

Privacy-Preserving Data Certification in the Internet of Things: Leveraging Blockchain Technology to Protect Sensor Data

Mathieu Chanson¹, Andreas Bogner², Dominik Bilgeri³, Elgar Fleisch⁴, Felix Wortmann⁵

¹ ETH Zurich, Switzerland, mchanson@ethz.ch

² ETH Zurich, Switzerland

³ ETH Zurich, Switzerland

⁴ ETH Zurich and University of St. Gallen, Switzerland

⁵ University of St. Gallen, Switzerland

Abstract

A constantly growing pool of smart, connected Internet of Things (IoT) devices poses completely new challenges for business regarding security and privacy. In fact, the widespread adoption of smart products might depend on the ability of organizations to offer systems that ensure adequate sensor data integrity while guaranteeing sufficient user privacy. In light of these challenges, previous research indicates that blockchain technology may be a promising means to mitigate issues of data security arising in the IoT. Building upon the existing body of knowledge, we propose a design theory, including requirements, design principles, and features, for a blockchain-based sensor data protection system (SDPS) that leverages data certification. We then design and develop an instantiation of an SDPS (CertifiCar) in three iterative cycles that prevents the fraudulent manipulation of car mileage data. Furthermore, we provide an ex-post evaluation of our design theory considering CertifiCar and two additional use cases in the realm of pharmaceutical supply chains and energy microgrids. The evaluation results suggest that the proposed design ensures the tamper-resistant gathering, processing, and exchange of IoT sensor data in a privacy-preserving, scalable, and efficient manner.

Keywords: Internet of Things, Big Data, Privacy, Security, Blockchain, Certification, Design Science Research, Design Theory.

1 Introduction

In recent years, new forms of information technology (e.g., sensors and mobile devices) have dramatically expanded what can be measured and analyzed, thereby posing completely new challenges regarding security and privacy (Lee, Cho, & Lim, 2018; Newell & Marabelli, 2015; Sicari, Rizzardi, Grieco, & Coen-Porisini, 2015; Weber, 2010). The potential for information systems (IS)-related security and privacy issues to affect customers in their daily lives and private spheres makes these challenges top business priorities (Sicari et al., 2015). In fact, the widespread adoption of smart products might depend on the ability of organizations to offer systems that ensure adequate security levels while guaranteeing sufficient user privacy (Sicari, Cappiello, De Pellegrini, Miorandi, & Coen-Porisini, 2016). Such Internet of Things (IoT) systems, referring to a constantly growing pool of smart, connected devices, including cars, health applications, and industry machinery, offer adversaries a whole new range of attack vectors for manipulating information systems (Lowry, Dinev, & Willison, 2017; Porter & Heppelmann, 2015). IoT systems are usually characterized by multi-party ecosystems, with data pipelines crossing organizational borders (Aggarwal, Ashish, & Sheth, 2013; Roman, Zhou, & Lopez, 2013). In such systems, malicious adversaries can manipulate “data at various stages in the [processing] pipeline”, from sensor to service, making data integrity a key concern (Aggarwal et al., 2013, p. 419). The IS research community is well aware of these challenges and has specifically called for more design research that can facilitate secure and reliable data processing and exchange in multi-party ecosystems (Bélanger & Crossler, 2011; Pavlou, 2011; Smith, Dinev, & Xu, 2011).

Previous research has indicated that blockchain technology is a promising means to mitigate issues of data security arising in the IoT and has some decisive advantages

over a conventional database system on central servers (Glaser, 2017; Hyvärinen, Risius, & Friis, 2017; Nærland, Müller-Bloch, Beck, & Palmund, 2017). More specifically, blockchains provide tamper-proof storage capabilities in the form of a distributed ledger that can be used to securely store and exchange IoT sensor data. However, core challenges, such as privacy, scalability, and potentially prohibitive transaction costs, remain to be addressed (Beck, Stenum Czepluch, Lollike, & Malone, 2016; Notheisen, Cholewa, & Shanmugam, 2017; Risius & Spohrer, 2017). While there are a variety of different blockchain-based IoT systems currently under development (Curtis, 2015; Mengelkamp et al., 2018; Modum, 2018), the corresponding academic research is still in its infancy (Avital, Beck, King, Rossi, & Teigland, 2016; Beck, Avital, Rossi, & Thatcher, 2017; Beck et al., 2016; Beck & Müller-Bloch, 2017; Lindman, Rossi, & Tuunainen, 2017).

In the IS community, privacy and security have been widely discussed as multidisciplinary, diverse concepts (Lowry et al., 2017; Oetzel & Spiekermann, 2014; Sicari et al., 2015). However, most studies do not provide actionable solutions. In this regard, Bélanger and Crossler (2011) note in their seminal literature review that scholars should “conduct design and action research with an eye towards actual implementation” (p. 1035). Similarly, Pavlou (2011) proposes that future IS security and privacy studies should adapt the design science perspective, “with emphasis on building actual implementable tools” (p. 980). While multiple technologies are available to realize IoT sensor data protection systems (SDPSs) (Ayoade, Karande, Khan, & Hamlen, 2018; Machado & Fröhlich, 2018; Margulies, 2015), limited prescriptive knowledge has been gathered to guide the development process of such systems. In addition, the potential of blockchain technology in SDPSs is, to the best of our knowledge, not yet reflected in the literature. Against this background, we

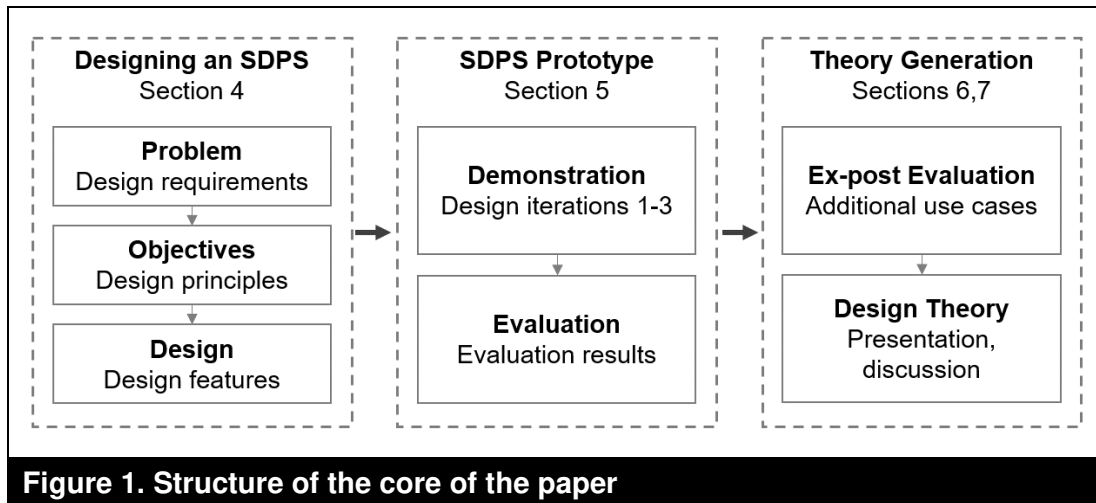
contribute to the IS literature by establishing theoretical insights into how to design an SDPS and by explicitly developing and evaluating a blockchain-based SDPS. More specifically, we aim to answer the following research questions:

- RQ1:** Which fundamental challenges arise in the context of IoT sensor data protection, and which requirements can be derived from these challenges for the design of information systems that facilitate IoT sensor data protection (i.e., SDPS)?
- RQ2:** Which actionable guidelines in the form of design principles and design features address these design requirements and inform the development of SDPS?
- RQ3:** What is the value proposition of blockchain technology in the realm of SDPSs, and what fundamental design implications of blockchain-based SDPSs must be considered?

Overall, our research is geared towards a design theory that guides the development of SDPSs that are able to protect IoT sensor data in a privacy-preserving manner. To answer our research questions, we follow the guidelines of design science research (DSR) (Gregor & Jones, 2007; March & Smith, 1995). Within the IS community, the development of design knowledge, be it in the form of design theories, principles, or guidelines, is of high significance for both research and practice (Baskerville, 2008; Hevner, March, Park, & Ram, 2004; Winter, 2008) and continues to attract a great deal of interest (Baskerville, Kaul, & Storey, 2015; Gregor & Hevner, 2013; Rai, 2017). We derive an artifact that consists of a set of interrelated design requirements, design principles, and design features. We demonstrate and refine our artifact on the basis of an instantiation that aims to prevent the fraudulent manipulation

of car mileage data. Finally, we provide an ex-post evaluation of the artifact and present our results in the form of a design theory.

The remainder of the paper is structured as follows. In Section 2, we introduce the practical issues that motivated this study and provide an overview of the related literature, thus laying the groundwork for addressing RQ1. Section 3 elaborates upon the design science research approach applied. The next four sections form the core of the paper and are depicted in Figure 1. In Section 4, we first describe the SDPS design requirements (RQ1) and proceed with the design principles and features (RQ2). In Section 5, we present the iterative development and evaluation of our artifact. Additionally, we evaluate the system ex-post in Section 6. Thereby, we confirm and refine the conclusions of RQ1 and RQ2 and form the foundation to answer RQ3. In Section 7, we present our results on RQ1 and RQ2 in the form of a design theory, focus on RQ3 and the design implications, and present our contributions. The paper concludes with Section 8, which reflects on the potential limitations and presents promising avenues for future research.



2 Foundations

2.1. *Internet of Things and Sensor Data*

By dramatically expanding what can be measured and analyzed, digitization is predicted to affect all areas of our lives (McAfee, Brynjolfsson, Davenport, Patil, & Barton, 2012; Newell & Marabelli, 2015). Digitization refers to the technical transformation of information processing from analog to digital and, in a broader sense, to the ever-increasing use of digital technology and its associated economic and social implications (Loebbecke & Picot, 2015; Nambisan, Lyytinen, Majchrzak, & Song, 2017; Negroponte, 1995). In the course of the ongoing digital transformation, a growing amount of intelligent, connected devices, including industrial machinery, cars, and health applications, will traverse the traditional separation of the physical and digital worlds (Porter & Heppelmann, 2015). The merger of these two worlds is widely referred to as the IoT and has recently gained significant attention in the IS literature and among practitioners (Iansiti & Lakhani, 2014; Loebbecke & Picot, 2015).

According to Atzori et al. (2010), the IoT refers to “a vision that virtually any physical object can be connected to the Internet”, a vision in which smart, connected devices generate unprecedented amounts of sensor data that can be classified as “big data” (H. Chen, Chiang, & Storey, 2012). Big data, in turn, is characterized by the ever-increasing volume, velocity, and variety of data combined with veracity-related challenges (Clarke, 2016; Goes, 2014; Schroeck, Shockley, Smart, Romero-Morales, & Tufano, 2012). This holds particularly for sensor data, which is increasing extraordinarily both in the size and speed of data generation (Abbasi, Sarker, & Chiang, 2016; Brynjolfsson & McAfee, 2012). In addition, sensor data is available in a variety of formats and from disparate sources (Brynjolfsson & McAfee, 2012;

Schroeck et al., 2012). Finally, veracity considers the varying degrees of reliability and credibility of sensor data sources (Abbasi et al., 2016). In light of these growing datasets and the corresponding technical and economic challenges, companies are increasingly relying on cloud solutions, which are typically operated by third parties (Lowry et al., 2017). In addition, more and more companies exchange and share sensor data to foster cross-organizational collaborations (Anderson, Baskerville, & Kaul, 2017).

2.2. Security and Privacy in the Internet of Things

The constantly growing pool of smart, connected IoT devices poses completely new challenges regarding security and privacy (Lee et al., 2018; Sicari et al., 2015; Weber, 2010). Companies are increasingly moving towards cloud solutions and sharing sensor data in multi-party ecosystems (Anderson et al., 2017; Lowry et al., 2017). However, distributed processing and sharing data with third parties is risky, as participating stakeholders (companies and end users) might misuse or lose control over data (Anderson et al., 2017; Moura & Serrão, 2016). Ultimately, the involvement of third parties significantly increases the risk of security and privacy breaches of IS systems (Lowry et al., 2017). In addition to intentional sharing in multi-party networks, unintentional access by malicious adversaries is a major security risk in the IoT, especially because of its “architecture of wireless transmitters and sensors that [...] connect into vast global networks” (Lowry et al., 2017, p. 556). For example, the internet connectivity of IoT devices can enable malware to quickly infect large populations around the globe (Kolias, Kambourakis, Stavrou, & Voas, 2017). Even the networking capabilities of devices that are not connected to the internet can be exploited to spread malware quickly and unobtrusively (Ronen, Shamir, Weingarten, & O’Flynn, 2017). This is because IoT sensors are usually unsupervised when

collecting data, leaving them particularly prone to various security threats (Aggarwal et al., 2013; Atzori et al., 2010; Ronen et al., 2017). The multilayered hardware and software stack of IoT solutions also makes these systems vulnerable to a variety of potential attacks (Sicari et al., 2016). For instance, malicious adversaries can manipulate “data at various stages in the [data processing] pipeline”, from sensor to service, making data integrity a key concern (Aggarwal et al., 2013, p. 419). Furthermore, many of the existing security principles that companies use to protect their systems, including routers, gateways, and firewalls, are not applicable to the IoT, as they “simply do not work for smaller and more mobile ‘things’” (Lowry et al., 2017, p. 556).

