and could be detected easily (*audit()*). The second storage component (DHT) also supports integrity by storing the respective hash values to the data. A manipulation of DHT data would also be detected by a mismatch between the hashes in the nodes (*audit()*). However, there remains the problem of adding non-valid data, which is a common issue in the area of DLT. Here, we rely on the parties' interest in sharing valid data and on mechanisms ensuring quality of input data that the respective responsible party applies.

In terms of data confidentiality, our approach ensures that the data is read only by authenticated and authorized parties. Authentication is ensured through lifecycle party login to the client application (*authenticate()*). Access control concerning the party and the data elements is realized through an ACL and encryption for off-chain data and an authorization smart contract for on-chain data (*authorize()*). In concrete terms, the ACLs specify access rights on a per-document basis, while the smart contract stores authorization information for all involved parties. Therefore, different confidentiality levels can be realized.

To conclude, our approach provides data integrity and confidentiality mechanisms (**R4**) – reinforcing data security in DT data sharing.

**R5. Read and Write Operations.** Read and write operations are managed through the Client application. For *read()* operations, the Client application fetches the requested data from the DHT and the ledger and presents the data in a comprehensive way adjusted for the demanding party. In case of a *write()* operation, the Client application triggers the right procedure to alter the smart contract with a transaction and uploads additional asset-relevant data beyond specification to the DHT. Consequently, our approach also fulfills **R5**.

## 6.2 Discussion

Keeping the requirements *variety* (**R2**) and *velocity* (**R3**) in mind, the question arises why data *volume* is not considered a requirement. As literature is currently not at consensus regarding the relevance of the Big Data feature *volume* [15] for Digital Twins, we consider explicit support for data volume to be non-necessary. Nevertheless, by storing documents off-chain, our approach can handle considerable amounts of data. Future implementations of our concept may conduct benchmark studies to explore scalability limits with regard to big data volumes.

It should be noted that our approach depends on multi-party participation. The more independent parties maintain the Distributed Ledger and DHT, the less vulnerable the data sharing is to manipulation. With regard to the access control capability, a decentralized identity management solution with a shared identity database could be an even more holistic, next-generation solution.

While we are aware that our approach currently lacks an implementation, we nevertheless believe that the use case shows suitability for practice. Future

work will focus on implementing the framework. Here, challenges might include adjusting a DHT framework to support authorization and data feeds (although Swarm shows promise in this regard [7]), as well as selecting a suitable Distributed Ledger framework.

The Distributed Ledger and the concomitant smart contracts could also be handled in a different way. For instance, the AML could be transformed into classes and types in the smart contract, similar to the BPMN to Solidity transformation in [27]. However, the effort clearly outweighs the utility as AML is a very powerful standard allowing very complex descriptions. Moreover, not all of the hypothetically generated classes and functions might be needed. Plus, functions or classes might be newly added later on, which results in the need to re-create the smart contract as they are currently not represented in the smart contract. This clearly increases effort and downgrades utility.

Another issue is entailed by the possibility to directly alter variable values referring to an actual function in our current version of the ledger. For instance, consider a PLC device with various functions such as setting a conveyor belt's velocity (with an integer parameter). Without constraints, the changed velocity could exceed safety bounds. Safety threats like this one, be they malicious or accidental, need to be mitigated in a production system. Therefore, we suggest integrating safety and security rules as proposed in [6]. They could be integrated as part of the specification contract, with the Device agent checking conformance of program calls on synchronization.

With respect to the current problems hampering secure DT data sharing, our approach tackles the issues stated in Sect. 3 in the following ways:

- The usage of different tools that can be connected with our main data sharing approach (External data sources, Fig. 3) is possible (*application of different tools*)
- Our approach is tailored for the integration of data in multiple formats and variety as stated in Sect. 4.3 (*usage of various data formats*)
- An agreement only on the standard describing the asset (e.g. AML) is required to transform the main description of the asset into the specification smart contract, while other standardized or non-standardized data can still be shared via the DHT (*missing standards*)
- The proposed shared collaborative data basis is distributed among all involved parties and the information flow is universal across the lifecycle phases (*broken information flow across lifecycle phases*)
- The Distributed Ledger registers the data as well as the involved party sharing the data, while mechanisms such as access control (Authorization contract, Fig. 3) support confidentiality issues (*clarification of the ownership of information*).

To sum up, the major part of the identified issues in the literature referring to DT data sharing are diminished or solved by our approach.

## 7 Conclusion

DT data not only ties physical and virtual twin [23], it also enables integration of the whole asset lifecycle, which is essential for realizing the DT paradigm. Moreover, the exchange of asset-relevant data (DT data) is vital for achieving the effects of a feedback loop. Closing the feedback loop in turn favors the development of a circular economy.

However, maintaining data security becomes a major requirement when sharing DT data between multiple parties, especially as the parties do not necessarily trust each other. Our approach of applying DLT can clearly solve this issue and enable secure multi-party data sharing. It provides confidentiality through access control arranged by usage of a smart contract. Moreover, data integrity is implicitly supported through the immutability of the original data in the ledger.

To conclude, our approach fulfills the requirements **R1–R5** for secure DT data sharing. Nevertheless, there remain minor drawbacks that need to be addressed in future research (see Sect. 6.2). Our upcoming work will focus on implementing our theoretical concept to demonstrate its feasibility in practice.

## References

1. Androulaki, E., et al.: Hyperledger fabric: a distributed operating system for permissioned blockchains. In: Proceedings of the Thirteenth EuroSys Conference, EuroSys 2018, pp. 30:1–30:15. ACM, New York (2018). <https://doi.org/10.1145/3190508.3190538>
2. Banerjee, A., Dalal, R., Mittal, S., Joshi, K.P.: Generating digital twin models using knowledge graphs for industrial production lines. In: Workshop on Industrial Knowledge Graphs, No. June, pp. 1–5 (2017). <http://ebiquity.umbc.edu/paper/html/id/779/Generating-Digital-Twin-models-using-Knowledge-Graphs-for-Industrial-Production-Lines>
3. Baumgartner, R.J.: Nachhaltiges Produktmanagement durch die Kombination physischer und digitaler Produktlebenszyklen als Treiber für eine Kreislaufwirtschaft. In: Interdisziplinäre Perspektiven zur Zukunft der Wertschöpfung (2018). [https://doi.org/10.1007/978-3-658-20265-1\\_26](https://doi.org/10.1007/978-3-658-20265-1_26)
4. Boschert, S., Heinrich, C., Rosen, R.: Next generation digital twin. In: Proceedings of TMCE 2018, No. May (2018). <https://www.researchgate.net/publication/325119950>
5. Di Francesco Maesa, D., Mori, P., Ricci, L.: Blockchain based access control. In: IEEE Blockchain Conference 2018, pp. 1379–1386 (2018). [https://doi.org/10.1007/978-3-319-59665-5\\_15](https://doi.org/10.1007/978-3-319-59665-5_15)
6. Eckhart, M., Ekelhart, A.: Towards security-aware virtual environments for digital twins. In: Proceedings of the 4th ACM Workshop on Cyber-Physical System Security - CPSS 2018, pp. 61–72 (2018). <https://doi.org/10.1145/3198458.3198464>
7. Ethereum Swarm Contributors: Swarm 0.3 documentation (2019). <https://readthedocs.org/projects/swarm-guide/downloads/pdf/latest/>
8. Glaessgen, E., Stargel, D.: The digital twin paradigm for future NASA and U.S. air force vehicles. In: 53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference (2012). <https://doi.org/10.2514/6.2012-1818>



9. Greengard, S.: Building a Better Iot (2017). <https://cacm.acm.org/news/218924-building-a-better-iot/fulltext>
10. ICS-CERT: Overview of cyber vulnerabilities. Technical report (2017). <https://ics-cert.us-cert.gov/content/overview-cyber-vulnerabilities>
11. Kieselmann, O., Wacker, A., Schiele, G.: k-rAC - a fine-grained k-resilient access control scheme for distributed hash tables. In: Proceedings of the 12th International Conference on Availability, Reliability and Security, ARES 2017, Reggio Calabria, Italy, pp. 1–43. ACM, New York (2017). <https://doi.org/10.1145/3098954.3103154>
12. Litke, A., Anagnostopoulos, D., Varvarigou, T.: Blockchains for supply chain management: architectural elements and challenges towards a global scale deployment. *Logistics* **3**(1) (2019). <https://doi.org/10.3390/logistics3010005>
13. Malakuti, S., Grüner, S.: Architectural aspects of digital twins in IIoT systems. In: Proceedings of the 12th European Conference on Software Architecture Companion Proceedings - ECSA 2018, pp. 1–2 (2018). <https://doi.org/10.1145/3241403.3241417>
14. Meroni, G., Plebani, P.: Combining artifact-driven monitoring with blockchain: analysis and solutions. In: Matulevičius, R., Dijkman, R. (eds.) CAiSE 2018. LNBI, vol. 316, pp. 103–114. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-92898-2\\_8](https://doi.org/10.1007/978-3-319-92898-2_8)
15. Negri, E., Fumagalli, L., Macchi, M.: A review of the roles of digital twin in CPS-based production systems. *Procedia Manuf.* **11**(June), 939–948 (2017). <https://doi.org/10.1016/j.promfg.2017.07.198>
16. Ovtcharova, J., Grethler, M.: Beyond the Digital Twin - Making Analytics come alive. *visIT [Industrial IoT - Digital Twin]*, pp. 4–5 (2018). <https://www.iosb.fraunhofer.de/servlet/is/81714/>
17. Ríos, J., Hernández, J.C., Oliva, M., Mas, F.: Product avatar as digital counterpart of a physical individual product: literature review and implications in an aircraft. In: *Advances in Transdisciplinary Engineering* (2015). <https://doi.org/10.3233/978-1-61499-544-9-657>
18. Rubio, J.E., Roman, R., Lopez, J.: Analysis of cybersecurity threats in industry 4.0: the case of intrusion detection. In: D’Agostino, G., Scala, A. (eds.) CRITIS 2017. LNCS (LNAI and LNB), vol. 10707, pp. 119–130. Springer, Heidelberg (2018). [https://doi.org/10.1007/978-3-319-99843-5\\_11](https://doi.org/10.1007/978-3-319-99843-5_11)
19. Sandhu, R.S., Samarati, P.: Access control: principles and practice. *IEEE Commun. Mag.* (1994). <https://doi.org/10.1109/35.312842>
20. Schroeder, G.N., Steinmetz, C., Pereira, C.E., Espindola, D.B.: Digital twin data modeling with automationML and a communication methodology for data exchange. *IFAC-PapersOnLine* **49**(30), 12–17 (2016). <https://doi.org/10.1016/j.ifacol.2016.11.115>
21. Talkhestani, B.A., Jazdi, N., Schloegl, W., Weyrich, M.: Consistency check to synchronize the Digital Twin of manufacturing automation based on anchor points. *Procedia CIRP* (2018). <https://doi.org/10.1016/j.procir.2018.03.166>
22. Tankard, C.: The security issues of the Internet of Things. *Comput. Fraud Secur.* **2015**(9), 11–14 (2015). [https://doi.org/10.1016/S1361-3723\(15\)30084-1](https://doi.org/10.1016/S1361-3723(15)30084-1)
23. Tao, F., Cheng, J., Qi, Q., Zhang, M., Zhang, H., Sui, F.: Digital twin-driven product design, manufacturing and service with big data. *Int. J. Adv. Manuf. Technol.* **94**(9–12), 3563–3576 (2018). <https://doi.org/10.1007/s00170-017-0233-1>
24. Uhlemann, T.H., Lehmann, C., Steinhilper, R.: The digital twin: realizing the cyber-physical production system for industry 4.0. *Procedia CIRP* (2017). <https://doi.org/10.1016/j.procir.2016.11.152>

300 M. Dietz et al.

25. Usländer, T.: Engineering of digital twins. Technical report, Fraunhofer IOSB (2018). <https://www.iosb.fraunhofer.de/servlet/is/81767/>
26. Voydock, V.L., Kent, S.T.: Security mechanisms in high-level network protocols. *ACM Comput. Surv.* (1983). <https://doi.org/10.1145/356909.356913>
27. Weber, I., Xu, X., Riveret, R., Governatori, G., Ponomarev, A., Mendling, J.: Untrusted business process monitoring and execution using blockchain. In: La Rosa, M., Loos, P., Pastor, O. (eds.) *BPM 2016*. LNCS, vol. 9850, pp. 329–347. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-45348-4\\_19](https://doi.org/10.1007/978-3-319-45348-4_19)
28. Wüst, K., Gervais, A.: Do you need a blockchain? In: *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*, pp. 45–54 (2018). <https://doi.org/10.1109/CVCBT.2018.00011>
29. Xu, X., Pautasso, C., Zhu, L., Gramoli, V., Ponomarev, A., Tran, A.B., Chen, S.: The blockchain as a software connector. In: *Proceedings - 2016 13th Working IEEE/IFIP Conference on Software Architecture, WICSA 2016*, pp. 182–191. IEEE (2016). <https://doi.org/10.1109/WICSA.2016.21>

## 7 EtherTwin: Blockchain-based Secure Digital Twin Information Management

---


Current status:	Published
Journal:	Information Processing & Management, Volume 58, January 2021
Date of acceptance:	October 26, 2020
Full citation:	PUTZ, B., DIETZ, M., EMPL, P. AND PERNUL, G. EtherTwin: Blockchain-based Secure Digital Twin Information Management. <i>Information Processing &amp; Management</i> 58,1 (2021), 102425.
Authors contributions:	Benedikt Putz 40% Marietheres Dietz 40% Philip Empl 10% Günther Pernul 10%

---

**Journal Description:** Information Processing and Management (IPM) presents cutting-edge original research at the intersection of computing and information science. The double-blind peer reviewed journal primarily targets the scholarly audience but also practitioners.

Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

## Information Processing and Management

journal homepage: [www.elsevier.com/locate/infoproman](http://www.elsevier.com/locate/infoproman)

# EtherTwin: Blockchain-based Secure Digital Twin Information Management

Benedikt Putz<sup>\*</sup>, Marietheres Dietz, Philip Empl, Günther Pernul

Universitätsstraße 31, Regensburg 93051, Germany

### ARTICLE INFO

#### Keywords:

Distributed ledgers  
Blockchain  
Digital twin  
Industry 4.0  
Decentralized application

### ABSTRACT

Digital Twins are complex digital representations of assets that are used by a variety of organizations across the Industry 4.0 value chain. As the digitization of industrial processes advances, Digital Twins will become widespread. As a result, there is a need to develop new secure data sharing models for a complex ecosystem of interacting Digital Twins and lifecycle parties. Decentralized Applications are uniquely suited to address these sharing challenges while ensuring availability, integrity and confidentiality. They rely on distributed ledgers and decentralized databases for data storage and processing, avoiding single points of trust. To tackle the need for decentralized sharing of Digital Twin data, this work proposes an owner-centric decentralized sharing model. A formal access control model addresses integrity and confidentiality aspects based on Digital Twin components and lifecycle requirements. With our prototypical implementation EtherTwin we show how to overcome the numerous implementation challenges associated with fully decentralized data sharing, enabling management of Digital Twin components and their associated information. For validation, the prototype is evaluated based on an industry use case and semi-structured expert interviews.

## 1. Introduction

Industrial control systems (ICS) such as supervisory control and data acquisition (SCADA) systems, human machine interfaces (HMI), programmable logic controllers (PLCs) and other field devices are able to control physical processes within industrial environments. Traditionally, they form the core of industrial infrastructures. In the course of the Industry 4.0, however, these industrial environments further converge with information technology [Rubio, Roman, and López \(2017\)](#). For instance, sensors measuring the conditions of the respective physical processes to control are increasingly installed. This sensor data as well as the ICS systems are integrated to corporate IT systems in order to centrally analyze and manage information about the industrial environment.

The Digital Twin (DT) presents one of the key concepts reflected in the Industry 4.0 movement. In Industry 4.0, the DT can generally be defined as a digital representation of an industrial asset over its entire lifecycle [Boschert, Heinrich, and Rosen \(2018\)](#). To represent and to further monitor its counterpart, the DT incorporates all kinds of asset-relevant information. This includes a multitude of generated sensor data from Industry 4.0 assets, which are united in DTs. Depending on the underlying asset, different lifecycles are covered by the digital twin. From this follows that different participants involved in the lifecycle might provide information for the DT

<sup>\*</sup> Corresponding author.

E-mail addresses: [benedikt.putz@ur.de](mailto:benedikt.putz@ur.de), [benedikt.putz@wiwi.uni-regensburg.de](mailto:benedikt.putz@wiwi.uni-regensburg.de) (B. Putz), [marietheres.dietz@ur.de](mailto:marietheres.dietz@ur.de) (M. Dietz), [philip.empl@ur.de](mailto:philip.empl@ur.de) (P. Empl), [guenther.pernul@ur.de](mailto:guenther.pernul@ur.de) (G. Pernul).

<https://doi.org/10.1016/j.ipm.2020.102425>

Received 15 May 2020; Received in revised form 26 October 2020; Accepted 26 October 2020

Available online 9 November 2020

0306-4573/© 2020 Elsevier Ltd. All rights reserved.

**Table 1**

Comparison of blockchain-based DT-related approaches by considering organizational as well as implementation characteristics. ○ not considered, ● partially considered, ● fully implemented.

	Huang et al. (2020)	Hasan et al. (2020)	Angrish et al. (2018)	Dietz et al. (2019)
DT definition	product	any asset	machine events	any asset
Components	○	○	○	○
Lifecycle phases	early & medium	early	medium	early & medium
BC suitability	○	○	○	●
Implementation	○	●	●	○
Open Source	○	●	○	○
Blockchain	unknown	Ethereum	Ethereum	unknown
Off-chain storage	○	●	●	●
Encryption	●	○	○	●
Access control	○	○	●	●
User Interface	○	○	○	●

or need to gather data managed by the DT (*DT data sharing*) Dietz and Pernul (2020a).