Against this background, the IoT fundamentally challenges the field of IS security and privacy, requiring the redefinition of well-established rules and organizational practices to protect sensor data (Fernandes, Rahmati, Eykholt, & Prakash, 2017; Singh, Pasquier, Bacon, Ko, & Evers, 2016). The confidentiality and integrity of data are essential to security and privacy to ensure that personal data cannot be viewed or manipulated by objectionable third parties (Anderson et al., 2017; Baskerville & Siponen, 2002; Chellappa & Pavlou, 2002). Specifically, privacy is commonly defined as “the ability of the individual to personally control information about oneself” (Stone, Gueutal, Gardner, & McClure, 1983, p. 460). Westin (1967) refers, in particular, to the possibility of data generators to determine the manner, scope, and time in which data is collected by, and transferred to, third parties. The existing IS studies on privacy cover a wide range of aspects and perspectives (Dinev, Hart, & Mullen, 2008; Malhotra, Kim, & Agarwal, 2004; Xu, Dinev, Smith, & Hart, 2011). However, despite the existing body of knowledge, there is a lack of actionable solutions, as Bélanger and Crossler (2011) conclude in their seminal literature review. Specifically, they

emphasize that beyond providing conceptual contributions towards the privacy debate, IS research should “conduct design and action research with an eye towards actual implementation” (p. 1035), developing tools to protect information privacy.

In summary, the IoT is advancing much faster than the related privacy and security measures and policies (Singh et al., 2016; Weber, 2010). The resulting security and privacy gaps are potentially dangerous loopholes that can be exploited by malicious actors to the detriment of consumers and organizations (Aggarwal et al., 2013; Lowry et al., 2017). In fact, the widespread adoption of IoT solutions might depend on organizations’ capabilities to offer systems that ensure adequate security levels while guaranteeing sufficient user privacy (Sicari et al., 2016). As such, the IoT, which is characterized by multi-stage data pipelines and big (sensor) data, has been identified by IS scholars as being “particularly compelling to security and privacy researchers”, as it carries “innate information and privacy risks” (Lowry et al., 2017, p. 546).

2.3. Existing Research on SDPSs and their Limitations

SDPSs, which aim to ensure the security and privacy of sensor data, are the subject of an extensive body of literature. In particular, the IS community has made considerable effort to investigate issues of security and privacy (Bulgurcu, Cavusoglu, & Benbasat, 2010; Chatterjee, Sarker, & Valacich, 2015; Y. Chen & Zahedi, 2016), which has resulted in various design theories (Heikka, Baskerville, & Siponen, 2006; Siponen & Iivari, 2006). A key research focus in the area of IS security is the use of organizational policies that define how the users of information systems should prevent, identify, and react in security incidents (Anderson et al., 2017; Cram, Proudfoot, & D’Arcy, 2017; Moody, Siponen, & Pahlila, 2018; Niemimaa & Niemimaa, 2017). An excellent review of the body of knowledge is provided by Cram et al. (2017),

who analyzed 114 security policy-related journal articles. From this research stream, the study of Anderson et al. (2017) is especially relevant for our work. They combine discussions of security with those of information privacy, focusing on the risks and rewards of either sharing or retaining full control over data. Thus, they cover a topic that is also fundamental to SDPSs, namely, security, privacy, and the necessity, or economic benefit, of sharing information. However, similar to the approach of other literature on organizational policies, Anderson et al. (2017) deliberately refrain from providing actionable guidelines for the implementation of information systems that would enable secure and privacy-preserving data exchange. Rather, they focus on how an organization and its personnel should behave in the vicinity of such systems. A lack of normative results can be similarly observed in most other examples of IS research on SDPSs (Crossler & Posey, 2017). This finding is in line with the seminal literature review by Bélanger and Crossler (2011) on information privacy, in which the authors conclude that “very few articles provide design and action contributions” (p. 1023). Moreover, the IS literature on privacy and security hardly addresses the specific design challenges that arise when processing IoT sensor data.

Beyond the domain of IS, there is a fruitful knowledge base of computer science literature that specifically addresses security and privacy issues in the IoT. Core insights from the latest research include the summation that “the task of affordably supporting security and privacy [in the IoT is] quite challenging” (Trappe, Howard, & Moore, 2015, p. 14) and the observation that while some known security principles should be adaptable to the IoT computing paradigm, “the nature of both physical processes and IoT devices lend themselves to the construction of new security mechanisms” (Fernandes et al., 2017, p. 83). Inspired by such statements, there has been an active stream of research developing specific solutions in the realm of SDPSs

(Kolias, Stavrou, Voas, Bojanova, & Kuhn, 2016; Margulies, 2015; Ronen et al., 2017), including work on the potential value contribution of blockchain technology. Ayoade, Karande, Khan, and Hamlen (2018), for example, present a system for the management of IoT data in which all permissions for data access are enforced by smart contracts on a blockchain, which also ensures traceability by the logging of all data access requests. Liang, Zhao, Shetty, and Li (2017) present a system that leverages a public blockchain to ensure the integrity of data collected by drones and additionally secures the communication between the drone and its control system. Machado and Fröhlich (2018) propose a system that uses blockchain technology to enable the verification of the data integrity of IoT devices. More specifically, they present a proof of concept and evaluate the performance of the implemented data pipeline. While these studies contain detailed descriptions of specific prototypes, they lack both the codifications and the abstractions of the interrelated set of requirements that the system needs to fulfill, as well as the design principles and features that address these requirements. Both types of research results, however, are necessary to allow for generalizability beyond a specific solution to a specific problem. The importance of such a thorough conceptualization has been extensively discussed among scholars and is a key aspect of DSR (Baskerville & Pries-Heje, 2010; Gregor & Hevner, 2013; Gregor & Jones, 2007; Meth, Mueller, & Maedche, 2015). Therefore, we suggest that the contributions of these existing studies could be expanded substantially by reflecting state-of-the-art DSR guidelines and providing a thorough conceptualization.

Taken together, there is a rich body of knowledge in the IS community on security and privacy. However, scholars have specifically called for studies that develop actionable guidelines to facilitate the design of practical tools. To the best of our

knowledge, there are no examples of prior research dedicated to the design and actual implementation of SDPSs. Outside of the IS community, there is an active stream of research focused on the development of SDPSs, describing the technical design of prototypes in detail. However, these studies provide very specific solutions to equally specific problems. Thereby, they lack well-defined conceptualizations and thus generalizable results addressing an entire problem class. Finally, due to the novelty of blockchain technology, there has been a lack of reflection on the specific advantages and limitations of blockchain technology in SDPSs.

2.4. Blockchain Technology

A blockchain is a distributed transactional database that is cryptographically secured and controlled by a consensus mechanism (Beck et al., 2017). From an operational perspective, a blockchain comprises an event log storing transactions in such a way that they are immutable once submitted to the system (Moyano & Ross, 2017). Instead of storing the transactions on a central server, various copies of the data exist across different computers, otherwise known as nodes, that participate in the blockchain (Tschorsch & Scheuermann, 2016). This decentralization enables a distributed governance, with a “consensus mechanism between the participating nodes in the system” (Hyvärinen et al., 2017, p. 445), thus eliminating the need to trust other participants of the system (Egelund-Müller, Elsmann, Henglein, & Ross, 2017; Nakamoto, 2008; Notheisen et al., 2017). Blockchains only accept new entries if they obey a predefined protocol and are thus deemed valid (Nærland et al., 2017; Risius & Spohrer, 2017). Since the introduction of the initial blockchain application Bitcoin in 2009, different forms of distributed ledger technologies, or incarnations of blockchains, have emerged (Lindman et al., 2017; Nakamoto, 2008). In the paper at hand, we focus on public permissionless blockchains that enable secure transactions

in open ecosystems where the participants are not limited to known players, trust is not granted, and all participants are treated equally (Beck, Müller-Bloch, & Ling, 2018). In addition to the generic properties outlined above, this blockchain type is characterized by a specific set of criteria. The protocols of public permissionless blockchains, such as Ethereum, allow anyone to see any transaction and every node to submit and validate transactions on the blockchain, “thus providing maximum transparency and replicability of transactions” (Hyvärinen et al., 2017, p. 444; Tschorsch & Scheuermann, 2016). Since they are open source, anyone can use these blockchains free of charge and legally (Nærland et al., 2017). In addition, as long as one follows the predefined protocol, there is no gatekeeper limiting access to the blockchain (Beck et al., 2018). Finally, permissionless blockchains are extraordinarily resistant to malicious attempts at manipulation, because the cryptographic logic driving the consensus mechanism and the storage of the transaction log both rely on a decentralized implementation (Gervais et al., 2016). Compared to traditional information systems, public permissionless blockchains “avoid the need for copious, often duplicate documentation, third-party intervention, and remediation” (Underwood, 2016, p. 15). Against this background, blockchain technology is often perceived as groundbreaking and is predicted to fundamentally affect how business is conducted (e.g., Chanson, Gjoen, Risius, & Wortmann, 2018; Gomber, Kauffman, Parker, & Weber, 2018), as many industries depend on the fact “that individuals and organizations trust other entities to create, store, and distribute essential records” (Beck et al., 2017, p. 381).

The above-outlined blockchain properties are particularly useful for mitigating issues of data security arising in the IoT and have some decisive advantages over a conventional database system on central servers (Bogner, Chanson, & Meeuw, 2016;

Glaser, 2017; Hyvärinen et al., 2017). Indeed, there are a variety of different blockchain-based IoT systems currently under development. Well-known examples address use cases in car leasing (Curtis, 2015), pharmaceutical supply chains (Modum, 2018), and energy markets (Meeuw, Schopfer, Ryder, & Wortmann, 2018; Mengelkamp et al., 2018). Applying blockchain to IoT use cases has the potential to ensure the “protection of critical infrastructure and data” (Hyvärinen et al., 2017, p. 443). More specifically, blockchains provide tamper-proof storage capabilities in the form of a distributed ledger that can be used to securely store IoT sensor data. In addition, they enable secure ledger access on the basis of well-defined protocols. Finally, blockchain solutions are not operated by one single party (Bogner et al., 2016); hence, they are neutral and particularly suitable in ecosystem settings with multiple parties and potentially diverging interests. However, recent research has often had a view of blockchain technology that is overly optimistic (Beck et al., 2017), and the core blockchain challenges in the field of IoT remain to be solved. First, simply writing IoT sensor data to a public permissionless blockchain is an unacceptable practice in light of the highly sensitive IoT data that is gathered across all areas of our lives (Beck et al., 2017; Lowry et al., 2017). The specific privacy challenges arising in the IoT (see Lowry et al., 2017; Sicari et al., 2015) require adequate countermeasures to ensure the data privacy of public permissionless blockchain-based IoT systems (Beck et al., 2017; Fabian, Ermakova, & Sander, 2016). Second, public permissionless blockchains are known for their restrictions with respect to scalability as well as for their potentially prohibitive transaction costs (Beck et al., 2016; Risius & Spohrer, 2017). In summary, permissionless blockchain technology is a promising means to mitigate issues of data integrity and availability arising in the IoT. However,

some core challenges, such as privacy, scalability, and the potentially prohibitive transaction costs, remain to be addressed.

3 Methodology

3.1. Overall Research Design

We address the problems discussed in Section 2 through design science research (Gregor & Jones, 2007; March & Smith, 1995), and we base our specific research approach on the guidelines of Peffers, Tuunanen, Rothenberger, and Chatterjee (2007). Design science has its roots in the seminal work of Herbert Simon (Simon, 1969) and is anchored in many disciplines, such as engineering, architectural science, computer science, and economics (Baskerville, 2008; March & Smith, 1995). Within the IS community, the development of design knowledge is of high significance for both research and practice (Baskerville, 2008; Hevner et al., 2004; Winter, 2008) and continues to attract considerable interest (Baskerville et al., 2015; Gregor & Hevner, 2013; Rai, 2017). The focus of design science is on the creation of the artificial and accordingly the rigorous construction and evaluation of innovative artifacts. It aims to generate new knowledge about a specific and relevant problem class and corresponding solutions to that problem class (Hevner & Chatterjee, 2010). Hence, the creation of utility for practical application through the resulting artifact is one of the core goals of design science research (Hevner et al., 2004; Winter, 2008). While some scholars put their emphasis on the artifact and its relevance (Hevner et al., 2004; March & Smith, 1995), others stress the importance of contributions to theory (Gregor & Jones, 2007; Kuechler & Vaishnavi, 2008; Walls, Widmeyer, & El Sawy, 1992). However, it is widely agreed that impactful design science research arises through synergies between relevance and rigor, that is, the contributions to the

application environment as well as to theory (Gregor & Hevner, 2013). We build upon this understanding and elaborate in the following on both the role of theory as well as the general design of the research process.