To achieve information management and sharing in Industry 4.0 with DTs, some obstacles arise. To manage DT data, involved lifecycle parties need access to the DT. Although the different parties participating in these processes work together, they each pursue different goals. Consider the lifecycle parties involved in an industrial plant, where a DT incorporates all relevant data. The manufacturer of the plant's motors should not gain access to the data about the plant's current status, but should get feedback whenever the motor is maintained in order to optimize the motor's construction and enhance its manufacturing process. In contrast, the maintainer of the plant's motors should only get access to the motor's current status and the components the maintainer is not responsible for, but not to any other component's status of the plant. Thus, the trust when sharing data via the DT is not given per default Malakuti and Grüner (2018). As a result, confidentiality and access control issues arise Dietz and Pernul (2020b). These issues cannot be resolved with a centralized authority, especially in multi-tenant and large-scale environments Esposito, Tamburis, Su, and Choi (2020).

This work addresses the lack of trust and security among multiple parties in DT data sharing by focusing on the following research question:

**RQ1.** How can the data of Digital Twins be shared among multiple untrusted lifecycle parties while ensuring confidentiality, integrity and availability?

Blockchains and their smart contracts possess various characteristics that can support the security of data sharing Berdik, Otoum, Schmidt, Porter, and Jararweh (2021). For instance, single and multi-party authentication can be implemented in a decentralized way Khan and Salah (2018) – without requiring trust in a central party. Moreover, blockchain solutions enable decentralized management of an asset's lifecycle and supply chain Khan and Salah (2018). Blockchain solutions rely on Decentralized Applications (DApps), user-friendly web-based interfaces to interact with blockchains and their smart contracts. These characteristics offer a novel opportunity to solve the aforementioned obstacles in DT information management.

In this work, we show why a blockchain-based solution is suitable for DT data sharing and propose a blockchain-based information management solution for the DT and the involved lifecycle parties. We go beyond the state-of-the-art research by including DT components with fine-grained access control and providing scalability for sensor data sharing. Finally, our approach is evaluated with a DApp prototype implementation (EtherTwin), an industry use case, expert interviews as well as performance and cost measurements.

The remainder of this work is organized as follows. We introduce related work in Chapter 2. The background of our research is laid in Chapter 3. Afterwards, we outline the logical design of our concept in Chapter 4. Chapter 5 describes the implementation of our EtherTwin DApp, which is subsequently evaluated in Chapter 6. Chapter 7 discusses our prototype in respect to the evaluation and future work. Finally, a conclusion is drawn in Chapter 8.

## 2. Related work

As DT research began to grow only during recent years, current works mainly propose theoretical frameworks. To date, various works mention the issue of the DT requiring strong security Kaur, Mishra, and Maheshwari (2020); Rubio et al. (2017); Uhlemann, Lehmann, and Steinhilper (2017), however applicable solutions are not provided yet. Especially, the secure management of DT data storage and exchange is important for practical use Malakuti and Grüner (2018).

There have been few other works exploring the blockchain-based accompaniment of assets in supply chain processes with DTs Mandolla, Petruzzelli, Percoco, and Urbinati (2019) and smart objects Meroni and Plebani (2018). Still, a comprehensive implementation of decentralized and secure data sharing for DTs is missing. Moreover, past works have shown the feasibility and advantages of blockchain-based access control for decentralized data sharing Di Francesco Maesa, Mori, and Ricci (2019); López-Pintado, Dumas, García-Bañuelos, and Weber (2019). However, there is no blockchain-based access control model tailored to the requirements of the DT lifecycle.

In the following, we compare previous works that focused on blockchain-based data management in connection with the DT. Table 1 summarizes the comparison by considering organizational aspects of data management as well as implementation characteristics: The first few characteristics are of organizational nature, the following are implementation-related.

**Table 2**

General lifecycle characteristics (involved parties and data) of an industrial asset. Potentially involved lifecycle parties are highlighted in *italic*.

	Early Phases	Medium Phases	Later Phases
<b>Lifecycle phases</b>	Idea, Planning, Manufacturing	Operation, Maintenance	Demolition, End of Existence
<b>Accruing Data</b>	<ul style="list-style-type: none"> <li>• Sketches</li> <li>• Blueprints</li> <li>• Manuals</li> <li>• Design models</li> </ul>	<ul style="list-style-type: none"> <li>• Sensor data</li> <li>• System logs</li> <li>• Maintenance reports</li> <li>• Simulations</li> </ul>	<ul style="list-style-type: none"> <li>• Condition of the components</li> <li>• Component's location</li> </ul>
<b>Involved parties</b>	Owner, Manufacturers, Distributors	Owner, <i>Manufacturers, Distributors, Maintainers</i>	Owner, <i>Manufacturers, Distributors, Maintainers</i>

So far, few works have tackled blockchain-based data management in connection with DTs. Angrish et al. develop a prototype for a peer-to-peer network of manufacturing nodes Angrish, Craver, Hasan, and Starly (2018). Hasan et al. propose a blockchain-based data management approach for the DT creation process Hasan et al. (2020). Huang et al. Huang, Wang, Yan, and Fang (2020) propose a management approach to store all relevant DT data on a custom blockchain. Dietz et al. propose a conceptual approach for blockchain-based DT data management Dietz, Putz, and Pernul (2019).

Thereby, the organizational aspects of the related works vary. While Huang et al. consider the DT being a product Huang et al. (2020), Angrish et al. define the DT as a mere collection of machine events Angrish et al. (2018). The majority of the works, as well as our work, see the DT as a representation of any asset Dietz et al. (2019); Hasan et al. (2020), be it a system, product or another physical object. Moreover, so far, none of the related works have tackled the DT as being a complex representation of an asset with sub-components. In contrast, we include components of the DT in our data model. In terms of lifecycle phases (cf. Table 2), we are the first to consider the DT management as beneficial for later lifecycle phases as well. Works to date have tackled either early Hasan et al. (2020), medium Angrish et al. (2018) or both of these phases Dietz et al. (2019); Huang et al. (2020). Additionally, we are the first to fully investigate the suitability of blockchains for DT data management by following a research methodology, while other works either neglect this aspect Angrish et al. (2018); Hasan et al. (2020); Huang et al. (2020) or only mention, but do not describe a method Dietz et al. (2019).

To date, prototypical implementations have been either neglected Dietz et al. (2019); Huang et al. (2020) or only partially accomplished Angrish et al. (2018); Hasan et al. (2020). Our work is the first to fully implement a proposed DT data sharing approach. Next to our work, only one other work has made the implementation open source Hasan et al. (2020). All works with an implementation part, however, make use of the Ethereum blockchain. In terms of off-chain storage, related work either do not suggest using it Huang et al. (2020), or suggest to use off-chain storage, but do not implement this part Angrish et al. (2018); Dietz et al. (2019); Hasan et al. (2020). In our EtherTwin prototype, a fully implemented off-chain storage is present. Encryption is proposed in two of the four related works Dietz et al. (2019); Huang et al. (2020), but is not described in detail and implemented – in comparison to our work. Likewise, access control mechanisms are mentioned in two works Angrish et al. (2018); Dietz et al. (2019) but are also not implemented. A user interface is only suggested by a single work Dietz et al. (2019), but we are the first to design and implement one.

The present work develops a component-based data model and an access control model for common lifecycle participants. To summarize, we contribute to DT and blockchain research by providing:

- **fine-grained access control** for DT data sharing in a decentralized setting without a trusted third party (TTP), ensuring confidentiality through encryption
- full-featured **open source prototype** EtherTwin based on blockchain design patterns and state of the art DApp technologies (Ethereum, Swarm) with performance/cost measurements
- evaluation based on an **industry use case** and expert interviews

### 3. Background

The background of this work is divided into three sections. Section 3.1 describes the foundations of DT research. Subsequently, the background of DApps is laid in Section 3.2.

#### 3.1. Digital twins

The DT is an emerging paradigm focusing on an enterprise asset – usually, a system, product or process, along its lifecycle Boschert et al. (2018). Its core goal is to virtually represent this asset as close to reality as possible Boschert et al. (2018). The lifecycle phases covered by a DT strongly depend upon its corresponding asset. Nevertheless, common early phases are *Idea, Planning* and *Design*, while an asset's *Operation* can be considered one of the medium phases and the asset's *Demolition* is one of the final phases Dietz and Pernul (2020b). Thereby, each phase can span many years. For instance, planning a complex asset like global satellite networks could take up to 10 years until *Operation*, while some legal regulations may command to safely store the asset after its decommission. Especially, these long and safety-oriented lifecycle phases require a tamper-proof data storage solution. By including various data sources and by

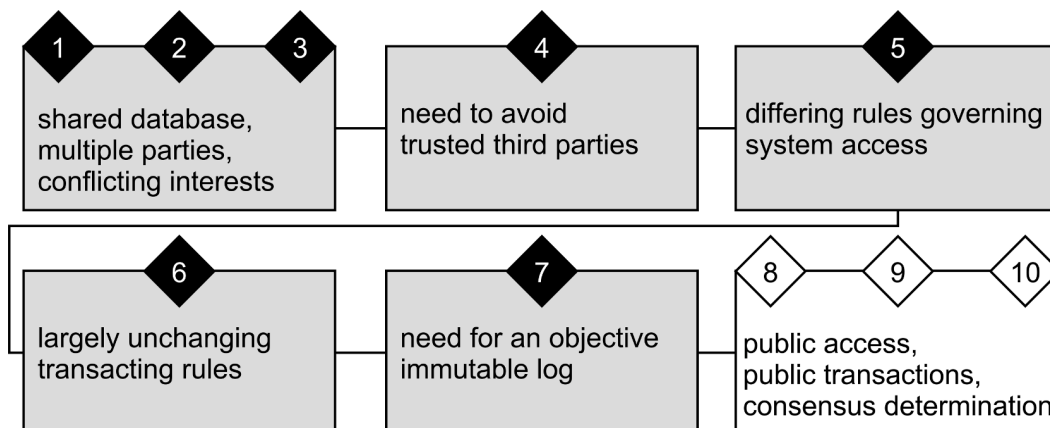


Fig. 1. Blockchain decision path by Pedersen et al. Pedersen et al. (2019).

integrating the multiple parties involved in these lifecycle phases, the DT unifies asset-specific data from previously separated domains Ríos, Hernández, Oliva, and Mas (2015). For instance, the asset's composition, sensor data of the asset's environment as well as simulation models can be included centrally in a DT Dietz and Pernul (2020a). This further promotes the complete traceability of assets and their components, especially if the assets (e.g. industrial plants, cars) comprise components from several manufacturers. Thereby, feedback loops across different lifecycle phases can be realized that support the concerned lifecycle parties when looking for improvement of their components Boschert et al. (2018). For instance, manufacturers can get insights from the operational phase of the asset and draw conclusions about the effectiveness of their components.

For the remainder of this work, we put the person that owns the physical asset in the role of the Owner. Whenever this kind of Owner is meant, it is written in capitals. We further claim that the ownership of the physical asset implies the ownership of the digital twin. Otherwise, two different parties respectively owning either one of them would commonly not trust each other. Thus, the interaction of digital twin and physical twin would not be achieved.

Table 2 summarizes the common lifecycle phases and points out the potentially involved lifecycle parties and the accruing data in the respective phases. Note that the data is continuously transformed along the lifecycle phases. For example, sketches of an industrial asset might exist from the *Idea* phase, transform into a blueprint in the *Design* phase. Also, design models might be created in the *Design* phase and elaborated towards fully-fledged simulations in the *Operation* and *Maintenance* phase. In terms of the involved parties, italicized parties are only potentially involved. For instance, consider the Owner sketching the asset during the *Idea* phase. Afterwards, the manufacturer elaborates this sketch towards a blueprint (*Design*) and manufactures the asset (*Manufacturing*). Later on, the Owner commissions the maintainer to put the asset into *Operation*.

Nevertheless, there are still some obstacles to overcome. Commonly, an industrial asset represents a complex system, product or process. As a consequence, a multitude of parties are involved. Consider an industrial plant consisting of various ICSs. Each of these systems potentially has its own manufacturer and in business life, they might be competitors. This leads to enormous trust issues, and towards current practices of each lifecycle party building their own DT Malakuti and Grüner (2018). Meanwhile, this practice contradicts the very idea of DTs. Furthermore, it results in the disappearance of the DT's core benefits like feedback loops to other lifecycle phases and parties. To overcome current malpractices and to motivate users to share their data among parties with different trust levels, our research aims to provide a strong platform with sufficient security (i.e. access control mechanisms) among untrusted parties.

### 3.2. Blockchain and decentralized applications

To address the complex issues of the DT sharing ecosystem, we investigate if blockchain technology is suitable. Pedersen et al. propose a ten-step decision path to determine if blockchain is a good fit Pedersen, Risius, and Beck (2019). The ten requirements are outlined in Fig. 1. For the DT lifecycle, there are multiple parties with the need for a shared database, which may have conflicting interests and thus, varying trust levels (steps 1–3). While in theory the lifecycle parties could rely on a TTP service, the dynamics and variety of DT data sharing hamper the management through a TTP. As Table 2 highlights, various data and data types are involved with varying velocity and integrity requirements. Integrity of stored data is an especially important security concern in IoT environments Zhao, Chen, Liu, Baker, and Zhang (2020). A TTP represents a single point of failure and an attack could interfere with the integrity of the data, making it preferable to avoid third parties. Related research on data auditing has shown that blockchain technology is able to remove the need for trusted third parties Li, Wu, Jiang, and Srikanthan (2020), which suggests that it could be a good fit for our work. (step 4). Moreover, the participants of the lifecycle require different access privileges depending on their role and characteristics, which means there are differing rules governing system access (step 5). Although system access rules differ in practice, the rules of transacting with DT data do not change frequently (step 6). The blockchain's immutable log is helpful to ensure integrity and traceability of all

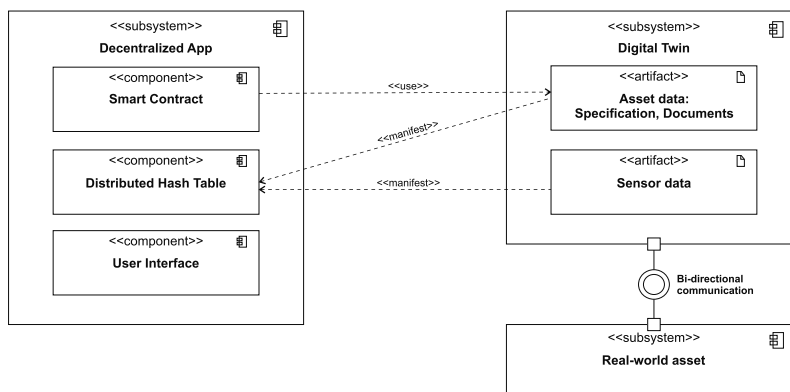


Fig. 2. Component Diagram describing the Digital Twin Sharing Context.

changes to the DT data, for example, in case of security issues or malfunctions. Furthermore, the documentation of changes made to data items is required to meet compliance requirements for some DTs (step 7). The need for public access largely depends on the DT's underlying asset and industry (steps 8 - 10). We do not make an assumption in this regard and design our application to work with both permissionless and permissioned networks.

After determining that blockchain is suitable, we explain the software components required for building a DT information management application. DApps are a new paradigm for developing distributed applications Xu, Weber, and Staples (2019). Application logic is fully decentralized, since front end code runs in the user's browser and back end code runs in smart contracts on the blockchain nodes. Decentralization comes with the advantage of full transparency of the application code as well as auditability of changes to a smart contract state. Dynamic smart contract access control models can be used to authorize state changes Di Francesco Maesa et al. (2019).

Full replication of blockchain data necessitates storing complex data elsewhere Baig and Wang (2019), leading to the concept of off-chain storage. A common approach is to use Distributed Hash Tables (DHTs), since they fit the decentralized paradigm well. Data items are content-addressed and replicated within the network based on a routing layer. Modern DHTs such as Swarm<sup>1</sup> are based on the established and secure DHT routing technology S/Kademlia Baumgart and Mies (2007) and integrate well with blockchains such as Ethereum<sup>2</sup>.

Blockchain smart contracts also need to ensure sufficient access control to prevent unauthorized modification of smart contract state. Numerous authors have developed access control concepts based on smart contracts. These are based on the existing access control models role-based (RBAC) Cruz, Kaji, and Yanai (2018) or attribute-based access control (ABAC) Rouhani, Belchior, Cruz, and Deters (2020), but there are also proposals for ciphertext-policy attribute-based encryption Badsha, Vakiliinia, and Sengupta (2020). Zhang et al. present an access control framework for the Internet of Things (IoT) supporting flexible access control methods Zhang, Kasahara, Shen, Jiang, and Wan (2019). Rouhani et al. also provide a comprehensive overview of smart contract based access control approaches Rouhani et al. (2020).

#### 4. System model

The following sections describe the logical structure of our DApp. Section 4.1 provides an overview of the DApp's entities. Section 4.2 explains the twin and its subparts, while Section 4.3 focuses on the authorization of participating parties.

##### 4.1. Overview

To capture context, the component diagram in Fig. 2 provides an overview of the DT sharing approach. In our system model, a component diagram defines physical as well as logical components and their dependencies. Therefore, it is well suited to put software architectures like our DApp into context.