Concerning the role of theory, we draw on Gregor and Jones (2007), who extend the work of Walls, Widemeyer, and El Sawy (1992) and note that theorizing is a key goal in DSR that may culminate in establishing an IS design theory. On the one hand, existing theory can serve, in the form of kernel theories, as justificatory knowledge and inputs for design cycles (Gregor & Jones, 2007). In particular, the design principles derived from such kernel theories may guide the implementation of an artifact (Walls et al., 1992). On the other hand, design theorizing should contribute to a novel design theory with the aim of formalizing knowledge in DSR (Gregor & Hevner, 2013; Gregor & Jones, 2007). This type of theory provides instructions that link design principles and features with actions. It is prescriptive in the sense that it provides rules and actionable guidelines and hence belongs to the theories of type five in Gregor's taxonomy (Gregor, 2006; Gregor & Hevner, 2013). Communicating such a design theory can be enabled by an artifact instantiation that embodies the related design principles and features (Gregor & Jones, 2007). An ex-post evaluation, in which additional slices of data are gathered after the original design cycles and the corresponding evaluations and are then used in an evaluation process to generate further theoretical insight, can be an important and constructive step to reach a sufficient abstraction level and theoretical saturation (Beck, Weber, & Gregory, 2013).

Concerning the general design of the research process, there is wide agreement that an iterative procedure of well-defined steps is most applicable for DSR (Hevner et al., 2004; Nunamaker Jr, Chen, & Purdin, 1990; Takeda, Veerkamp, & Yoshikawa, 1990). Since the recognition of DSR in the mainstream of IS with the publication of

Hevner et al.'s article (2004), the discourse within the IS community has been intense and ongoing regarding the specific structuring of this process. Many different approaches and improvements and derivatives thereof have been suggested by renowned scholars (Beck et al., 2013; Hevner, 2007; Peffers et al., 2007; Vaishnavi & Kuechler, 2015). Our project's research design is based on the guidelines of Peffers et al. (2007) and informed by the design approach of Meth et al. (2015). We extend Peffers et al.'s guidelines by considering an additional phase of ex-post evaluation (Pries-Heje, Baskerville, & Venable, 2008; Venable, Pries-Heje, & Baskerville, 2016) after finalizing the prototype, as suggested by Beck et al. (2013), which facilitates the generation of additional insight. Finally, to summarize the knowledge gathered, we follow Gregor and Jones (2007) and present our results in the form of a design theory.

3.2. Design Cycles

Based on the theoretical and procedural reflections above, we design our research project in three design cycles, each composed of five phases, which are followed by two final steps of evaluation and communication. This research design, the output of each phase, and the according iteration between conceptualization, development, instantiation, and evaluation, is outlined in Figure 2.

The first design cycle was initiated with an intensive literature review to identify the problem at hand and reflect on RQ1. Our examination of the topic was triggered by a report of the prevalence of odometer fraud (TÜV Rheinland, 2015). Developing systems that are able to securely process and exchange odometer sensor data arose as a main challenge in this study. Our literature review quickly expanded to similar issues regarding IoT sensor data present in other industries, such as pharma

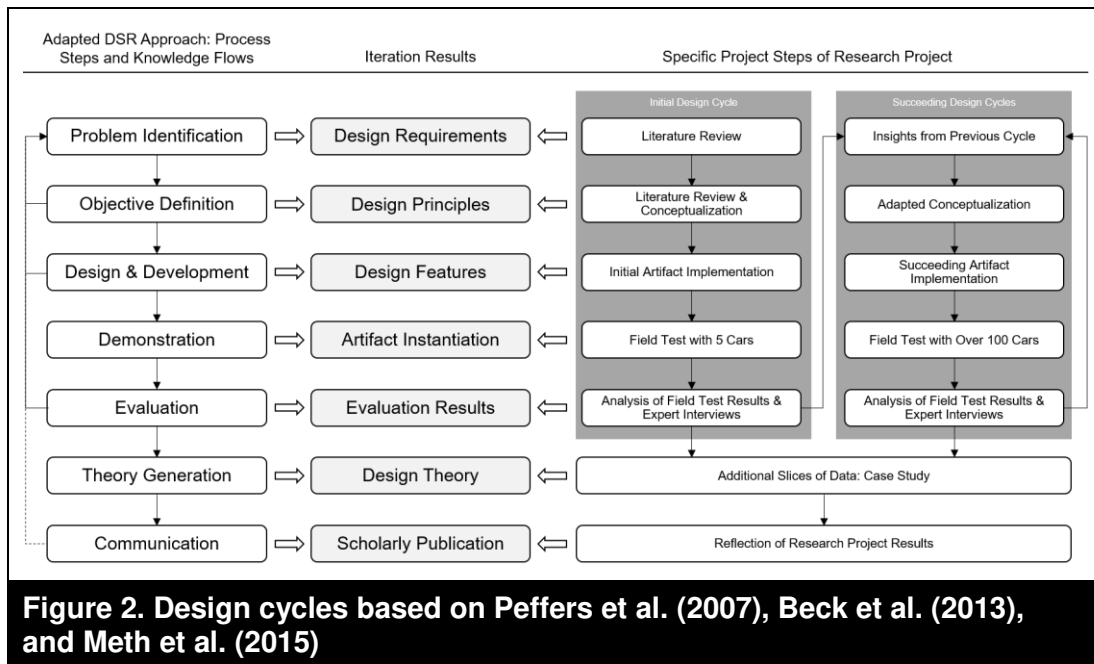


Figure 2. Design cycles based on Peffers et al. (2007), Beck et al. (2013), and Meth et al. (2015)

(Modum, 2018) and energy (Mengelkamp et al., 2018). This initial literature review allowed us to develop the first preliminary requirements for the artifact to be built. We then conducted a second literature review to find reference points in theory and the extant body of knowledge to refine these preliminary requirements, deepening the findings concerning RQ1. Based on this, we then derived design principles in the objective definition phase and identified the design features that are required to address these design principles, hence addressing RQ2. All these steps focused on the generalized problem class. In the next step, we instantiated the developed design with respect to a specific use case (prevention of odometer fraud) and developed the first version of our prototype CertifiCar. We evaluated this initial version of CertifiCar in a field test with five cars as well as on the basis of expert interviews. We used the results of this evaluation to adapt the artifact design in the second design cycle and, based on these changes, implemented a new version of our artifact. Again, we evaluated the artifact in a field test and on the basis of expert feedback. We integrated these findings into the third design cycle, which was run similarly to the second design

cycle and resulted in the final version of the artifact. The final version of CertifiCar was deployed in a field test with 100 cars, and the subsequent evaluation was based on the results of this field test and expert interviews. During these loops of development and evaluation, we iteratively refined the design requirements, principles, and features, enhancing the results to RQ1 and RQ2. Furthermore, the knowledge acquired in this phase built the foundation to approach RQ3. Ultimately, we gathered additional slices of data for a detailed ex-post evaluation of the derived design requirements, principles, and features of the artifact (Beck et al., 2013; Pries-Heje et al., 2008). This helped to confirm the validity of our responses to RQ1, RQ2, and RQ3 and led to diverse additional insights into RQ3.