Fig. 2 illustrates the connection between real-world asset, DT and the developed DApp. These components represent a greater architectural unit (subsystems). The first two subsystems present the sole DT paradigm, consisting of the DT and its real-world asset connected by the bi-directional communication interface. One of the two artifacts within the DT is asset data, e.g. the specification of the asset with its compositional structure and documents about the asset. The other artifact is the sensor data produced in the asset's environment. To enable data sharing, the DApp is added. It includes the components Smart Contract, DHT and User Interface. The dependency relations show the association of the DT data to the DApp. For instance, the Smart Contract requires the specification data of the asset in order to be built (usage dependency). Moreover, the shared DT data is stored in the DHT: The manifest dependency shows

<sup>1</sup> swarm.ethereum.org

<sup>2</sup> ethereum.org






Fig. 3. Entity Relationship Model of the DApp.

that the logical DT artifacts physically manifest in the DHT of our proposed DApp. Finally, the user interface component provides access to the data for all participating lifecycle parties.

#### 4.2. Entity relationship model

Fig. 3 illustrates entities and relationships for DT sharing with a DApp. The dashed borders show a logical grouping into three main components: *Registry*, *Authorization* and *Twin Data*. To keep track of all available twins, a single access point is needed, referred to as the *Registry*. Similarly, the *Authorization* group of entities represents the access control model, which is explained in detail in Section 4.3. *Twin Data* is derived from DT sharing requirements elaborated in our previous work Dietz et al. (2019).

The on-chain entities contain metadata about the DT. The main entity of a DT is the *Specification*, comprising the *Components* of the real-world asset it is representing. *Sensors* and other data (abstracted with the term *Document*) are managed by associating them with the corresponding component. Moreover, for each DT *External Sources* such as legacy systems can be integrated. These can provide additional data to the already incorporated documents and sensor feeds.

Off-chain entities (*Specification Version*, *Sensor Entries*, *Document Version*) contain full data and are linked to on-chain entities, as indicated by 1:n relationships in Fig. 3.

#### 4.3. Access control

In order to share data securely, an authorization and access control policy is required. This way access to data items can be restricted to certain parties. For instance, a maintenance report of an asset's component (e.g. of a PLC) should only be shared with the lifecycle parties of this component (e.g. the PLC's manufacturer). Access control addresses this need by restricting the user operations for data objects. In our approach we follow a hybrid access control model, combining RBAC and ABAC. While a role refers to a certain organizational function, a particular attribute refers to a specific characteristic of a user. During the DT lifecycle, each user interacts with certain twin components, which constitute the user's attributes in our model. While roles are predefined, these attributes allow access control on-the-fly.

Our proposed approach is modeled after the RBAC-A (role-centric) combination strategy, where attributes are applied to constrain RBAC Coyne and Weil (2013); Kuhn, Coyne, and Weil (2010). Thereby, the user's assigned role defines the base permissions, while the user's additional attributes can further limit these permissions. The exclusive use of ABAC would create an unnecessary overhead of rules, which control the access of the user. This would further increase complexity, both in terms of attribute combination for the user and the subsequent access granting decisions. Our hybrid access control model for DT data sharing upholds essential RBAC advantages (e.g. ease of user provisioning) and enhances flexibility by integrating attributes.

To provide a profound basis for later implementation, we elaborate a formalism of the access control used in our DT sharing approach. Italicized terms refer to entities from Fig. 3. Every sharing party is considered a *User*  $U := \{u_1, \dots, u_n\}$ . In our hybrid access control approach, each user can have one *Role*  $R := \{r_1, \dots, r_n\}$  as well as several *Attributes*  $A := \{a_1, \dots, a_n\}$  per DT. *Components*  $C := \{c_1, \dots, c_n\}$  serve a special purpose in this access control model, as they are used for modeling *Attributes*:  $A := a_1, \dots, a_n \mid a_i = c_1 \vee \dots \vee c_n$ .

To continue, *Permissions*  $P := \{p_1, \dots, p_n\}$  are mainly derived from the user's role but also from its attribute(s). This underlines the hybrid RBAC-A mode in a role-centric realization Kuhn et al. (2010): Roles determine the basic permissions, while some of these permissions are limited by the users' attribute(s). The permissions usually specify the access to an object  $O := \{o_1, \dots, o_n\}$  and the allowed operation  $Op := \{op_1, \dots, op_n\}$ . Objects are always associated to a component to link the asset-relevant data to the component they belong:  $o \rightarrow c \mid o \in O \wedge c \in C$ . This results in the following definition of Permissions:  $P = Op \times O$ , whereby it can be concluded that  $p \rightarrow c$ .

The n-m relation of users to roles is expressed by  $UR = U \times R$ . Likewise, the user to attributes relation can be described as  $UA = U \times$

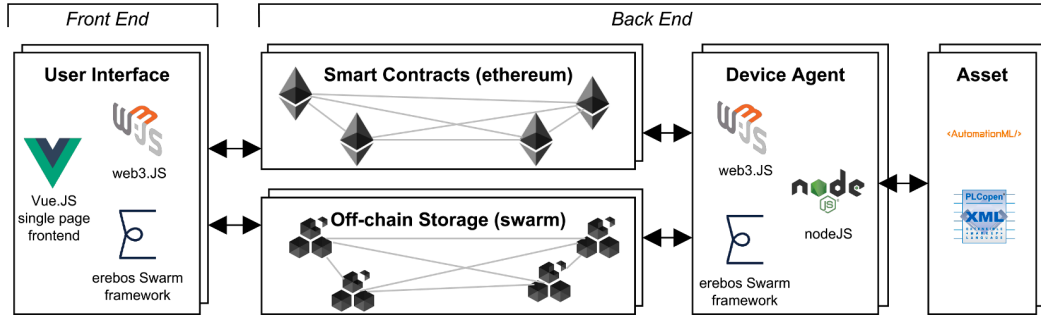


Fig. 4. Technologies used in our prototypical implementation of DT data sharing.

A. Mapping the role-attribute combination to a user results in:

$$\text{assign\_users}(r, a) = u \in U \mid (u, r) \in UR \wedge (u, a) \in UA$$

Thereby, the mapping  $M^{UR}$  describes the set of actual assignments of users to roles, while  $M^{UA}$  determines the set of assigned attributes to users. Similar to the mapping above, the m-n relation between permissions and roles are specified by  $PR = P \times R$ . Finally, permissions are mapped to role-attribute combinations while the attribute restricts the role permission and  $M^{PR}$  describes the set of actual assignments of users to permissions:

$$\text{assign\_permissions}(r, a) = p \in P \mid (p, r) \in PR \wedge p \rightarrow c \mid c \equiv a$$

## 5. Decentralized application architecture

Based on the system model elaborated in Section 4, we choose appropriate technologies and standards to implement a DApp for DT data sharing in Section 5.1. We implement our entities by leveraging several blockchain design patterns (Section 5.2). To showcase the inner workings of the DApp, the most important data flows for DT management are described in Section 5.3. The concomitant access control implementation is detailed in Section 5.4.

### 5.1. Technology selection

For the DApp prototype we rely on the Ethereum blockchain, which is commonly used for research, e.g. in blockchain-based business process management Haarmann, Batoulis, Nikaj, and Weske (2018). It offers the Turing complete smart contract programming language Solidity and has a large developer community, resulting in advanced development tools and vulnerability scanners Ayman, Aziz, Alipour, and Laszka (2019).

Fig. 4 depicts the technical architecture of the EtherTwin DApp. A **User Interface** simplifies the interactions of the DT lifecycle participants, such as creating twins and uploading data. For trustless interaction with the blockchain it is implemented using the single page application JavaScript framework Vue.js<sup>3</sup> – a server is only needed to serve static assets. The module ethereumjs-wallet is used for managing the user's blockchain account, providing access to the user's public and private key. Key pairs are dynamically created on first access and stored in the browser's local storage for future visits.

Web3.JS is used to send transactions signed with the private key to the **Smart Contracts** on the Ethereum blockchain. The front end is connected to an Ethereum blockchain node through a WebSocket connection. WebSockets improve performance over HTTP connections by providing a two-way communication channel between the client and the Ethereum node. This avoids the need to set up individual HTTP connections for each request Fette and Melnikov (2011). WebSockets also enable subscription to smart contract events (publish-subscribe style), which is utilized by the Device Agent for synchronization purposes. During development, we observed a significant speed up in page load times after switching to an RPC connection based on WebSockets.

The erebos module<sup>4</sup> reuses the blockchain account to upload data to the **Off-chain Storage** based on the Swarm DHT. While Swarm is mostly known for its permissionless test network, it can also be deployed as a private DHT with a fixed set of peers. Swarm reuses Ethereum accounts as its identity system, which simplifies its integration as off-chain storage. Additionally, the data types used in both systems are compatible: References to Swarm data are encoded as 32 byte SHA3 hashes, which can be stored in a Solidity bytes32 variable in the smart contract. For dynamic content, Swarm provides Feeds. Feeds have a fixed address specified by user (Ethereum account) and topic (any SHA3 hash). They can only be updated by their owner with a public-key signature. Any Swarm user can read-access the most current and past updates. This concept is useful for sharing file keys and real-time sensor data under a fixed address, despite Swarm's content-addressed storage. Ethereum Swarm Contributors (2019)

<sup>3</sup> vuejs.org

<sup>4</sup> erebos.js.org

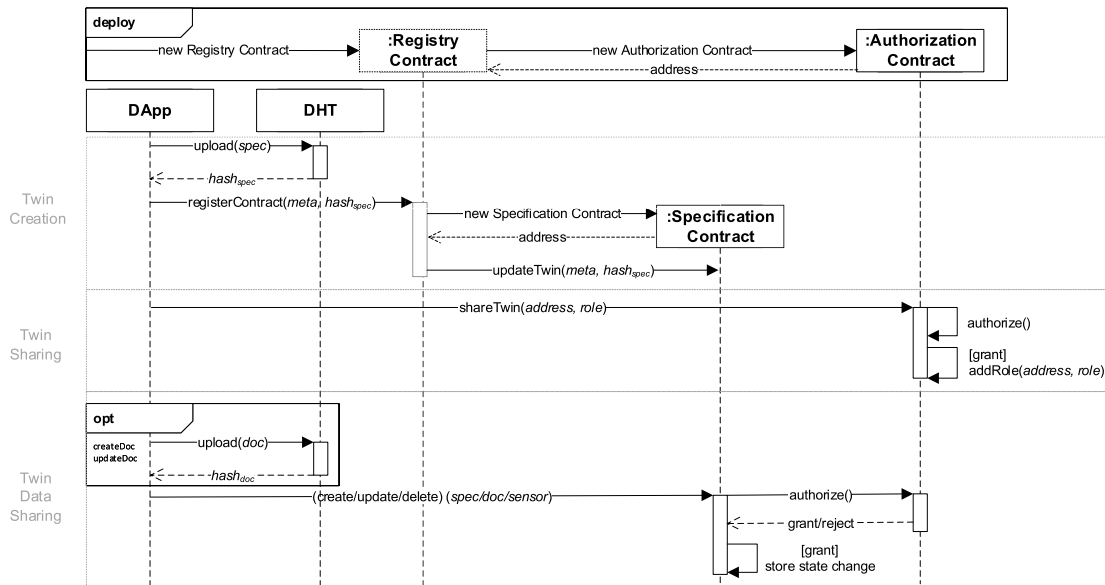


Fig. 5. Sequence diagram showing initial deployment and user interactions with the smart contracts and DHT.

The **Device Agent** bidirectionally synchronizes the DT's underlying **Asset** with the decentralized DT on Ethereum and Swarm. It runs as a node.js<sup>5</sup> background process and monitors new sensor data from the asset, which is then uploaded to Swarm.

Like other authors creating DTs [Eckhart and Ekelhart \(2018\)](#); [Schroeder, Steinmetz, Pereira, and Espindola \(2016\)](#), our work relies on Automation Markup Language (AML), which is defined in the industrial standard IEC 62714. AML describes the specification of the asset including its components and their logic. Components are derived by parsing the AML-based asset specification.

## 5.2. Design patterns

Several blockchain application design patterns [Xu, Pautasso, Zhu, Lu, and Weber \(2018\)](#) are used in our prototype to address the requirements of DT data sharing. A *Contract Registry* pattern keeps track of individual DT contracts. A *Factory Contract* pattern is used to instantiate individual DT sharing instances. The access control model from [Section 4.3](#) is implemented using the *Embedded Permission* pattern and implemented in the separate Authorization contract. The *Multiple Authorization* pattern is used to ensure that all sharing parties agree before changes to a DT contract are made. The *Off-chain Data Storage* pattern is used to meet the data volume and latency requirements. The Device Agent implements the *Reverse Oracle* pattern to mediate between the industrial asset and the distributed ledger. It monitors events occurring on the asset and publishes sensor data for authorized parties. Additionally, the agent is responsible for managing and distributing the symmetric file keys used for encrypting off-chain data, as detailed in [5.4](#).

## 5.3. Data flow

[Fig. 5](#) shows how the contracts interact during the deployment, twin creation and sharing phases of the DT lifecycle.

**Deployment.** Initially, the Registry and Authorization contracts are deployed by the blockchain consortium initiator. The Specification contract template is deployed, but not yet instantiated as it is twin-specific.

**Registration.** When a user first opens the app, a new Ethereum account is created, represented by an Ethereum public-private key pair. The public key is shared off-chain by publishing it on the account's Swarm Feed. This avoids on-chain storage costs and allows anyone to retrieve the public key from the corresponding Swarm Feed. To improve usability and to avoid the need to share addresses out-of-band, we also register a mapping of the user's Ethereum address to a username on the Authorization contract.

**Twin Creation and Sharing.** On twin creation, the Owner provides a specification, which is parsed to extract the twin's components. A transaction is sent to the Registry Contract, which creates a new Specification Contract instance based on the provided data. In the authorization contract, the access control attributes of the newly created twin are initialized with the provided components. The AML-formatted specification is stored on the DHT and included with a hash reference. After a twin has been created, the Owner may share it by adding a role to the lifecycle participant's blockchain account.

**Twin Data Sharing.** Each transaction intending to create, update or delete an entity of the twin must first be authorized through the Authorization contract. It should be noted that deletion only removes the entry from the current state; the state's history is

<sup>5</sup> nodejs.org

**Table 3**

Role mapping for entity Create/Read/Update/Delete permissions. ~: Permission depends on presence of component attribute.

Permission	Twin				Document				Sensor			
	C	R	U	D	C	R	U	D	C	R	U	D
Device	x	✓	x	x	x	✓	x	x	x	✓	✓	x
Owner	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Manufacturer	x	✓	x	x	~	~	~	x	x	~	x	x
Maintainer	x	✓	x	x	~	~	~	x	~	~	~	x
Distributor	x	✓	x	x	~	~	~	x	x	x	x	x

preserved on the blockchain. If the action is authorized, the corresponding state change is registered in the smart contract state. For documents and specification version updates, the same procedure applies, except that metadata such as the filename remains unchanged.

**Sensor Feed Updates.** Due to latency requirements and to reduce the number of costly smart contract transactions, sensor data is shared off-chain. After a sensor is registered on-chain, the Device Agent connects to the corresponding component sensor and subscribes to its sensor data. Each new sensor entry is encrypted with the component's sensor encryption key and published to the sensor's Swarm Feed.

#### 5.4. Access control implementation

In the following the decentralized implementation of the formal access control model detailed in Section 4.3 is described.

**Authentication.** Authentication is based on blockchain accounts, which consist of a private key and an address. Identities are represented by addresses, which are created by hashing the public key. They are used for signing transactions and sharing confidential data intended for specific participants.

**Authorization.** Data stored on-chain is implicitly accessible to all participants storing the blockchain. For this reason, only metadata and off-chain references are stored in smart contract state. State change transactions require authorization by the Authorization contract authorization, with component-based entities (documents, sensors) also requiring the corresponding component attribute. The append-only nature of the blockchain ensures traceability of all changes.

The default mapping of permissions to roles is shown in Table 3. Permissions comprise the CRUD operations for each of the main sharing objects *Twin*, *Document* and *Sensor*. The entities *External Source* and *Specification* do not have separate permissions and instead inherit the *Twin* permissions. Role and attribute mappings are controlled by the DT Owner and can be modified for each individual DT. For example, permissions may be removed from a role or attributes added to a user. These permissions are enforced on-chain by the Authorization contract. Read permissions for off-chain data are enforced by encrypting all data related to off-chain entities (*Specification Version*, *Sensor Entry* and *Document Version*). The encryption key is shared only with authorized users.

**Encryption.** All data is AES-256-encrypted before being uploaded to the Swarm DHT. Permissions are enforced by sharing a public-key encrypted version of the symmetric file key. Since Ethereum uses public keys based on elliptic curve cryptography, we rely on the Elliptic Curve Integrated Encryption Scheme (ECIES). However, Ethereum addresses are hashes of the public key and not the public key itself, which means they cannot be used for encryption. Therefore, participants additionally share their public key on their personal Swarm Feed (identified by their account address).

The file keys are then distributed on a Swarm Feed, which allows dynamic off-chain updates when new users gain permission. For the specification file, the asset Owner manages the file keys. For component-based entities, the file keys are managed by the Device Agent. The Device Agent must be trusted, since it has full access to the asset. It is thus able to enforce on-chain permissions for off-chain data continuously.

The Device Agent creates two unique symmetric keys for each component (for documents and sensors). File key recipients are determined based on roles and attributes stored on-chain. The corresponding algorithm for creating file keys is shown in Algorithm 1. The formal notation is based on Section 4.3. The algorithm must be executed for each twin before any files can be uploaded, since it distributes the symmetric keys needed for encryption. For this reason the Device Agent continuously monitors the blockchain for newly created *twins* managed by its address and associated permission updates. This is achieved by subscribing to contract events emitted by the *Authorization* contract. The Device Agent also subscribes to attribute and role change events. On each event, on-chain permissions are retrieved and the corresponding file keys are added/removed accordingly.

## 6. Evaluation

To evaluate the proposed DApp architecture, we follow a methodological approach based on Venable et al.'s framework for evaluation in Design Science Research Venable, Pries-Heje, and Baskerville (2012). The goal is to ensure both rigor and efficiency of our research. In our ex-post evaluation, we utilize both artificial (prototype, technical simulation) and naturalistic (case study, expert interviews) evaluation methods.

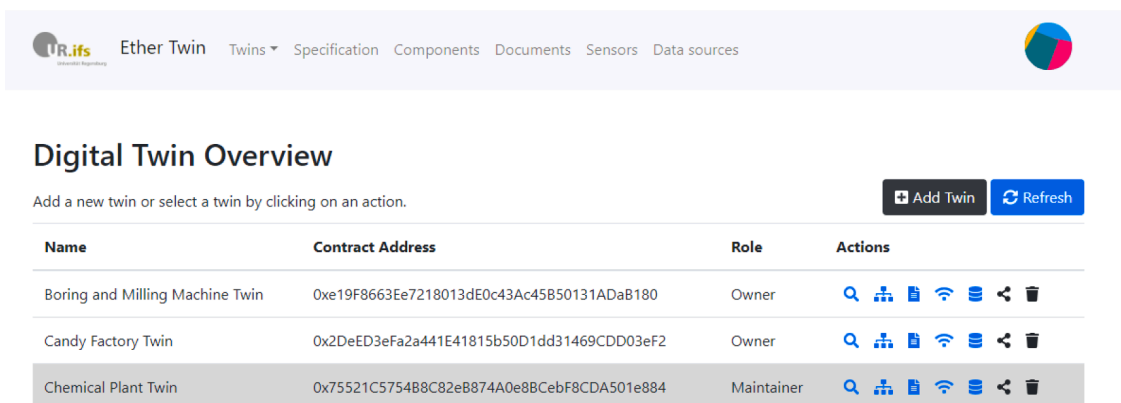


Fig. 6. Screenshots of the prototype's home menu.

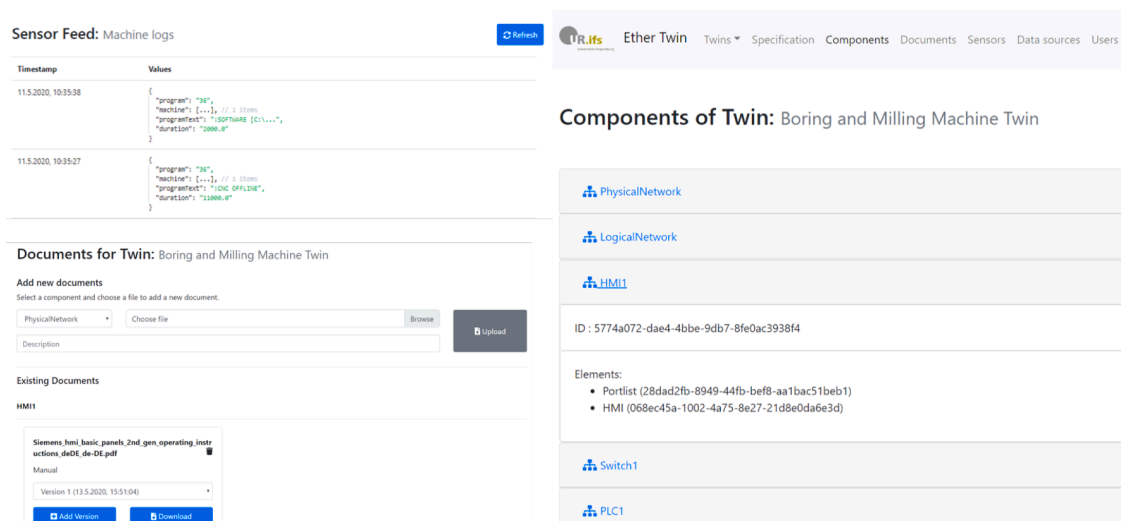


Fig. 7. Screenshots of the prototype's component-based structure and information management.

We first describe the EtherTwin prototype and its user interface in Section 6.1. The prototype is evaluated with several technical experiments concerning latency and cost in Section 6.1. Its practical application is explained via an industry use case in Section 6.3. Finally, we interview several industry experts regarding the prototype's benefits and remaining challenges in Section 6.4.

### 6.1. Prototype

The EtherTwin prototype is available on GitHub<sup>6</sup>, including a video demonstrating the use case illustrated in Section 6.1. It consists of about 3000 single lines of code (SLOC) for the DApp and Device Agent, as well as 400 SLOC for the smart contracts. We analyzed all smart contracts for vulnerabilities using the SmartCheck vulnerability scanner Tikhomirov et al. (2018). Hereafter, screenshots are presented to show the prototype's functionality.

The prototype's start page is illustrated in Fig. 6. It gives an overview of the twins the user is involved with, and shows the role of the user for each twin. The navigation bar shows the available pages for the selected twin that is highlighted in gray. Navigation to the respective pages is handled by clicking on the respective icon in the twin's row. The icon shown on the very right of the navigation bar provides a visual representation of the user's network address. It leads to the account page, containing information about the network and current user.

Fig. 7 contains three screenshots that show the component-based organization of the prototype per twin. The screenshot on the

<sup>6</sup> <https://github.com/sigma67/ethertwin>

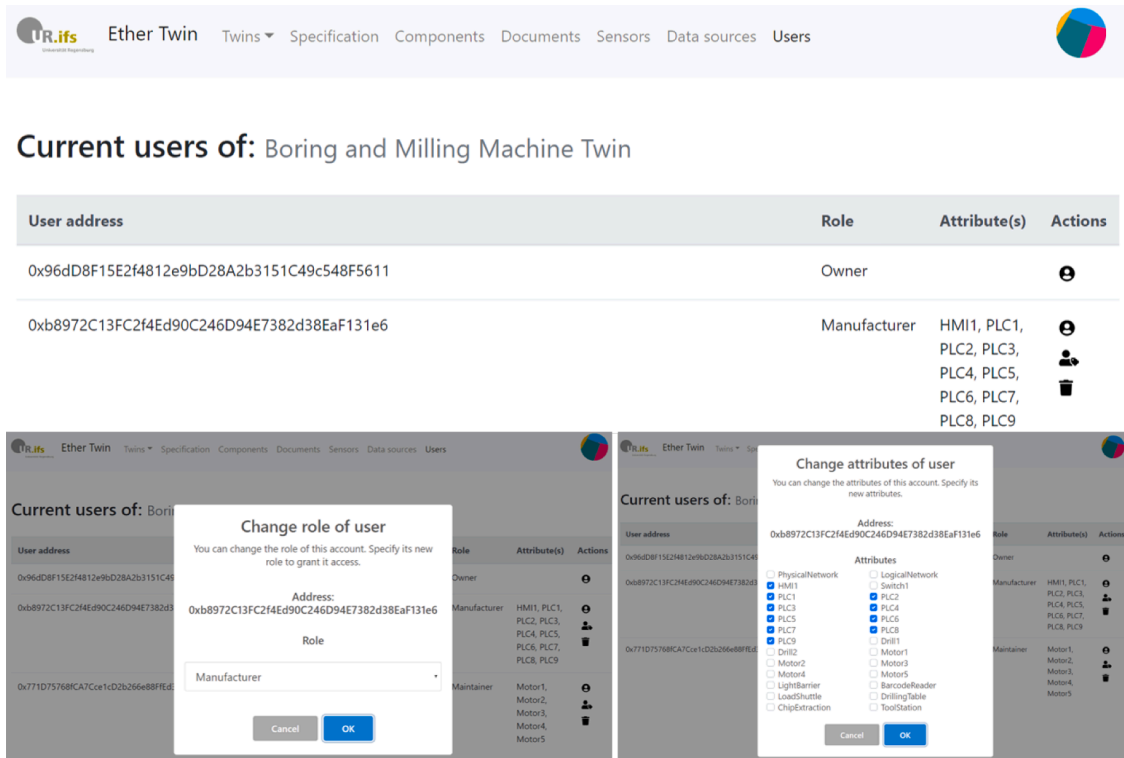


Fig. 8. Screenshots of implemented access control mechanisms.

right shows the composition of the selected asset/twin with its components and sub-components. This structure is parsed from the AML specification, which is required for twin creation. The upper left screenshot illustrates the sensor feed capabilities. The screenshot on the lower left shows the existing documents per twin. Each document is thereby assigned to a component. For each component, documents from each lifecycle-phases can be uploaded. However, users can only download, upload or update a document to a component if they have the respective component attribute in the smart contract. In practice, each user should be assigned the component attributes that the user is involved with in the lifecycle.

Fig. 8 shows the prototype's role and attribute management page. In the EtherTwin prototype, the Owner of a twin can see all other involved users and their lifecycle involvement. Furthermore, the Owner can handle the access to the resources as shown in the screenshots below. The screenshot on the bottom left side shows how the user's role can be changed, while the screenshot on the right side illustrates the adjustment of the user attributes.

Further screenshots of the prototype can be found in Appendix A (Figure A1, Figure A2, Figure A3) and in our GitHub-repository<sup>7</sup>.

## 6.2. Technical experiments

To evaluate the performance of our prototype, we first consider latency of the interactions described in the prototype. Our prototype environment is set up on a Raspberry Pi using Parity Ethereum 2.7.2 and Swarm 0.5.7. The DApp and Device Agent were run on an i7-8550U CPU.

When a new twin is created, the Device Agent must create the twin's symmetric encryption keys before any data can be shared. To evaluate this latency, we benchmarked the runtime of Algorithm 1. The algorithm runs every time a DT is created or its permissions are updated. It only runs once the transaction is included in the blockchain, since it is triggered by smart contract events. The results in Fig. 9 show that the runtime is on the order of one to three seconds. This is sufficient for real-world scenarios, since sharing interactions are not immediate. The runtime is not significantly affected by the number of users the DT is shared with. It increases only slightly with the complexity of the asset specification (number of components).

To ensure user adoption, interactions with the user interface should have low latency. Each time a smart contract transaction is issued, the user needs to wait for a blockchain confirmation. Therefore, we measured the latency of interactions with private and public blockchains in Table 4.

Another aspect relevant for public blockchain deployments are transaction costs for the Ethereum smart contracts. The

<sup>7</sup> <https://github.com/sigma67/ethertwin/tree/master/misc/Screenshots>




Fig. 9. Runtime values for Algorithm 1 for varying numbers of users and components.

Table 4

Latency (ms) and cost (ETH, €) for contract deployment and interactions. Gas price: 10 Gwei, 120 € /ETH.

Action	ms	Gas	ETH	€
Initial Deployment	-	14,548k	0.14548	17.46
Twin Creation	896	4576k	00.4576	5.49
Twin Sharing	353	144k	00.0144	0,17
Specification Version Creation	262	94k	00.0094	0,18
Document Creation	485	254k	00.0254	0,50
Document Version Creation	365	99k	00.0099	0,19
Sensor Creation	374	95k	00.0095	0,18
Attribute Update	276	50k	00.0050	0,10

cryptocurrency costs are shown in Table 4. Sub-second latencies demonstrate that the user experience is fluid, despite client-side encryption/decryption and network delays. Costs are also quite low except for *Twin Creation*, since a new contract is instantiated. The use case costs include each action once, except for *Document Creation*, since two documents are created. In practice, lifecycle participants should decide based on usage cost projections to either jointly run a private network with operational costs, or use the public Ethereum network with the associated transaction fees.

### 6.3. Use case

This use case illustrates how our EtherTwin DApp is used in practice by creating a DT based on real enterprise data that is provided during the Secure Industrial Semantic Sensor Cloud (SISSeC) project<sup>8</sup>. The SISSeC project focuses on introducing Industry 4.0 in small and medium enterprises (SME) and aims at securely unifying and analyzing machine data. The industrial assets targeted in the project are part of the manufacturing process of printed circuit board (PCB) panel prototypes of a small German enterprise. The central goal for the PCB panel manufacturer is to gather all data available about the machines, to unite and analyze the data. Thereby, novel insights such as the determination of flaws in the manufacturing process present the desirable outcome.

Our DApp prototype unifies the available data of an industrial asset throughout all lifecycle phases. In this use case, we create a DT for the boring and milling machine that gouges holes into the PCB panels. The demonstration of the implemented use case can be found online<sup>9</sup> and its manifestation can be gathered from the screenshots of the prototype (Section 6.1 and Appendix A).

At first, the machine specification in the form of an AML-file is implemented to set up the respective smart contracts for the use case (cf. Fig. 2). Then the feed data from the machine is integrated from sensors, ranging from sensors determining the position of the drill to logs of the PLCs concerning the running program. Moreover, we unified asset-relevant documents like manuals of the machines' ICSs<sup>10</sup>

Thereby, the documents are assigned to their corresponding component. For instance, a manual of a Siemens S5 PLC is assigned to the PLCs of the boring and milling machine. Currently, we created user accounts for the PLC's manufacturer, the machine operator and the maintainer of the machine's motors for demonstration. However, there are other users that can be included, e.g. the manufacturer of the motors, the maintainer of the PLCs and HMI or the distributor of the barcode reader.

Based on an interview with the CEO and the CIO of the firm that currently operates the boring and milling machine, we gather that our EtherTwin prototype meets their current needs for central collection of data about their machine. For example, when service is required, the operator usually has trouble providing the right information to the maintainer. However, this information is needed for the maintenance service to bring the right tools and rapidly assesses the machine's state and problem. In their view, EtherTwin poses a solution to this issue. Moreover, they consider the component-based data management a useful strategy that facilitates their search of

<sup>8</sup> <https://www.it-logistik-bayern.de/produktionslogistik/projekt-sissec>

<sup>9</sup> <http://ethertwin.ur.de>. The use case can be tested with a demo Owner account with the private key 1bed7-c10358ece007522558c4801b84424750f5a626ce5c9093411c9fc197a6f, to be entered on the account page (top right icon)

<sup>10</sup> Please note that the SISSeC project is at an early stage, where more data about the machine is still to be gathered.

information about sub parts. Nevertheless, the interviewees also state that EtherTwin's access control mechanisms are very valuable to prevent knowledge drain. Nevertheless, it is uncertain whether lifecycle participants have the required knowledge to install the proposed solution.

This use case shows that our DApp supports our goal of unifying asset-relevant data among its lifecycle with its participants. This results in enabling a feedback loop among the machine's lifecycle phases. The participants of the lifecycle phases can harness this information to optimize their own business.

#### 6.4. Expert interviews

To validate our prototypical implementation of blockchain-based DT information management, we conducted semi-structured interviews with industry experts. The goal of the interviews is to determine the prototype's conformance to practical requirements and to identify potential adoption barriers.

*Participants* We conducted semi-structured interviews with ten industry experts from six different enterprises. The industrial domains the experts have experience with include engineering industries (4 experts), manufacturing (2 experts), logistics (1 expert) and IT firms and blockchain corporations (3 experts). Four of the investigated experts have a security background, while two of these are security information architects, whereby one is designing secure blockchain architectures. Another expert is responsible for security lifecycle and governance and the last one is tackling IoT security in particular. Two of the remaining experts work exclusively on blockchain technology and another works as an information architect. The last three experts are IT consultants.

The experts have a cumulative 101 years of experience, ranging from 2 to 25 years with an average value of 10.1 years and a median of 7.5 years. This experience was gained in companies of various sizes, including both SMEs (with up to 249 employees) and large enterprises (up to 500,000 employees). The average enterprise size the interviewees are familiar with is 164,583 with the median at 30,000 employees.

*Procedure* To identify the opportunities and challenges of using our blockchain-based DT data management approach and to evaluate the implemented prototype, we develop three categories of questions for the interview. These categories are based on DT **lifecycle aspects (1)**, the suitability of the **blockchain approach (2)** and the characteristics of the developed **prototype (3)**. The questions for the interview are based on relevant literature. We follow [Dietz et al. \(2019\)](#) and [Dietz and Pernul \(2020a\)](#) to identify DT lifecycle aspects **(1)**. For **(2)**, we rely on [Malakuti and Grüner \(2018\)](#) and [Rubio et al. \(2017\)](#) that provide the problem area to which our approach poses a solution. To derive the questions for category **(3)**, we derive the questions from the distinct features of our prototype (cf. [Table 1](#)).

To evaluate and gain additional practical insights on the categories **(1)**, **(2)** and **(3)**, we conduct a semi-structured expert interview according to [Lazar, Feng, and Hochheiser \(2017\)](#). The interview is structured in the following phases:

- Phase 1) Introduction. At the start, the participants are questioned about their expertise and practical experience. Subsequently, an introduction to our research problem and approach is given. Additionally, we guide each interviewee through our EtherTwin prototype. Before the experts are interviewed, we encourage them to mention any issues that emerge during the following phases.
- Phase 2) Interview. In this phase, the set of questions corresponding to the three categories are posed. Thereby, the interview questions are deliberately stated in a generic way to enable experts to share their individual experience [Lazar et al. \(2017\)](#). The questions start off with **lifecycle aspects (1)**, which represent the most generic questions, followed by requesting the experts' opinion on the underlying **blockchain approach (2)**. The last category contains the least generic questions and tackles our EtherTwin **prototype (3)**.
- Phase 3) Wrap-up. We summarize the experts' main feedback. The expert is encouraged to state additional feedback on our research and EtherTwin prototype to help validate our approach. Moreover, areas requiring revision can be identified.

The guideline to the expert interview, including the interview questions, procedure and research purpose, can be found in [Appendix B](#). Each interview participant received a copy of the guideline in advance of the interview.

*Results* We briefly describe the results of the interviews below, before discussing the experts' suggestions for improvement in [Section 7](#).