In our conceptualization efforts, we follow three core design steps to derive the design requirements, principles, and features (Hevner & Chatterjee, 2010; March & Smith, 1995). In the first step, we develop design requirements based on the input from the problem identification step. The design requirements are generic requirements that should be met by any artifact aiming to create a solution for the underlying problem class. This notion of design requirements is closely related to the meta-requirements described by Walls et al. (1992) and the general requirements introduced by Baskerville and Pries-Heje (2010). In the second design step, we identify design principles based on the input of the suggestion step, for instance, by drawing on the extant information asymmetry literature. Our concept of design principles corresponds to the generic capabilities of an artifact through which the design requirements are addressed and relates these requirements indirectly with design features containing the technical specifics of the solution. This notion of design principles is closely linked to the meta-design introduced by Walls et al. (1992) and the relationship between general requirements and general components that

Baskerville and Pries-Heje (2010) emphasize. In the third step, we derive design features on the basis of the design principles and implement them in an instantiation of the artifact. These design features capture the technical specifics of the solution and are closely related to the general components described by Baskerville and Pries-Heje (2010). A design principle that is instantiated by an explicit design feature can be understood as an explanation (design principle) of why a specified piece (design feature) leads to a predefined goal (design requirement) (Kuechler & Vaishnavi, 2012). These explanations will assist us in abstracting the results of the instantiation of our prototype (CertifiCar) to a more generalized level and in creating a better understanding of the conceptual foundation of the design theory we propose.

As we reported above, we attempted to ensure the appropriate grounding and viability of the proposed design and its corresponding artifact instantiation in multiple iterations of our research design. Thereby, we distinguish between the interim evaluations at the end of each design cycle and the ultimate ex-post evaluation after finalizing the artifact development. In practice, in each design cycle, we use the last two phases to demonstrate and evaluate the current instantiation of the prototype, as the guidelines of Peffers et al. (2007) suggest. This procedure is detailed in Section 5, where we depict the iterative development of the prototype and the corresponding demonstrations and evaluations. Subsequently, we perform an additional ex-post evaluation (Pries-Heje et al., 2008), as suggested by Beck et al. (2013), to facilitate the generation of a novel theory. Specifically, we perform semi-structured interviews with nine experts on different security and privacy topics regarding IoT data to generalize and verify the viability of our proposed actionable guidelines, resulting in our final design theory. We only briefly discuss the interim evaluations and emphasize the ex-post evaluation because it focuses on the generalized problem class defined

by the design requirements derived and, contrary to the interim evaluations, not on the specifics of the prototype implemented in this study.

4 Designing an IoT Sensor Data Protection System

4.1. *Developing Design Requirements*

To derive the specific design requirements for an SDPS that enables the process of IoT sensor data generation, processing, and exchange, we built upon practically motivated problems that are outlined in the existing literature. More specifically, as outlined in the foundations section, studies of interest include the following: (1) research regarding the Internet of Things and sensor data (core key words: Internet of Things, IoT, cyber-physical systems, sensor data, big data, digital and digitization¹), (2) research regarding security and privacy (core key words: protection, security, secure, privacy, private, privacy-preserving, data, information and system¹), and (3) specific research focusing on systems that protect sensor data (core key words: Internet of Things, IoT, cyber-physical systems, sensor data, security, cybersecurity, attack, protection, privacy, private and privacy-preserving¹). To consolidate the existing research, we considered prestigious IS journals (i.e., the AIS basket of journals), international IS conferences (AMCIS, ECIS, ICIS, MCIS, PACIS), and high-quality journals with a specific focus on practical relevance (the Harvard Business Review, MIS Quarterly Executive, and MIT Sloan Management Review). Additional IS outlets were considered through the AIS eLibrary. With respect to research focused on systems that protect sensor data, we included the ACM Digital Library, as well as the IEEE Xplore Digital Library. Finally, we conducted a backward and forward search based on the gathered literature (Webster & Watson, 2002).

¹ Using respective combinations

A core challenge in IoT is security and data manipulation (Lowry et al., 2017). The IoT creates new security challenges, for instance, that the data collection nodes are typically left unattended for long periods of time (Aggarwal et al., 2013; Ronen et al., 2017). In addition, a data recipient cannot be sure if the received data is valid, because a malicious adversary, potentially the data owner himself, has the possibility to manipulate the data at several stages in the data pipeline (Aggarwal et al., 2013). Additional problems are introduced by the fact that the progress in deploying and developing the IoT is much faster than the accompanying security practices (Singh et al., 2016). Therefore, a recipient of IoT sensor data often encounters the problem that the data integrity cannot be taken for granted (Miorandi, Sicari, De Pellegrini, & Chlamtac, 2012; Sicari et al., 2015). Consequently, we derive the following design requirement:

DR1: Enable tamper-resistant data generation, processing, and exchange.

The process of IoT sensor data generation, processing, and exchange should be supported by systems that ensure tamper resistance throughout the whole data pipeline.

A second challenge in the realm of IoT sensor data is privacy (Lee et al., 2018; Sicari et al., 2015). More specifically, there is a lack of well-established privacy-preserving mechanisms (Bélanger & Crossler, 2011). This is especially striking because IoT sensors often have access to very detailed personal data (Lowry et al., 2017). In addition, users are often not able to determine which data is recorded and transmitted (Davenport, 2013; Westin, 1967). Home assistance devices, such as Amazon Alexa and Google Home, are always on, although most of the time they are neither supposed to store nor transmit recorded information. Similar thoughts apply to other devices deployed inside the home of a user. Therefore, the goal of any data

processing system in the realm of IoT is to preserve privacy (Alqassem & Svetinovic, 2014; Sicari et al., 2016). Consequently, we derive the following design requirement:

DR2: Enable privacy-preserving data generation, processing, and exchange.

The process of IoT sensor data generation, processing, and exchange should be supported by systems that are capable of preserving the privacy of the corresponding data owner.