In terms of relevant **lifecycle aspects (1)**, the experts mentioned that there are additional roles at each operator that need separate permissions, for example engineers, managers, developers, analysts and security employees. Half of the experts believe that including relevant participants such as auditors and regulatory authorities could be beneficial. Moreover, 30% of the experts think it would be beneficial to include roles for public authorities, e.g. to manage the compliance to environmental law. Two experts mention that the Owner and operator may not be the same. For instance, the operator might only have leased the industrial asset, while the Owner might still be the integrator or another lifecycle participant. Moreover, some of the roles should be further distinguished between manufacturers of the components, the integrator of the components (the manufacturer of the machine) and the operator. Another helpful remark, mentioned by two experts, is that modeling sub-roles might be required.

Six out of ten experts see the data for managing an industrial asset as dependent on various aspects, including the industrial asset itself, the lifecycle phases involved as well as the use case, as also other non-industrial devices could be modeled with our approach. However, the most important data was:






Fig. 10. Expert evaluations regarding the suitability of the presented solution.

- sensor and operating data (50%)
- cumulative information & analytics including dashboards (50%)
- master data categorized into mechanical, electrical and IT-related (30%)
- the relations among the components, e.g. dependencies, network (20%)
- as well as information about hardware- and software and their modifications (10%)

Three experts reckon the sensitivity of the data as a very important aspect, esp. when critical infrastructures and functions are involved.

Information about an industrial asset is currently shared, but only in a limited way. According to the experts, Industry 4.0 sharing practices are currently still in an early stage. For example, ICSs are integrated to corporate IT systems in order to communicate with enterprise resource planning (ERP) and e-commerce systems. Moreover, according to the experts, current sharing practices involve only strategic partners. Nevertheless, the creation of greater collaboration platforms is planned – leaning towards the notion of an ecosystem of DTs. In practice, our sharing approach is thus seen as a future issue, while security is and will remain a pressing challenge.

Thereby, the most important advantages of information sharing are collaboration opportunities and product improvement (70%), followed by transparency and recording (40%). Using the blockchain, the experts expect a very low rate of failure (availability) and manipulation (integrity). The greatest disadvantage are data ownership issues and the potential loss of valuable corporate trade secrets (70%). Four experts expect the risk of industrial espionage (esp. among the supply chain), resulting in an increase of the substitutability potential through rivals and in a reduction of lock-in effects. Nevertheless, two experts emphasized that the benefits certainly outweigh the downsides.

Regarding **blockchain suitability (2)**, the answers are summarized in Fig. 10. Most experts agreed that blockchain is suitable for managing DT data along its lifecycle. Nevertheless, some experts were skeptical and preferred a TTP over a blockchain solution for its simplicity. One expert noted that this solution should not be used for machine operations since that would require millisecond latencies. Another interviewee suggested that it should be used to track machine interactions, i.e. firmware upgrades and part changes.

Finally, the user interface of the EtherTwin **prototype (3)** was received positively. All participants agreed that it was well suited to the task at hand. Adjectives used for description were *intuitive*, *clear* and *modern*. One expert argued that the developed user interface is not needed in practice, since the backend should be fully integrated with existing systems, such as condition-based maintenance systems, ERP and product data management (PDM).

When asked for estimates of practical performance requirements, the experts provided varying estimations based on their experience. While an SME with 10 manufacturing machines may create two twins of these machines per year, an automotive manufacturer may create one per produced car, or as many as 10,000 per month. Estimates for shared documents for a twin also ranged from one document per day to a few documents per year, depending on the amount of shared documentation (i.e. aircraft production requires a large number of accompanying documentation material). For sensor data, raw sensor logs can result in significant data volume and velocity (up to terabytes/day), but not all of this data requires sharing. Experts suggested that only non-nominal or aggregated sensor data needs to be shared, resulting in a volume around hundreds of entries per hour.

## 7. Discussion

Hereafter we discuss the results of the evaluation, the resulting limitations and how the experts' feedback can be used to improve the prototype in the future. We start with discussing the lifecycle aspects in Section 7.1. At last, performance (Section 7.2) and security (Section 7.3) aspects are discussed.

### 7.1. Lifecycle

**Access Control.** Additional lifecycle roles (e.g. an auditor or government authorities) could be implemented by updating our Authorization Contract. This includes the possibility for sub-roles and inheritance, for example to separate permissions for a technician and financial controller at the manufacturer. Delegation of rights could be achieved by including permission delegations in ABAC, for which several strategies have been proposed [Servos and Osborn \(2016\)](#). Another suggestion concerned the need for time or

event-based access to data by lifecycle parties. The access control model could be expanded to include an expiry time for each attribute, which is validated whenever access is requested. EtherTwin provides a starting point that can be extended and specialized to fit practical use cases individually. Another recurring suggestion made by experts was a role-specific user interface. In addition to tailoring the available roles to the practical use case, a role-specific twin overview page could help users find the needed information faster.

**Data Governance.** For collaboratively run applications, governance aspects are important to consider. Future software updates to the deployed smart contracts may be necessary to incorporate additional DT features. Code changes to deployed smart contracts are not trivial and require specific application patterns to avoid data loss. Upgradability of smart contracts can be achieved using a Contract Registry or a Data Contract pattern Xu et al. (2018). To create new twins with enhanced functionality, the existing Registry contract can be upgraded to allow for modular Specification contract templates.

**Data structure.** Additionally, a data structure might be established that is not only based on the components (cf. Fig. 3), but categorizes mechanical, electrical and IT information as well. Future work could investigate how to integrate this categorization, e.g. as an alternative structure or as sub-structure for each component.

**Additional Data.** Our prototype has few restrictions regarding data volume and variety. The additional information deemed relevant by the experts could thus be easily integrated. Additionally, simulation is a key part of DTs. EtherTwin currently supports upload of simulation data, but future work could investigate how simulations can be directly deployed in the user interface. For better usability, future work could also extend our prototype by including analytical dashboards. Experts suggested that each role should be able to get an at-a-glance overview of the asset's state. Such a dashboard could include out-of-range sensor values, recently updated documents, asset performance metrics and risk indicators.

**Asset Control.** Similarly, DTs should provide some control over the industrial asset. Firmware management and updates were suggested by experts as a potential use case for EtherTwin. Program code of PLCs could be uploaded through the user interface by permissioned users and automatically installed by the Device Agent, documenting all actions in the smart contract. This enables traceability and accountability of participants for each modification made to the physical asset.

## 7.2. Performance

**On-chain.** The twin and document creation rates estimated by the experts do not present a challenge for a prototype, as even the maximum values are within the performance limits of Ethereum and Swarm. Private Ethereum blockchains support between 50 and 100 transactions/second Dinh et al. (2017), which implies that more than 4 million twin interactions are possible per day (i.e. document creation, sensor creation).

**Off-chain.** Experts mentioned that multiple events might need to be shared per second for a specific sensor. However, Swarm is currently restricted to one update per feed per second. To deal with this restriction, sensor feeds are updated once per second with batched sensor updates from the Device Agent. Thus, no data is lost and failure data is shared in a timely fashion. This restriction precludes real-time monitoring and control of assets, as pointed out by an expert.

## 7.3. Security

**Research Question.** With our research question we aim to develop secure lifecycle information management for DTs:

**RQ1.** How can the data of Digital Twins be shared among multiple untrusted lifecycle parties while ensuring confidentiality, integrity and availability?

Our prototype provides *confidentiality* by including fine-grained access control as well as encryption. All experts agreed that access control and the concomitant encryption are essential for business adoption (cf. Fig. 10). On-chain data *integrity* is assured by the immutability of the blockchain and the full replication of data among the participating nodes. Off-chain data integrity is provided by maintaining the DHT encryption key and sensor feeds through the Device Agent controlled by the Owner. Additionally, Swarm feed updates must be signed and are append-only, so integrity of past entries is maintained. In terms of *availability*, our decentralized approach enables the participants to manage their own nodes to maintain fully replicated copies of on-chain and off-chain storage. The proposed architecture relies on three distinct systems to function properly: the blockchain network, the DHT and at least one Device Agent per organization. Due to the resulting complexity, consequences of failure should be properly considered:

- **Blockchain node failure:** If an organization's blockchain node crashes, it will be unable to access the DApp as it relies on the blockchain node as a data source. This would not affect other organizations. If  $> \frac{1}{3}$  of all blockchain nodes in the network fail, write transactions are no longer available for all participants
- **DHT node failure:** If the DHT node is unavailable, the organization will be unable to retrieve the DT specification, documents and sensor data. Other organizations are unaffected, unless they are trying to retrieve Twin data of the crashed organization for the first time

- **Device Agent failure:** A failed Device Agent implies that encryption keys will not be updated on permission changes while it is down. Thus, newly shared data for its twins will not be available to the sharing recipients. Additionally, sensor data will not be updated.

**Encryption.** Despite these already built-in security measures, sharing business data with external entities bears risks for enterprise security. While data is encrypted, loss of the encryption key or compromise of the elliptic curve/AES encryption schemes could lead to access by unauthorized parties. Since data cannot be removed from other nodes once uploaded to the DHT, there is an inevitable loss of control that comes with sharing encrypted data. Due to the distributed nature of the DHT, read access to shared data cannot be revoked, and it is not possible to determine which users actually accessed DHT data. Safeguarding the Device Agent and the encryption keys are thus paramount to data security in our approach. The prototype could be improved in this regard by hiding private key information in the user interface and using the Web Cryptography API<sup>11</sup> instead of the browser's local storage. In addition, the sharing enterprise must rely on the recipients to protect the encryption keys and data as well. Future research could also investigate future-proofing the encryption procedures by utilizing quantum-proof schemes.

**Misuse.** Another consideration mentioned by an expert is misuse potential at the time of data entry. For a decentralized application, besides signature checks there is no way of checking the authenticity of uploaded data. A malicious lifecycle participant could thus upload false information that cannot be deleted. Nevertheless, the versioning system in EtherTwin ensures that past versions of data remain available.

**Public blockchains.** If a public blockchain is used, confidentiality of on-chain metadata becomes a concern. Since on-chain data is not encrypted, participants should avoid including confidential information in metadata. If this is maintained, the contracts can be deployed on the public Ethereum blockchain and there is no need for participants to operate a blockchain infrastructure. To ensure infrastructure control and data confidentiality, both the Ethereum blockchain and the Swarm DHT can also be set up as private networks. Permissioned Ethereum networks may use a more resource-efficient byzantine-fault tolerant consensus algorithm such as IBFT 2.0 [Saltini and Hyland-Wood \(2019\)](#).

**Identity Management.** Usability could be improved by mapping Ethereum addresses to human-readable names. Organizations may associate employee identities in existing identity management systems with Ethereum key pairs to enable single sign-on. Future research could investigate how to best implement a mapping of enterprise identity to blockchain identity.

## 8. Conclusion

To conclude, the EtherTwin DApp implements the complex DT sharing requirements of the Industry 4.0 landscape without the need for a TTP. This is achieved through a fine-grained blockchain-based access control model coupled with encrypted off-chain data storage. The open source prototypical implementation is based on Ethereum and Swarm. Additionally, we evaluate our model through use case elaboration and performance testing. Interview responses by industry experts validate the prototype's practical suitability and provide avenues for future research.

For example, our work on blockchain-based information sharing and access control may be extended to other areas, i.e. health DT data sharing, data marketplaces and machine certifications. Business processes can also be interpreted as DTs [Dietz and Pernul \(2020a\)](#). Approaches for blockchain-based business process management involving physical assets could thus be integrated with blockchain-based DTs modeled in our work. Additionally, our prototype could be enhanced by enabling data flow from the twin to the industrial asset. These interactions could involve calling PLC functions through the smart contract, or installing firmware updates. Finally, simulation environments could be directly integrated in the decentralized sharing platform, instead of only sharing simulation results as documents.

## Declaration of Competing Interest

The authors declare that they do not have any financial or nonfinancial conflict of interests

## Acknowledgements

We would like to thank the interviewed experts for their time and valuable contributions. Furthermore, we would like to express our thanks to our reviewers for their helpful suggestions. Part of this work was performed under the ZIM SisseC project<sup>12</sup>, which is supported under contract by the German Federal Ministry for Economic Affairs and Energy (16KN085725).

<sup>11</sup> <https://www.w3.org/TR/WebCryptoAPI/>

<sup>12</sup> <https://www.it-logistik-bayern.de/produktionslogistik/projekt-sissec>

Appendix A. Screenshots of the prototype

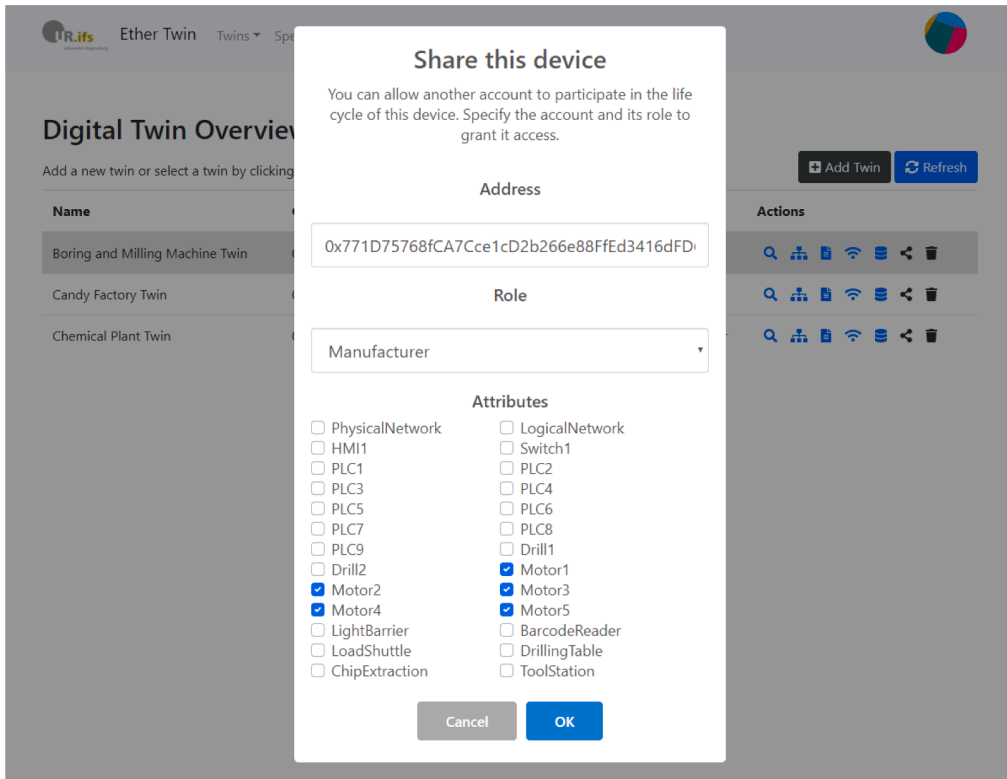


Fig. A1. Screenshot of the "share twin"-functionality.

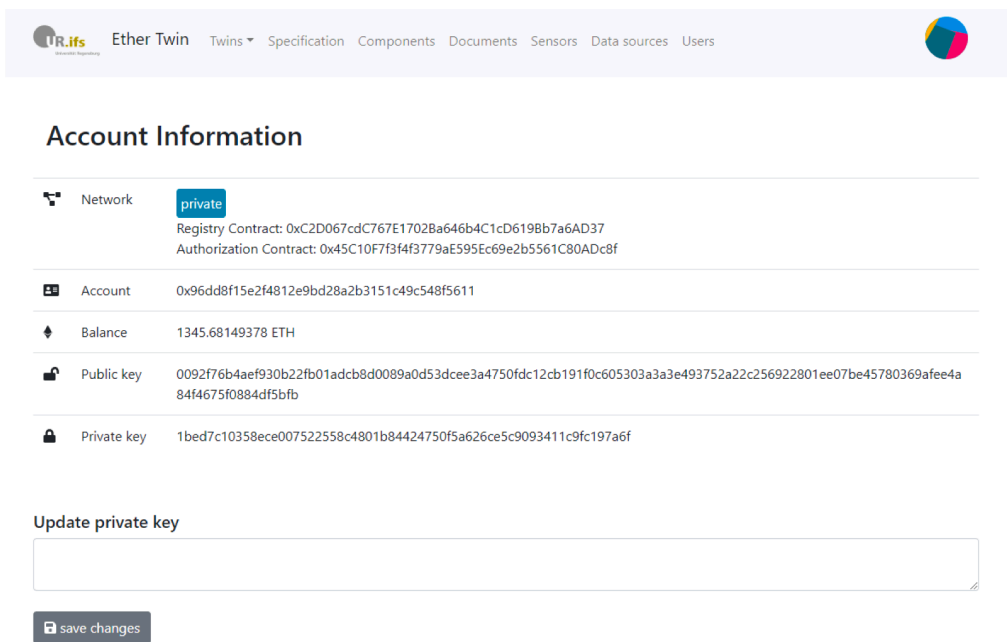
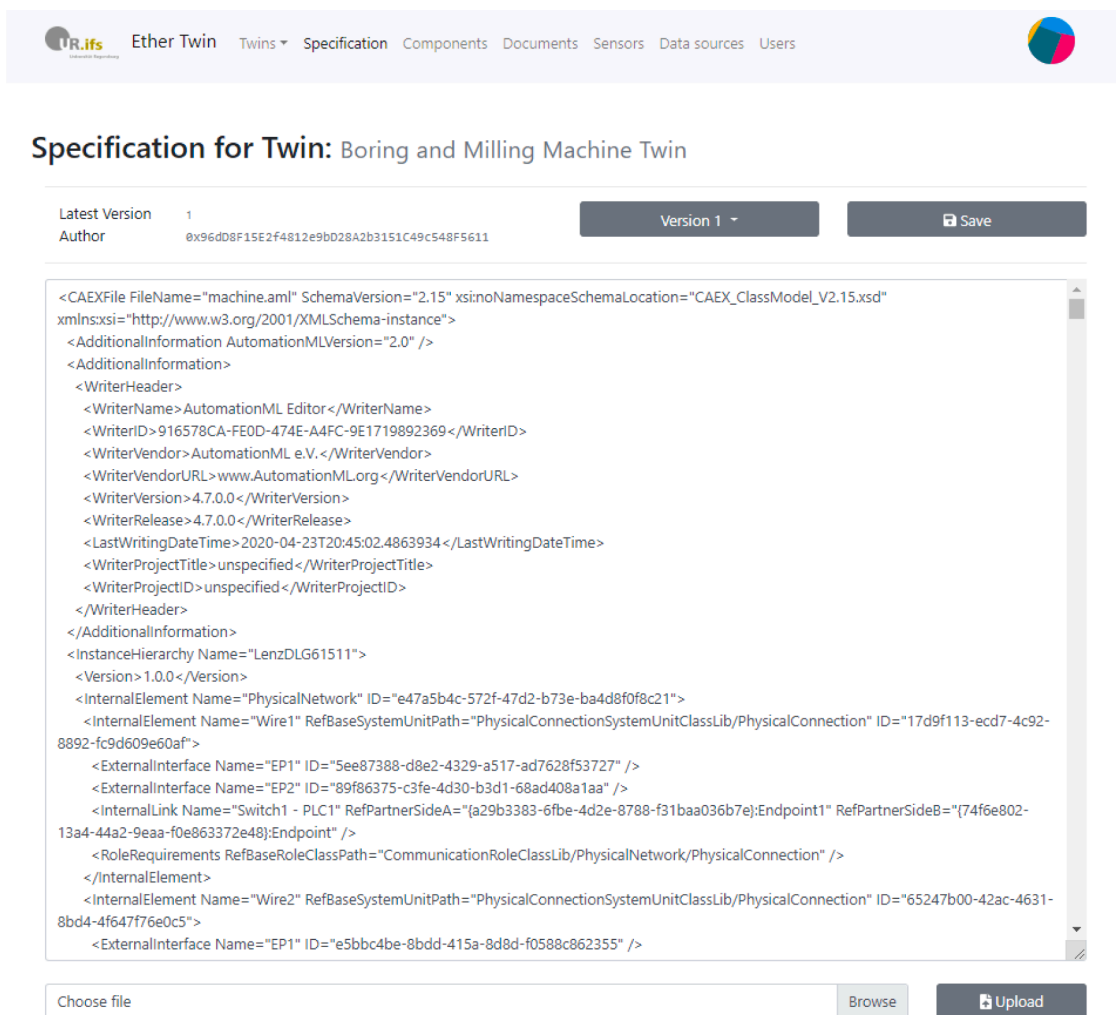


Fig. A2. Screenshot of the user's account page.



The screenshot displays the 'Ether Twin' web application interface. At the top, there is a navigation menu with links for 'Twins', 'Specification', 'Components', 'Documents', 'Sensors', 'Data sources', and 'Users'. The main heading is 'Specification for Twin: Boring and Milling Machine Twin'. Below the heading, there are controls for the 'Latest Version' (1) and 'Author' (0x96d08f15e2f4812e9bd028a2b3151c49c548f5611). A 'Version 1' dropdown and a 'Save' button are also present. The central part of the interface is a large text area containing XML Schema (AML) code for a machine model. The code includes metadata like 'AutomationML Editor' and 'AutomationML e.V.', and defines elements for a physical network with two external interfaces (EP1 and EP2) and an internal link (Switch1 - PLC1). At the bottom, there are 'Choose file', 'Browse', and 'Upload' buttons.

```
<CAEXFile FileName="machine.aml" SchemaVersion="2.15" xmlns: xsi:noNamespaceSchemaLocation="CAEX_ClassModel_V2.15.xsd"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
<AdditionalInformation AutomationMLVersion="2.0" />
<AdditionalInformation>
<WriterHeader>
<WriterName>AutomationML Editor</WriterName>
<WriterID>916578CA-FE0D-474E-A4FC-9E1719892369</WriterID>
<WriterVendor>AutomationML e.V.</WriterVendor>
<WriterVendorURL>www.AutomationML.org</WriterVendorURL>
<WriterVersion>4.7.0.0</WriterVersion>
<WriterRelease>4.7.0.0</WriterRelease>
<LastWritingDateTime>2020-04-23T20:45:02.4863934</LastWritingDateTime>
<WriterProjectTitle>unspecified</WriterProjectTitle>
<WriterProjectID>unspecified</WriterProjectID>
</WriterHeader>
</AdditionalInformation>
<InstanceHierarchy Name="LenzDLG61511">
<Version>1.0.0</Version>
<InternalElement Name="PhysicalNetwork" ID="e47a5b4c-572f-47d2-b73e-ba4d8f0f8c21">
<InternalElement Name="Wire1" RefBaseSystemUnitPath="PhysicalConnectionSystemUnitClassLib/PhysicalConnection" ID="17d9f113-ecd7-4c92-8892-fc9d609e60af">
<ExternalInterface Name="EP1" ID="5ee87388-d8e2-4329-a517-ad7628f53727" />
<ExternalInterface Name="EP2" ID="89f86375-c3fe-4d30-b3d1-68ad408a1aa" />
<InternalLink Name="Switch1 - PLC1" RefPartnerSideA="{a29b3383-6fbe-4d2e-8788-f31baa036b7e};Endpoint1" RefPartnerSideB="{74f6e802-13a4-44a2-9eaa-f0e863372e48};Endpoint" />
<RoleRequirements RefBaseRoleClassPath="CommunicationRoleClassLib/PhysicalNetwork/PhysicalConnection" />
</InternalElement>
<InternalElement Name="Wire2" RefBaseSystemUnitPath="PhysicalConnectionSystemUnitClassLib/PhysicalConnection" ID="65247b00-42ac-4631-8bd4-4f647f76e0c5">
<ExternalInterface Name="EP1" ID="e5bbc4be-8bdd-415a-8d8d-f0588c862355" />
```

Fig. A3. Screenshot of the AML-structured specification of the asset.

**Data:** A set of twins  $T = (t_1, \dots, t_n)$  with mappings for roles  $M_t^{UR}$ , permissions  $M_t^{PR}$ , attributes  $M_t^{UA}$  and a set of components  $C_t$ .  
**Result:** A set of encrypted file keys  $fk_{tcu} \forall t \in T, c \in C, u \in U_t$  used to decrypt data  $D_{tcn} \forall n \in 1..N$  and uploaded to DHT feeds owned by the Device Agent.

```

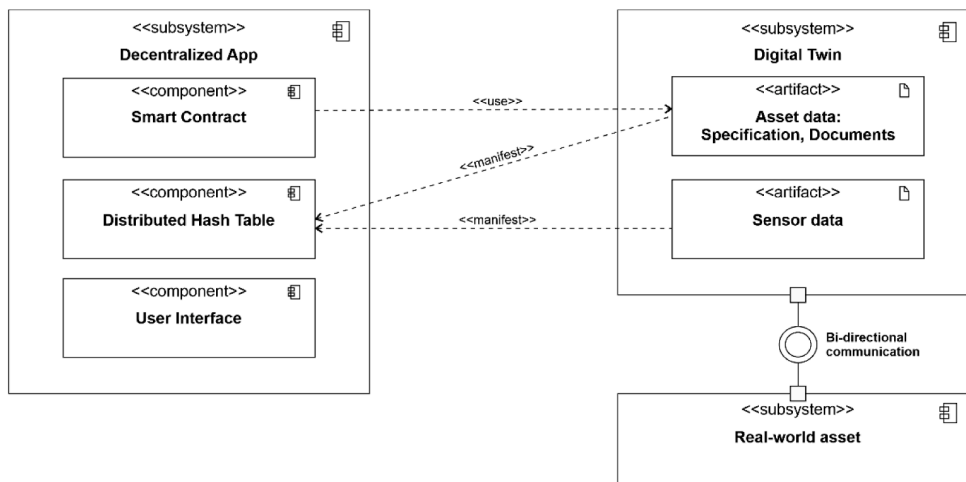
1: function CREATEFILEKEYS(t)
2:    $C_t \leftarrow \text{getComponents}(t)$                                 > retrieve permissions from smart contract
3:    $M_t^{UR} \leftarrow \text{getRoleAssignment}(t)$ 
4:    $M_t^{PR} \leftarrow \text{getPermissionAssignment}(t)$ 
5:    $M_t^{UA} \leftarrow \text{getAttributeAssignment}(t)$ 
6:   for each  $c \in C_t$  do                                       > generate two symmetric keys sk per component
7:      $(sk_{tc}^{doc}, sk_{tc}^{sensor}) = \mathcal{G}_{enc}$ 
8:   end for each
9:   for each  $u \in (M_t^{UR} \cap M_t^{UA})$  do                             > prepare file keys fk for users
10:     $pk_u \leftarrow (\text{DHT}) \text{getUserPublicKey}(u)$ 
11:    for each  $c \in C_t$  do
12:       $r \leftarrow M_t^{UR}$ 
13:      for each  $d \in \{doc, sensor\}$  do
14:        if  $(c \equiv a \mid a \in M_t^{UA}) \wedge (p_{read}^d \in M_t^{PR})$  then
15:           $fk_{tcu}^d = \text{ECIES\_enc}(pk_u, sk_{tc}^d)$ 
16:           $(\text{DHT}) \text{updateFeed}(c, fk_{tcu}^d)$ 
17:        end if
18:      end for each
19:    end for each
20:  end for each
21: end function

```

Algorithm 1. Create off-chain file keys based on read permissions.

## Appendix B. Guideline of the expert interview


The present research focuses on blockchain-based Digital Twin information management (“*EtherTwin: Blockchain-based Digital Twin Information Management*”). It aims at developing an approach to securely share and store data about an Industry 4.0 asset (Digital Twin data) along its entire lifecycle. As the involved lifecycle parties that share data usually pursue different goals, trust issues hamper data sharing in practice. To resolve these trust issues and to strengthen security, blockchain technology is harnessed in the form of a Decentralized Application (DApp).



**Figure 1:** The concept of DT data sharing and storage realized with a Decentralized App (DApp), with its basis on the blockchain technology.

The *Digital Twin* describes a concept to virtually represent an Industry 4.0 asset (e.g. a conveyor belt). Thereto, a central collection of data of the industrial asset is of utmost importance. In **Figure 1**, the core data artifacts of a *Digital Twin* as well as its connection to the *Real-world asset* the *Digital Twin* represents are shown on the right-hand side. Moreover, the connectors show where each data artifact of the *Digital Twin* is represented in the *DApp*-architecture. First, the information of the asset’s *Specification* is used to create a *Smart Contract*. For instance, this includes data about the asset’s composition (*Specification* data). Asset data like *Documents* and *Sensor data* are stored in the *Distributed Hash Table* off-chain and all interactions with these data (upload, delete, update etc.) are linked in the blockchain. A *User Interface* builds on top and enables a user-friendly interaction from all lifecycle parties, including less technically savvy users.

The goal of this research is a full-featured implementation of Digital Twin data management, the EtherTwin prototype, and its evaluation. To enable secure data sharing, our approach ensures that integrity, confidentiality and availability of data are maintained.



**Figure 2:** The selected technologies for the EtherTwin prototype.

**Figure 2** shows the technologies used to build the EtherTwin prototype. The asset's *Specification* is based on the industry format [AutomationML](#). With this information, the components (industrial control systems) of the asset can be derived and a Digital Twin management space is built. The blockchain part is realized with Ethereum. An off-chain *Distributed Hash Table* is integrated. The Device Agent manages the connection to the sensors etc. of the industrial assets and directly incorporates sensor data and system log data.

Next to implementing the approach in form of the EtherTwin prototype, we conducted performance tests and proposed a use case from practice that shows how our prototype can be used. In addition, expert interviews are to be conducted, where individual questions are examined in greater depth to gather the experience and position of the experts.



The following questions are based on the three categories corresponding to our research: (1) Lifecycle Data Sharing, (2) Blockchain Suitability and (3) Prototype.

### **(1) Lifecycle Data Sharing**

- What roles and which participants can be found involved in an industrial assets' lifecycle?
- What data is relevant to manage a Digital Twin of an industrial asset?
- Which participants share (respectively require) which type of data?
- Is data currently shared at all about assets?
- What would be potential benefits for the firm that could be realized through sharing data about assets?

### **(2) Blockchain Suitability**

- Is the solution suitable to share data over an industrial asset's lifecycle?
- To what extent do the following aspects benefit the adoption of the solution:
  - Access control,
  - Encryption and
  - Decentralization?

### **(3) Prototype**

- Is the user interface satisfactory? How could the usability be improved?
- What are the current practical requirements regarding throughput and latency:
  - How many Digital Twins are created per month?
  - How many asset-related documents are shared per day?
  - How many sensor data of an industrial asset occur per second?
- Would a self-hosted solution be preferred over a public solution with transaction costs?


## **References**

- Angrish, A., Craver, B., Hasan, M., & Starly, B. (2018). A case study for blockchain in manufacturing: æfabrecg: A prototype for peer-to-peer network of manufacturing nodes. *Procedia Manufacturing*, 26, 1180–1192. <https://doi.org/10.1016/j.promfg.2018.07.154>. 46th SME North American Manufacturing Research Conference, NAMRC 46, Texas, USA
- Ayman, A., Aziz, A., Alipour, A., & Laszka, A. (2019). Smart contract development in practice: Trends, issues, and discussions on stack overflow. *CoRR*.
- Badsha, S., Vakiliinia, L., & Sengupta, S. (2020). BloCyNfo-Share: Blockchain based Cybersecurity Information Sharing with Fine Grained Access Control. *10th annual computing and communication workshop and conference, {ccwc} 2020, las vegas, nv, usa, january 6–8, 2020* (pp. 317–323). IEEE. <https://doi.org/10.1109/CCWC47524.2020.9031164>.
- Baig, F., & Wang, F. (2019). Blockchain enabled distributed data management - A vision. *35th IEEE international conference on data engineering workshops, ICDE workshops 2019, macao, china, april 8–12, 2019* (pp. 28–30). IEEE. <https://doi.org/10.1109/ICDEW.2019.00-39>.
- Baumgart, I., & Mies, S. (2007). S/Kademlia: A practicable approach towards secure key-based routing. *International conference on parallel and distributed systems* (pp. 1–8). <https://doi.org/10.1109/ICPADS.2007.4447808>.
- Berdik, D., Otoum, S., Schmidt, N., Porter, D., & Jararweh, Y. (2021). A survey on blockchain for information systems management and security. *Information Processing and Management*, 58(1), 102397. <https://doi.org/10.1016/j.ipm.2020.102397>.
- Boschert, S., Heinrich, C., & Rosen, R. (2018). Next Generation Digital Twin. *Proceedings of tmce* (pp. 209–217).
- Coyne, E. J., & Weil, T. R. (2013). ABAC And RBAC: Scalable, flexible, and auditable access management. *IT professional*, 15(3), 14–16. <https://doi.org/10.1109/MITP.2013.37>.
- Cruz, J. P., Kaji, Y., & Yanai, N. (2018). RBAC-SC: Role-based access control using smart contract. *IEEE Access*, PP, 1. <https://doi.org/10.1109/ACCESS.2018.2812844>.
- Di Francesco Maesa, D., Mori, P., & Ricci, L. (2019). A blockchain based approach for the definition of auditable access control systems. *Computers & Security*, 84, 93–119. <https://doi.org/10.1016/j.cose.2019.03.016>.