A third challenge is related to IoT and big data. As we have outlined, the technical transformation of information processing from analog to digital and the according merger of the physical and digital worlds are expected to generate unprecedented amounts of data (Lowry et al., 2017; Porter & Heppelmann, 2015). Hence, systems that enable tamper-resistant data generation and exchange must be able to cope with “big data” (H. Chen et al., 2012). To operate in such a context, a corresponding system should have sufficient throughput to handle the expected amounts of data the IoT will generate. This aspect becomes particularly relevant when using blockchain technology, as many of the existing blockchain technologies are still struggling with scalability problems (Hyvärinen et al., 2017; Tschorsch & Scheuermann, 2016). Consequently, we derive the following design requirement:

DR3: Enable large data volume throughput. *The process of IoT sensor data generation, processing, and exchange should be supported by systems that are capable of processing the large amounts of data that are typical of IoT applications.*

Finally, the advantages of information systems must always be weighed against their disadvantages (Delone & McLean, 2003). In light of this fundamental economic principle, the IS-related costs are of particular importance in a business environment. Although this holds true for any IS, it is of special importance for solutions that rely on blockchain technology (Risius & Spohrer, 2017). As discussed above, the currently

unsolved issues regarding the scalability of different blockchain technologies and high transaction costs have the potential to generate substantial financial expenditures (Beck et al., 2016; Hyvärinen et al., 2017). Consequently, we derive the following design requirement:

DR4: Ensure economic feasibility. *The process of IoT sensor data generation, processing, and exchange should be supported by systems that ensure economic feasibility.*

Summing up, based on the fundamental SDPS challenges, we derived four general design requirements (see Table 1). These design requirements determine our design theory's purpose and scope that the design principles and design features must address to overcome or reduce the existing challenges (see Figure 3).

Table 1: General SDPS challenges and design requirements			
ID	SDPS challenge	SDPS design requirement	Main corresponding literature
1	Adversaries have the possibility to manipulate sensor data at several stages in the processing pipeline, so data integrity cannot be taken for granted.	SDPS should ensure tamper resistance throughout the whole data pipeline.	(Aggarwal et al., 2013; Lowry et al., 2017; Sicari et al., 2015)
2	IoT sensors can capture detailed and very sensitive personal data.	SDPS should be capable of preserving the privacy of the data owner.	(Bélanger & Crossler, 2011; Davenport, 2013; Lee et al., 2018; Sicari et al., 2016)
3	IoT sensors are able to generate vast amounts of data.	SDPS should provide sufficient data throughput to process large amounts of data.	(H. Chen et al., 2012; Hyvärinen et al., 2017; Porter & Heppelmann, 2015)
4	The protection of IoT sensor data can require substantial resources and induce significant costs.	SDPS should ensure economic feasibility, that is, the protection benefits have to outweigh the protection costs.	(Beck et al., 2016; Hyvärinen et al., 2017; Risius & Spohrer, 2017)

4.2. *Deriving Design Principles*

To address the design requirements, we build upon theory and the existing body of knowledge to derive design principles. With respect to DR1 (tamper-resistant data generation, processing, and exchange), theory on information asymmetry provides a fruitful basis to derive design principles. The (neo-)classical market model suggests that participants are fully informed about all goods (Albersmeier, Schulze, Jahn, & Spiller, 2009). However, business transactions are often characterized by fundamental information deficits (information asymmetries) that favor opportunistic behavior and restrict the smooth functioning of markets (Akerlof, 1970; Spence, 1976). To overcome these information deficits and avoid opportunistic behavior, certain measures such as certification, guarantees, or well-established brand names have been identified (Akerlof, 1970; Bond, 1982; Genesove, 1993).

With regard to the protection of sensor data, certification, in particular, appears to be a suitable measure to prevent opportunistic behavior (manipulation), as it is not restricted to companies that have high credibility or a strong brand name. Certification indicates the attainment of a certain quality level and is based on auditing (Akerlof, 1970). It most often relies on protection and investigation schemes that cover the whole supply (e.g., food business) chain or information (e.g., financial auditing) chain, as certain product and information qualities cannot be judged by inspections that are limited to the end of the chain (Albersmeier et al., 2009). This is particularly relevant for sensor data. Only in the case of very obvious manipulations is it possible to detect manipulated sensor data by means of a single inspection at a certain point in the information processing chain (e.g., when the odometer value of a car is equal to or even smaller than zero). Hence, the entire information chain from source (sensor) to sink (final data consumer) must be protected from manipulation, e.g., by applying an

appropriate means of encryption. By protecting the data along the entire information chain, it can be certified that the data was not manipulated on the way from the source to the sink.

DP1: Sensor data is certified on the basis of source to sink protection.

If data is protected from source to sink, data producers can be made accountable for the data they provide. However, in the case of sensor data, even if the information chain is protected from source to sink, data manipulation can still occur. More specifically, the data producer can focus on the source and manipulate the sensor or its environment. For example, anecdotal evidence and a corresponding patent² suggest that temperature sensors in cold chains are regularly covered with insulation material to hide shorter periods of irregularities. In cars, as a second example, mileage sensors (odometers) are multi-component systems that are connected by cables so that manipulating devices (“CAN filters”, “CAN blockers”) can be placed between them. More specifically, small sensing units often do not have the computing power for encryption or processing and hence communicate their raw sensor values to more powerful control units over wires that can be intercepted. Therefore, sensors are not per se monolithic components that are well protected and cannot be manipulated. To account for the corresponding manipulation risk, additional means might be required to enable trustworthy certification. More specifically, cross-validation and plausibility checks are common means in auditing (Whittington & Pany, 2015) that might also be used with sensor data to reveal manipulations. In the case of car mileage manipulation, for example, GPS data can be used to cross-validate the mileage data of a car.

DP2: Sensor data is certified on the basis of cross-validation.

² <https://patents.google.com/patent/DE10228648A1/de>

However, cross-validation and plausibility checks can only reduce the manipulation risk. Similar to financial auditing, a “detection risk” (Dong, Liao, & Zhang, 2018; Hogan & Wilkins, 2008) remains, which depicts the probability that manipulations are not detected. In summary, with the implementation of DP1, it can be certified that data was not manipulated on its way from the source to the sink, so that data producers can be made accountable for the data they provide. In addition, with the availability of cross-validation data and the implementation of DP2, it can be certified with an associated detection risk that the sensor or its environment were not manipulated.