- Dietz, M., & Pernul, G. (2020a). Digital twin: Empowering enterprises towards a system-of-Systems approach. *Business & Information Systems Engineering*, 62(2), 179–184.
- Dietz, M., & Pernul, G. (2020b). Unleashing the digital Twin's potentials for ICS security. *IEEE Security & Privacy*.
- Dietz, M., Putz, B., & Pernul, G. (2019). A distributed ledger approach to digital twin secure data sharing. In S. N. Foley (Ed.), *Data and applications security and privacy xxxiii* (pp. 281–300). Springer International Publishing. [https://doi.org/10.1007/978-3-030-22479-0\\_15](https://doi.org/10.1007/978-3-030-22479-0_15).
- Dinh, T. T. A., Wang, J., Chen, G., Liu, R., Ooi, B. C., & Tan, K.-L. (2017). BLOCKBENCH: A Framework for Analyzing Private Blockchains. In *SIGMOD'17 Proceedings of the 2017 acm international conference on management of data* (pp. 1085–1100). New York, NY, USA: ACM. <https://doi.org/10.1145/3035918.3064033>.
- Eckhart, M., & Ekelhart, A. (2018). Towards Security-Aware Virtual Environments for Digital Twins. *Proceedings of the 4th acm workshop on cyber-physical system security* (pp. 61–72). <https://doi.org/10.1145/3198458.3198464>.
- Eposito, C., Tamburis, O., Su, X., & Choi, C. (2020). Robust decentralised trust management for the internet of things by using game theory. *Information Processing and Management*. <https://doi.org/10.1016/j.ipm.2020.102308>.
- Ethereum Swarm Contributors (2019). Swarm Documentation. <https://swarm-guide.readthedocs.io/en/latest>.
- Fette, I., & Melnikov, A. (2011). RFC6455 - The WebSocket protocol. *IETF Standards Track*. <https://doi.org/10.17487/RFC6455>.
- Haarmann, S., Batoulis, K., Nikaj, A., & Weske, M. (2018). DMN Decision Execution on the Ethereum Blockchain. In J. Krogstie, & H. A. Reijers (Eds.), *Caise 2018* (pp. 327–341). Cham: Springer International Publishing.
- Hasan, H. R., Salah, K., Jayaraman, R., Omar, M., Yaqoob, I., Pesic, S., ... Boscovic, D. (2020). A blockchain-Based approach for the creation of digital twins. *IEEE Access*.
- Huang, S., Wang, G., Yan, Y., & Fang, X. (2020). Blockchain-based data management for digital twin of product. *Journal of Manufacturing Systems*.
- Kaur, M. J., Mishra, V. P., & Maheshwari, P. (2020). The convergence of digital twin, IoT, and machine learning: Transforming data into action. In M. Farsi, A. Daneshkhan, A. Hosseinian-Fa, & H. Jahankhani (Eds.), *Digital twin technologies and smart cities* (pp. 3–17). Springer International Publishing.
- Khan, M. A., & Salah, K. (2018). IoT Security: Review, blockchain solutions, and open challenges. *Future Generation Computer Systems*, 82, 395–411.
- Kuhn, D. R., Coyne, E. J., & Weil, T. R. (2010). Adding attributes to role-Based access control. *Computer*, 43(6), 79–81. <https://doi.org/10.1109/MC.2010.155>.
- Lazar, J., Feng, J., & Hochheiser, H. (2017). *Research methods in human-Computer interaction, 2nd edition*. Morgan Kaufmann.
- Li, J., Wu, J., Jiang, G., & Srikanthan, T. (2020). Blockchain-based public auditing for big data in cloud storage. *Information Processing and Management*, 57(6), 102382. <https://doi.org/10.1016/j.ipm.2020.102382>.
- López-Pintado, O., Dumas, M., García-Bañuelos, L., & Weber, I. (2019). Dynamic Role Binding in Blockchain-Based Collaborative Business Processes. In P. Giorgini, & B. Weber (Eds.), *Caise 2019* (pp. 399–414). Cham: Springer International Publishing.
- Malakuti, S., & Grüner, S. (2018). Architectural aspects of digital twins in IIoT systems. *Proceedings of the 12th European Conference on Software Architecture Companion Proceedings - ECSA '18*, 1–2.
- Mandolla, C., Petruzzelli, A. M., Percoco, G., & Urbinati, A. (2019). Building a digital twin for additive manufacturing through the exploitation of blockchain: A case analysis of the aircraft industry. *Computers in Industry*, 109, 134–152. <https://doi.org/10.1016/j.compind.2019.04.011>.
- Meroni, G., & Plebani, P. (2018). Combining Artifact-Driven Monitoring with Blockchain: Analysis and Solutions. *Caise 2018* (pp. 103–114). Cham: Springer International Publishing.
- Pedersen, A. B., Risius, M., & Beck, R. (2019). A ten-step decision path to determine when to use blockchain technologies. *MIS Quarterly Executive*, 18(2), 99–115. <https://doi.org/10.17705/2msqe.00010>.
- Ríos, J., Hernández, J. C., Oliva, M., & Mas, F. (2015). Product avatar as digital counterpart of a physical individual product: Literature review and implications in an aircraft. In R. Curran, N. Wognum, M. Borsato, J. Stjepandic, & W. J. C. Verhagen (Eds.), *Advances in Transdisciplinary Engineering: 2. Transdisciplinary lifecycle analysis of systems - proceedings of the 22nd ISPE inc. international conference on concurrent engineering, delft, the netherlands, july 20–23, 2015* (pp. 657–666). IOS Press. <https://doi.org/10.3233/978-1-61499-544-9-657>.
- Rouhani, S., Belchior, R., Cruz, R. S., & Deters, R. (2020). Distributed attribute-based access control system using a permissioned blockchain. *CoRR, abs/2006.04384*.
- Rubio, J. E., Roman, R., & López, J. (2017). Analysis of cybersecurity threats in industry 4.0: The case of intrusion detection. In G. D'Agostino, & A. Scala (Eds.), *Lecture Notes in Computer Science: 10707. Critical information infrastructures security - 12th international conference, CRITIS 2017, lucca, italy, october 8–13, 2017, revised selected papers* (pp. 119–130). Springer. [https://doi.org/10.1007/978-3-319-99843-5\\_11](https://doi.org/10.1007/978-3-319-99843-5_11).
- Saltini, R., & Hyland-Wood, D. (2019). IBFT 2.0: A Safe And live variation of the IBFT blockchain consensus protocol for eventually synchronous networks. *CoRR, abs/1909.1*.
- Schroeder, G. N., Steinmetz, C., Pereira, C. E., & Espindola, D. B. (2016). Digital twin data modeling with automationml and a communication methodology for data exchange. *IFAC-PapersOnLine*, 49(30), 12–17.
- Servos, D., & Osborn, S. L. (2016). Strategies for incorporating delegation into attribute-based access control (ABAC). In F. Cuppens, L. Wang, N. Cuppens-Boulahia, N. Tawbi, & J. García-Alfaro (Eds.), *Lecture Notes in Computer Science: 10128. Foundations and practice of security - 9th international symposium, FPS 2016, québec city, qc, canada, october 24–25, 2016, revised selected papers* (pp. 320–328). Springer. [https://doi.org/10.1007/978-3-319-51966-1\\_21](https://doi.org/10.1007/978-3-319-51966-1_21).
- Tikhomirov, S., Voskresenskaya, E., Ivanitskiy, I., Takhaviev, R., Marchenko, E., & Alexandrov, Y. (2018). SmartCheck: Static Analysis of Ethereum Smart Contracts. *2018 IEEE/ACM 1st international workshop on emerging trends in software engineering for blockchain (wetseb)* (pp. 9–16).
- Uhlemann, T. H.-J., Lehmann, C., & Steinhilper, R. (2017). The digital twin: Realizing the cyber-physical production system for industry 4.0. *Procedia CIRP*, 61, 335–340. <https://doi.org/10.1016/j.procir.2016.11.152>. The 24th CIRP Conference on Life Cycle Engineering
- Venable, J. R., Pries-Heje, J., & Baskerville, R. L. (2012). A comprehensive framework for evaluation in design science research. In K. Peffers, M. A. Rothenberger, & W. L. K. Jr. (Eds.), *Lecture Notes in Computer Science: 7286. Design science research in information systems. advances in theory and practice - 7th international conference, DESRIST 2012, las vegas, nv, usa, may 14–15, 2012. proceedings* (pp. 423–438). Springer. [https://doi.org/10.1007/978-3-642-29863-9\\_31](https://doi.org/10.1007/978-3-642-29863-9_31).
- Xu, X., Pautasso, C., Zhu, L., Lu, Q., & Weber, I. (2018). A Pattern Collection for Blockchain-based Applications. *Proceedings of the 23rd european conference on pattern languages of programs* (pp. 3:1–3:20). ACM.
- Xu, X., Weber, I., & Staples, M. (2019). *Architecture for blockchain applications*. Springer. <https://doi.org/10.1007/978-3-030-03035-3>.
- Zhang, Y., Kasahara, S., Shen, Y., Jiang, X., & Wan, J. (2019). Smart contract-based access control for the internet of things. *IEEE Internet of Things Journal*. <https://doi.org/10.1109/JIOT.2018.2847705>.
- Zhao, Q., Chen, S., Liu, Z., Baker, T., & Zhang, Y. (2020). Blockchain-based privacy-preserving remote data integrity checking scheme for IoT information systems. *Information Processing and Management*. <https://doi.org/10.1016/j.ipm.2020.102355>.

# Appendix A

## Literature



**Figure 16:** Overview: publications on digital twin (security) over the last years

Year	ScienceDirect	IEEE	ACM	AIS	Springer
2015	1	0	0	0	0
2016	1	1	0	0	1
2017	13	8	1	1	2
2018	21	29	3	4	12
2019	80	84	11	6	34
2020	150	152	16	15	102
2021	287	298	10	33	131
	553	572	41	59	282

**Table 6:** Results on digital twin research\*

\* "digital twin" in title; years: 2000-2020

Year	ScienceDirect	IEEE	ACM	AIS	Springer
2018	0	0	0	0	1
2019	2	5	0	1	1
2020	2	9	1	2	3
2021	6	21	1	3	11
	10	35	2	6	16

**Table 7:** Results on digital twin security research\*\*

\*\* "digital twin" in title + "security" in abstract; years: 2000-2020

## Appendix B

# Curriculum vitae

### Marietheres Dietz

\* Sept 01, 1993    ✉ marietheres.dietz@ur.de

#### Employment History





- 11/2017 – present    📌 **Research assistant.** Chair of Information Systems I (Prof. Dr. Günther Pernul), Faculty of Business, Economics, and Management Information Systems, University of Regensburg.
- 08/2016 – 10/2016    📌 **Graduate student research assistant.** b2b e-commerce, ibi research, Regensburg.
- 04/2016 – 03/2017    📌 **Graduate student research assistant.** Chair of Information Systems I (Prof. Dr. Günther Pernul), Faculty of Business, Economics, and Management Information Systems, University of Regensburg.
- 03/2015 – 04/2015    📌 **Internship e-commerce.** ibi research, Regensburg.
- 10/2013 – 09/2015    📌 **Student research assistant.** Chair of Information Systems – Quality Management and Assurance (Prof. Dr. Matthias Klier), Faculty of Business, Economics, and Management Information Systems, University of Regensburg.
- 08/2013 – 10/2013    📌 **Internship software development.** Wiso GmbH, Ingolstadt.
- 05/2013 – 09/2015    📌 **Student research assistant.** Department of Information Systems (Dr. Norbert Meckl), Faculty of Business, Economics, and Management Information Systems, University of Regensburg.

#### Education




- 11/2017 – present    📌 **Ph.D., University of Regensburg, Germany** in Wirtschaftsinformatik.  
Thesis title: *A two-fold Perspective on Enterprise Security in the Digital Twin Context.*
- 10/2015 – 09/2017    📌 **M.Sc. with Honors, University of Regensburg, Germany** in Wirtschaftsinformatik.  
Thesis title: *Prototypische Umsetzung einer Applikation zur Detektion von Anomalien in Datenströmen.*
- 08/2015 – 01/2016    📌 **M.Sc., Linnaeus University, Sweden** in Wirtschaftsinformatik. Semester abroad.
- 10/2012 – 07/2015    📌 **B.Sc., University of Regensburg, Germany** in Wirtschaftsinformatik.  
Thesis title: *B2B-E-Commerce: Eine Analyse des Online-Ein- und Verkaufsverhaltens auf Basis einer Expertenbefragung.*

## Academic Experience






### Project involvement

- 2022     **DEVISE** *Data Quality Management for Improving Information Security*, Federal Ministry of Education and Research.
- 2021     **INSIST** *Industrial IoT Security Operations Center*, Bavarian Ministry of Economic Affairs, Regional Development, and Energy (StMWi).
- 2017-2021     **SISseC** *Secure Industrial Semantic Sensor Cloud*, German Federal Ministry for Economic Affairs and Energy (BMWi).
- 2017     **DINGfest** *Detection, Visualization, and Forensic Analysis of Security Incidents*, Federal Ministry of Education and Research (BMBF).


### Service to the research community

- 2021 - present     **Invited reviewer.** HICCS 2021. ECIS 2021. IEEE S&P 2022.
- 2020 - present     **External journal reviewer.** IEEE TII 2020. Information Systems 2022.
- 2018 - present     **External conference reviewer.** ARES 2018. CAiSE 2018. ESORICS 2018. ISPEC 2018. SECURE 2018. ACM UMAP 2019. CAiSE 2019. DBSec 2019. ESORICS 2019. ICISSP 2019. SECURE 2019. ESORICS 2020. ICCCN 2020. ICISSP 2020. IEEE DSC 2020. ARES 2021. CC Grid 2021. CAiSE 2022. COMPSAC 2022.

### Teaching





- 2018 – present     **Part-lecturer.** "Security of data-intensive Applications" (M.Sc.)
- 2017 – 2020     **Tutor and module responsible.** "Informationssysteme: Entwicklungen und Trends" (M.Sc.)
- 2016 – 2017     **Student tutor.** "IT-Security I" (B.Sc.) and "Internettechnologien" (B.Sc.)
- 2013 – 2015     **Student tutor.** "Management und Methoden der Softwareentwicklung" (B.Sc.)
-  **Student tutor.** "Grundlagen der Wirtschaftsinformatik" (B.Sc.)

### Other

- 2017 – present     **Study program coordinator.** Honors program in Business, Economics, and Management Information Systems, University of Regensburg.

## Research Publications

### Journal Articles

- 1    Böhm, F., **Dietz, M.**, Preindl, T., & Pernul, G. (2021). Augmented Reality and the Digital Twin: State-of-the-Art and Perspectives for Cybersecurity. *Journal of Cybersecurity and Privacy*, 1(3), 519–538.   
  doi:10.3390/jcp1030026
- 2    Putz, B., **Dietz, M.**, Empl, P., & Pernul, G. (2021). EtherTwin: Blockchain-based Secure Digital Twin Information Management. *Information Processing & Management*, 58(1), 102425.   
  doi:https://doi.org/10.1016/j.ipm.2020.102425
- 3    **Dietz, M.**, & Pernul, G. (2020a). Digital Twin: Empowering Enterprises Towards a System-of-Systems Approach. *Business & Information Systems Engineering*, 62(2), 179–184.   
  doi:10.1007/s12599-019-00624-0
- 4    **Dietz, M.**, & Pernul, G. (2020b). Unleashing the Digital Twin's Potential for ICS Security. *IEEE Security & Privacy*, 18(4), 20–27.   
  doi:10.1109/MSEC.2019.2961650

### Conference Proceedings

- 1 **Dietz, M.**, Hagemann, L., von Hornung, C., & Pernul, G. (2022). Employing Digital Twins for Security-by-Design System Testing. In *Proceedings of the 2022 ACM Workshop on Secure and Trustworthy Cyber-Physical Systems*. doi:10.1145/3510547.3517929
- 2 **Dietz, M.**, Schlette, D., & Pernul, G. (2022). Harnessing Digital Twin Security Simulations for systematic Cyber Threat Intelligence. In *2022 IEEE 46th Annual Computers, Software, and Applications Conference (COMPSAC)*.
- 3 **Dietz, M.**, Englbrecht, L., & Pernul, G. (2021). Enhancing Industrial Control System Forensics Using Replication-based Digital Twins. In G. Peterson & S. Sheno (Eds.), *Advances in Digital Forensics XVII* (pp. 21–38). doi:10.1007/978-3-030-88381-2\_2
- 4 Vielberth, M., Glas, M., **Dietz, M.**, Karagiannis, S., Magkos, E., & Pernul, G. (2021). A Digital Twin-Based Cyber Range for SOC Analysts. In K. Barker & K. Ghazinour (Eds.), *Data and Applications Security and Privacy XXXV* (pp. 293–311). doi:10.1007/978-3-030-81242-3\_17
- 5 **Dietz, M.**, Vielberth, M., & Pernul, G. (2020). Integrating Digital Twin Security Simulations in the Security Operations Center. In *Proceedings of the 15th International Conference on Availability, Reliability and Security*. doi:10.1145/3407023.3407039
- 6 **Dietz, M.**, Putz, B., & Pernul, G. (2019). A Distributed Ledger Approach to Digital Twin Secure Data Sharing. In S. N. Foley (Ed.), *Data and Applications Security and Privacy XXXIII* (pp. 281–300). doi:10.1007/978-3-030-22479-0\_15
- 7 **Dietz, M.**, & Pernul, G. (2018). Big Log Data Stream Processing: Adapting an Anomaly Detection Technique. In S. Hartmann, H. Ma, A. Hameurlain, G. Pernul, & R. R. Wagner (Eds.), *Database and Expert Systems Applications* (pp. 159–166). doi:10.1007/978-3-319-98812-2\_12
- 8 **Dietz, M.**, Klier, J., Klier, M., & Wiesneth, K. (2016). Bürgerzufriedenheit durch E-Government? – Eine Analyse auf Basis des Kano-Modells. In D. Stelzer, V. Nissen, & S. Straßburger (Eds.), *Multikonferenz Wirtschaftsinformatik, MKWI 2016, Illmenau, Deutschland, 23.-25.2.2016, Proceedings* (pp. 505–516).

# Bibliography

- [1] ALMEAIBED, S., AL-RUBAYE, S., TSOURDOS, A., AND AVDELIDIS, N. P. Digital twin analysis to promote safety and security in autonomous vehicles. *IEEE Communications Standards Magazine* 5, 1 (2021), 40–46.
- [2] ATALAY, M., AND ANGIN, P. A digital twins approach to smart grid security testing and standardization. In *2020 IEEE International Workshop on Metrology for Industry 4.0 IoT* (2020), pp. 435–440.
- [3] AZZAOU, A. E., KIM, T. W., LOIA, V., AND PARK, J. H. Blockchain-based secure digital twin framework for smart healthy city. *Advanced Multimedia and Ubiquitous Engineering* 716 (2020), 107 – 113.
- [4] BADIA, R. M., BAKER, M., COUGHLIN, T., FARABOSCHI, P., FRACHTENBERG, E., KAABUNGA, V., KASAHARA, H., KEETON, K., LANGE, D., LAPLANTE, P., MATWYSHYN, A., MENDELSON, A., METRA, C., MILOJICIC, D., PATEL, N., SARACCO, R., TUBB, M., AND VIANA, I. P. Technology Predictions 2022. <https://ieeecs-media.computer.org/media/tech-news/tech-predictions-report-2022.pdf>, 2022. [Online; accessed 09-Feb-2022].
- [5] BITTON, R., GLUCK, T., STAN, O., INOKUCHI, M., OHTA, Y., YAMADA, Y., YAGYU, T., ELOVICI, Y., AND SHABTAI, A. Deriving a cost-effective digital twin of an ics to facilitate security evaluation. In *Computer Security* (Cham, 2018), J. Lopez, J. Zhou, and M. Soriano, Eds., Springer International Publishing, pp. 533–554.
- [6] BOSCHERT, S., HEINRICH, C., AND ROSEN, R. Next Generation Digital Twin. In *Proceedings of the 12th International Symposium on Tools and Methods of Competitive Engineering* (2018), TMCE 2018, pp. 209–217.
- [7] BÉCUE, A., FOURASTIER, Y., PRAÇA, I., SAVARIT, A., BARON, C., GRADUSOFS, B., POUILLE, E., AND THOMAS, C. Cyberfactory1 — securing the industry 4.0 with cyber-ranges and digital twins. In *2018 14th IEEE International Workshop on Factory Communication Systems (WFCS)* (2018), pp. 1–4.
- [8] CHEN, M., SHAO, J., GUO, S., SU, L., AND DU, H. Convoy\_dtn: A security interaction engine design for digital twin network. In *2021 IEEE Globecom Workshops (GC Wkshps)* (2021), pp. 1–5.

- [9] DANILCZYK, W., SUN, Y., AND HE, H. Angel: An intelligent digital twin framework for microgrid security. In *2019 North American Power Symposium (NAPS) (2019)*, pp. 1–6.
- [10] DANILCZYK, W., SUN, Y. L., AND HE, H. Blockchain checksum for establishing secure communications for digital twin technology. In *2021 North American Power Symposium (NAPS) (2021)*, pp. 1–6.
- [11] DANILCZYK, W., SUN, Y. L., AND HE, H. Smart grid anomaly detection using a deep learning digital twin. In *2020 52nd North American Power Symposium (NAPS) (2021)*, pp. 1–6.
- [12] DIETZ, M., ENGLBRECHT, L., AND PERNUL, G. Enhancing industrial control system forensics using replication-based digital twins. In *Advances in Digital Forensics XVII* (Cham, 2021), G. Peterson and S. Sheno, Eds., Springer International Publishing, pp. 21–38.
- [13] DIETZ, M., AND PERNUL, G. Big log data stream processing: Adapting an anomaly detection technique. In *Database and Expert Systems Applications* (Cham, 2018), S. Hartmann, H. Ma, A. Hameurlain, G. Pernul, and R. R. Wagner, Eds., Springer International Publishing, pp. 159–166.
- [14] DIETZ, M., AND PERNUL, G. Digital Twin: Empowering Enterprises Towards a System-of-Systems Approach. *Business & Information Systems Engineering* 62, 2 (2020), 179–184.
- [15] DIETZ, M., AND PERNUL, G. Unleashing the digital twin’s potential for ics security. *IEEE Security Privacy* 18, 4 (2020), 20–27.
- [16] DIETZ, M., PUTZ, B., AND PERNUL, G. A distributed ledger approach to digital twin secure data sharing. In *Data and Applications Security and Privacy XXXIII* (Cham, 2019), S. N. Foley, Ed., Springer International Publishing, pp. 281–300.
- [17] DIETZ, M., VIELBERTH, M., AND PERNUL, G. Integrating digital twin security simulations in the security operations center. In *Proceedings of the 15th International Conference on Availability, Reliability and Security* (New York, NY, USA, 2020), ARES ’20, ACM.
- [18] DONG, W., YANG, B., WANG, K., YAN, J., AND HE, S. A dual blockchain framework to enhance data trustworthiness in digital twin network. In *2021 IEEE 1st International Conference on Digital Twins and Parallel Intelligence (DTPI) (2021)*, pp. 144–147.
- [19] ECKHART, M., AND EKELHART, A. A Specification-Based State Replication Approach for Digital Twins. In *Proceedings of the 2018 Workshop on Cyber-Physical Systems Security and Privacy* (2018), CPS-SPC ’18, ACM, p. 36–47.



- [20] ECKHART, M., AND EKELHART, A. Towards Security-Aware Virtual Environments for Digital Twins. In *Proceedings of the 4th ACM Workshop on Cyber-Physical System Security (CPSS '18)* (2018), pp. 61–72.
- [21] ECKHART, M., AND EKELHART, A. *Digital Twins for Cyber-Physical Systems Security: State of the Art and Outlook*. Springer International Publishing, Cham, 2019, pp. 383–412.
- [22] ENDERS, M. R., AND HOSSBACH, N. Dimensions of digital twin applications - a literature review. In *25th Americas Conference on Information Systems, AMCIS 2019* (2019), AIS.
- [23] FENG, H., CHEN, D., AND LV, H. Sensible and secure iot communication for digital twins, cyber twins, web twins. *Internet of Things and Cyber-Physical Systems 1* (2021), 34–44.
- [24] GEHRMANN, C., AND GUNNARSSON, M. A digital twin based industrial automation and control system security architecture. *IEEE Transactions on Industrial Informatics 16* (2020), 669–680.
- [25] GUO, Y., YAN, A., AND WANG, J. Cyber security riskanalysis of physical protection systems of nuclear power plants and research on the cyber security test platform using digital twin technology. In *2021 International Conference on Power System Technology (POWERCON)* (2021), pp. 1889–1892.
- [26] HADAR, E., KRAVCHENKO, D., AND BASOVSKIY, A. Cyber digital twin simulator for automatic gathering and prioritization of security controls' requirements. In *2020 IEEE 28th International Requirements Engineering Conference (RE)* (2020), pp. 250–259.
- [27] HEVNER, A. R., MARCH, S. T., PARK, J., AND RAM, S. Design Science in Information Systems Research. *MIS Quarterly 28*, 1 (2004), 75–105.
- [28] HOSSEN, T., GURSOY, M., AND MIRAFZAL, B. Digital twin for self-security of smart inverters. In *2021 IEEE Energy Conversion Congress and Exposition (ECCE)* (2021), pp. 713–718.
- [29] JIANG, L., ZHENG, H., TIAN, H., XIE, S., AND ZHANG, Y. Cooperative federated learning and model update verification in blockchain empowered digital twin edge networks. *IEEE Internet of Things Journal* (2021), 1–1.
- [30] JONES, D., SNIDER, C., NASSEHI, A., YON, J., AND HICKS, B. Characterising the digital twin: A systematic literature review. *CIRP Journal of Manufacturing Science and Technology 29* (2020), 36–52.
- [31] KANAK, A., UGUR, N., AND ERGUN, S. A visionary model on blockchain-based accountability for secure and collaborative digital twin environments. In *2019*

- IEEE International Conference on Systems, Man and Cybernetics (SMC)* (2019), pp. 3512–3517.
- [32] KARAARSLAN, E., AND BABIKER, M. Digital twin security threats and countermeasures: An introduction. In *2021 International Conference on Information Security and Cryptology (ISCTURKEY)* (2021), pp. 7–11.
- [33] KAUR, M. J., MISHRA, V. P., AND MAHESHWARI, P. The Convergence of Digital Twin, IoT, and Machine Learning: Transforming Data into Action. 2020, pp. 3–17.
- [34] KOLIAS, C., KAMBOURAKIS, G., STAVROU, A., AND VOAS, J. DDoS in the IoT: Mirai and Other Botnets. *Computer* 50, 7 (2017), 80–84.
- [35] KRITZINGER, W., KARNER, M., TRAAAR, G., HENJES, J., AND SIHN, W. Digital Twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine* 51, 11 (2018), 1016 – 1022.
- [36] LAAKI, H., MICHE, Y., AND TAMMI, K. Prototyping a digital twin for real time remote control over mobile networks: Application of remote surgery. *IEEE Access* 7 (2019), 20325–20336.
- [37] LANGNER, R. Stuxnet: Dissecting a cyberwarfare weapon. *IEEE Security Privacy* 9, 3 (2011), 49–51.
- [38] LEE, R. M., ASSANTE, M. J., AND CONWAY, T. Analysis of the Cyber Attack on the Ukrainian Power Grid, 2016. [https://ics.sans.org/media/E-ISAC\\_SANS\\_Ukraine\\_DUC\\_5.pdf](https://ics.sans.org/media/E-ISAC_SANS_Ukraine_DUC_5.pdf), accessed 2021-03-15.
- [39] LIAO, S., WU, J., BASHIR, A. K., YANG, W., LI, J., AND TARIQ, U. Digital twin consensus for blockchain-enabled intelligent transportation systems in smart cities. *IEEE Transactions on Intelligent Transportation Systems* (2021), 1–11.
- [40] LIU, H., TU, J., LIU, J., ZHAO, Z., AND ZHOU, R. Generative adversarial scheme based GNSS spoofing detection for digital twin vehicular networks. In *Wireless Algorithms, Systems, and Applications - 16th International Conference, WASA 2021, Nanjing, China, June 25-27, 2021, Proceedings, Part III* (2021), Z. Liu, F. Wu, and S. K. Das, Eds., vol. 12939 of *Lecture Notes in Computer Science*, Springer, pp. 367–374.
- [41] LIU, J., ZHANG, S., LIU, H., AND ZHANG, Y. Distributed collaborative anomaly detection for trusted digital twin vehicular edge networks. In *Wireless Algorithms, Systems, and Applications - 16th International Conference, WASA 2021, Nanjing, China, June 25-27, 2021, Proceedings, Part II* (2021), Z. Liu, F. Wu, and S. K. Das, Eds., vol. 12938 of *Lecture Notes in Computer Science*, Springer, pp. 378–389.
- [42] LIU, M., FANG, S., DONG, H., AND XU, C. Review of digital twin about concepts, technologies, and industrial applications. *Journal of Manufacturing Systems* 58 (2021), 346–361. Digital Twin towards Smart Manufacturing and Industry 4.0.

- [43] LU, Y., HUANG, X., ZHANG, K., MAHARJAN, S., AND ZHANG, Y. Communication-efficient federated learning and permissioned blockchain for digital twin edge networks. *IEEE Internet of Things Journal* 8, 4 (2021), 2276–2288.
- [44] LU, Y., HUANG, X., ZHANG, K., MAHARJAN, S., AND ZHANG, Y. Low-latency federated learning and blockchain for edge association in digital twin empowered 6g networks. *IEEE Transactions on Industrial Informatics* 17, 7 (2021), 5098–5107.
- [45] MALAKUTI, S., AND GRÜNER, S. Architectural Aspects of Digital Twins in IIoT Systems. In *Proceedings of the 12th European Conference on Software Architecture: Companion Proceedings* (2018), ECSA '18, ACM, pp. 1 – 2.
- [46] MCLAUGHLIN, S., KONSTANTINOVA, C., WANG, X., DAVI, L., SADEGHI, A., MANIATAKOS, M., AND KARRI, R. The cybersecurity landscape in industrial control systems. *Proceedings of the IEEE* 104, 5 (2016), 1039–1057.
- [47] MILLER, S., BRUBAKER, N., KAPPELMANN ZAFRA, D., AND CABAN, D. Triton actor ttp profile, custom attack tools, detections, and attck mapping. <https://www.fireeye.com/blog/threat-research/2019/04/triton-actor-ttp-profile-custom-attack-tools-detections.html>, 2019. Accessed: 2020-02-23.
- [48] MINERVA, R., LEE, G. M., AND CRESPI, N. Digital twin in the iot context: A survey on technical features, scenarios, and architectural models. *Proceedings of the IEEE* 108, 10 (2020), 1785–1824.
- [49] NEGRI, E., FUMAGALLI, L., AND MACCHI, M. A Review of the Roles of Digital Twin in CPS-based Production Systems. *Procedia Manufacturing* 11 (2017), 939–948.
- [50] ÖSTERLE, H., BECKER, J., FRANK, U., HESS, T., KARAGIANNIS, D., KRCMAR, H., LOOS, P., MERTENS, P., OBERWEIS, A., AND SINZ, E. J. Memorandum on design-oriented information systems research. *European Journal of Information Systems* 20 (2011), 7–10.
- [51] PEDERSEN, A. B., RISIUS, M., AND BECK, R. A ten-step decision path to determine when to use blockchain technologies. *MIS Quarterly Executive* 18, 2 (2019), 99–115.
- [52] PORTER, M. E., AND HEPPELMANN, J. E. How smart, connected products are transforming competition. *Harvard business review* 92, 11 (2014), 64–88.
- [53] PUTZ, B., DIETZ, M., EMPL, P., AND PERNUL, G. Ethertwin: Blockchain-based secure digital twin information management. *Information Processing Management* 58, 1 (2021), 102425.
- [54] RUBIO, J. E., ALCARAZ, C., ROMAN, R., AND LOPEZ, J. Current cyber-defense trends in industrial control systems. *Computers & Security* 87 (2019), 101561.

- [55] RUBIO, J. E., ROMAN, R., AND LOPEZ, J. Analysis of cybersecurity threats in industry 4.0: The case of intrusion detection. In *Critical Information Infrastructures Security* (Cham, 2018), G. D'Agostino and A. Scala, Eds., Springer International Publishing, pp. 119–130.
- [56] SAAD, A., FADDEL, S., YOUSSEF, T., AND MOHAMMED, O. A. On the implementation of iot-based digital twin for networked microgrids resiliency against cyber attacks. *IEEE Transactions on Smart Grid* 11, 6 (2020), 5138–5150.
- [57] SARACCO, R. The rise of Digital Twins. <https://cmte.ieee.org/futuredirections/2018/01/16/the-rise-of-digital-twins/>, 2018. [Online; accessed 09-Feb-2022].
- [58] SARACCO, R. Digital Twins in Industry. <https://cmte.ieee.org/futuredirections/2021/09/21/digital-twins-in-industry/>, 2021. [Online; accessed 09-Feb-2022].
- [59] SARACCO, R. Technology Predictions 2022: Digital Twins in Manufacturing. <https://cmte.ieee.org/futuredirections/2022/01/29/technology-predictions-2022-digital-twins-in-manufacturing/>, 2022. [Online; accessed 09-Feb-2022].
- [60] SELLITTO, G. P., ARANHA, H., MASI, M., AND PAVLESKA, T. Enabling a zero trust architecture in smart grids through a digital twin. In *Dependable Computing - EDCC 2021 Workshops* (Cham, 2021), R. Adler, A. Bennaceur, S. Burton, A. Di Salle, N. Nostro, R. L. Olsen, S. Saidi, P. Schleiss, D. Schneider, and H.-P. Schwefel, Eds., Springer International Publishing, pp. 73–81.
- [61] SELLITTO, G. P., MASI, M., PAVLESKA, T., AND ARANHA, H. A cyber security digital twin for critical infrastructure protection: The intelligent transport system use case. In *The Practice of Enterprise Modeling* (Cham, 2021), E. Serral, J. Stirna, J. Ralyté, and J. Grabis, Eds., Springer International Publishing, pp. 230–244.
- [62] SHAFTO, M., CONROY, M., DOYLE, R., GLAESSGEN, E., KEMP, C., LEMOIGNE, J., AND WANG, L. Modeling, simulation, information technology & processing roadmap. *National Aeronautics and Space Administration* (2012).
- [63] SHEN, W., HU, T., ZHANG, C., AND MA, S. Secure sharing of big digital twin data for smart manufacturing based on blockchain. *Journal of Manufacturing Systems* 61 (2021), 338–350.
- [64] SUSILA, N., SRUTHI, A., AND USHA, S. Chapter ten - impact of cloud security in digital twin. In *The Digital Twin Paradigm for Smarter Systems and Environments: The Industry Use Cases*, P. Raj and P. Evangeline, Eds., vol. 117 of *Advances in Computers*. Elsevier, 2020, pp. 247–263.
- [65] TANKARD, C. Advanced Persistent threats and how to monitor and deter them. *Network Security* 2011, 8 (2011), 16 – 19.

- [66] TAO, F., ZHANG, H., LIU, A., AND NEE, A. Y. C. Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics* 15, 4 (2019), 2405–2415.
- [67] UHLEMANN, T. H., LEHMANN, C., AND STEINHILPER, R. The Digital Twin: Realizing the Cyber-Physical Production System for Industry 4.0. In *Procedia CIRP* (2017), vol. 61, Elsevier B.V., pp. 335–340.
- [68] VAN DER VALK, H., HASSE, H., MÖLLER, F., ARBTER, M., HENNING, J.-L., AND OTTO, B. A taxonomy of digital twins. In *26th Americas Conference on Information Systems, AMCIS 2019* (2020), AIS.
- [69] VIELBERTH, M., BÖHM, F., FICHTINGER, I., AND PERNUL, G. Security operations center: A systematic study and open challenges. *IEEE Access* 8 (2020), 227756–227779.
- [70] VIELBERTH, M., GLAS, M., DIETZ, M., KARAGIANNIS, S., MAGKOS, E., AND PERNUL, G. A digital twin-based cyber range for soc analysts. In *Data and Applications Security and Privacy XXXV* (Cham, 2021), K. Barker and K. Ghazinour, Eds., Springer International Publishing, pp. 293–311.
- [71] WEBSTER, J., AND WATSON, R. T. Analyzing the past to prepare for the future: Writing a literature review. *MIS Quarterly* 26 (2002).
- [72] WORTMANN, F., AND FLÜCHTER, K. Internet of Things: Technology and Value added. *Business & Information Systems Engineering* 57, 3 (2015), 221–224.
- [73] WÜST, K., AND GERVAIS, A. Do you Need a Blockchain? In *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)* (2018), pp. 45–54.
- [74] XU, J., HE, C., AND LUAN, T. H. Efficient authentication for vehicular digital twin communications. In *2021 IEEE 94th Vehicular Technology Conference (VTC2021-Fall)* (2021), pp. 1–5.
- [75] ZHENG, Y., LU, R., GUAN, Y., ZHANG, S., AND SHAO, J. Towards private similarity query based healthcare monitoring over digital twin cloud platform. In *2021 IEEE/ACM 29th International Symposium on Quality of Service (IWQOS)* (2021), pp. 1–10.