With regard to DR2 (privacy-preserving data generation, processing, and exchange), we build upon Westin’s (1967) theory of privacy to derive a corresponding design principle. Westin’s theory is one of the best articulated and best supported theories of privacy (Margulis, 2011). A fundamental cornerstone of Westin’s theory is the existence of the following four states of privacy (Margulis, 2011): (1) solitude is about being free from observation by others, (2) intimacy is about the seclusion required to form close associations, (3) anonymity is about the condition of being unknown, and (4) reserve is about limiting disclosure to others. In essence, for Westin (1967, p. 7), privacy “is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others”.

At the core of Westin’s definition is the right of a data owner to have full control over the communication and use of her data. With respect to the exchange of sensor data, the data owner should therefore determine when and to what extent data is communicated and to whom. However, the means of exchange that determines the “how” is a software system. Hence, the data owner is limited in her privacy by the

restrictions of the system. If the system restricts privacy too much, though, the data owner still has the option to not use the system. In summary, in the context of sensor data exchange, we derive the following design principle that addresses DR2:

DP3: Data owners determine when and to what extent their certified data is communicated to others.

While the first two design requirements focus on what a sensor data exchange system should enable (tamper resistance and privacy), DR3 (large data volume throughput) and DR4 (economic feasibility) further qualify how the system should operate (scalable and thereby also cost efficient) and shift the focus from positive system outcomes (prevent manipulation, assure privacy) to possible negative outcomes (system costs). The existence of such positive and negative system outcomes is well reflected in IS theory. The DeLone and McLean Model of Information Systems Success captures the idea that the system impact has to reflect the balance of positive and negative impacts (Delone & McLean, 2003). The concept of “net benefits” depicts the rationale that “no outcome is wholly positive, without any negative consequences” (Delone & McLean, 2003, p. 22).

Applying the aforementioned rationale to tamper-resistant sensor data exchange, a potential solution has to ensure that the positive effects are not cancelled out by negative consequences. In respect to the design challenge at hand, the protection and certification of IoT sensor data can be resource-intensive and costly, especially in the context of large amounts of sensor data (Sicari et al., 2015). Hence, data has to be processed on a system architecture that is linearly scalable with respect to performance and costs. Thereby, the scalability captures how “well a particular solution fits a problem as the scope of that problem increases” (Schlossnagle, 2006). Linear scalability is an established concept that refers to the relationship between an

input (e.g., amount of sensor data) and an output (e.g., performance or cost) (Bonvin, 2012). While the term defines a very specific type of relationship (linear), it is often used in a broader sense. In contrast to negative or sublinear scalability (Williams & Smith, 2004), linear scalability depicts the idea that the performance does not erode and the costs of a system do not explode at scale.

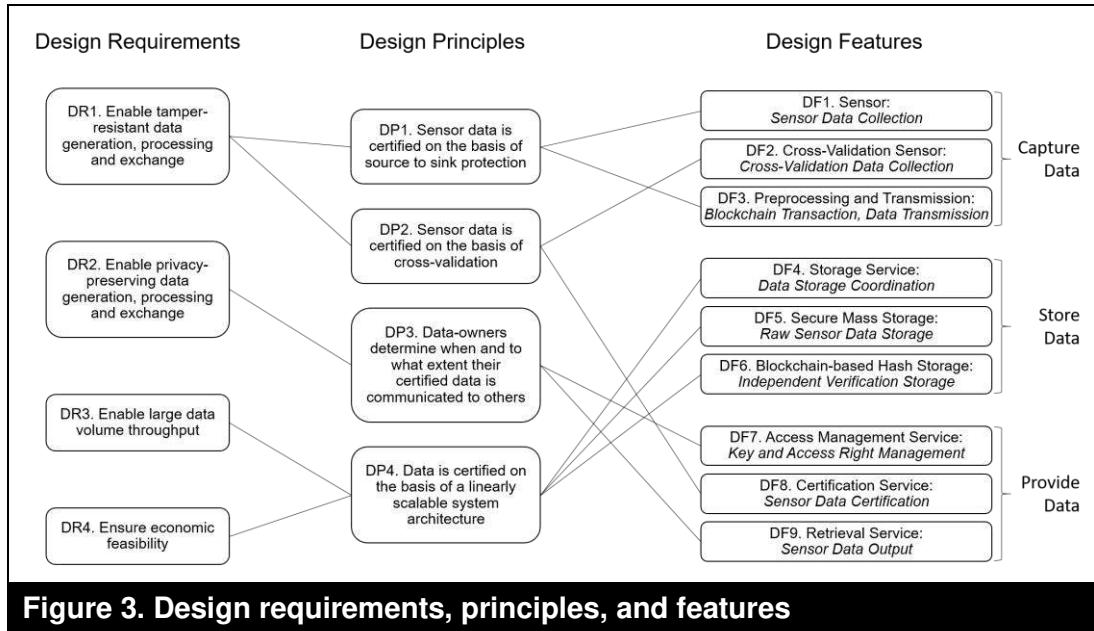
DP4: Data is certified on the basis of a linearly scalable system architecture.

4.3. *Mapping Design Principles to Design Features*

In the last step of the conceptualization, we map the identified design principles to design features. As we elaborated above, the design features are specific artifact capabilities designed to fulfil the design principles derived previously (Meth et al., 2015). An overview of these features, including the design principles and design requirements that we derived, is shown in Figure 3. The design features that we describe build on the fundamental premises (see Section 2.4) that (1) permissionless blockchain technology is a fruitful means to address the issues of data security and privacy arising in the IoT and (2) the limitations of the existing blockchain technology, with respect to privacy, scalability, and costs, have to be addressed appropriately. In the following discussion, we introduce the design features along with three fundamental system capabilities, namely, capture data, store data, and provide data.

To implement the first design principle, that is, certify that the data was not manipulated on the way from the source to the sink, two features are needed. First, we have to collect the data (DF1) and, second, we need to preprocess the data in a way that prevents data manipulation from this point on (DF3). To achieve this, we follow existing practices (Ayoade et al., 2018; Nærland et al., 2017) and save only the hash of the data (i.e., the “digital fingerprint” of the data) in a public permissionless blockchain. We can later use this hash to check that the data, which is stored in raw

format in a traditional database, has not changed by other parties since the transaction was signed. As only changes after the signature can be detected by this approach, it is essential to choose the earliest possible point in the data pipeline to create this signature and swiftly add the transaction to a blockchain.



The second design principle of cross-validation-based certification calls for two additional design features, namely, the collection of appropriate validation data (DF2) and a certification mechanism that performs the cross-validation (DF8). In the case of car mileage data, for example, GPS data can be collected for validation purposes in addition to odometer values. The GPS data can then be used to calculate the mileage data, which can be compared to the mileage values received from the odometer sensor.

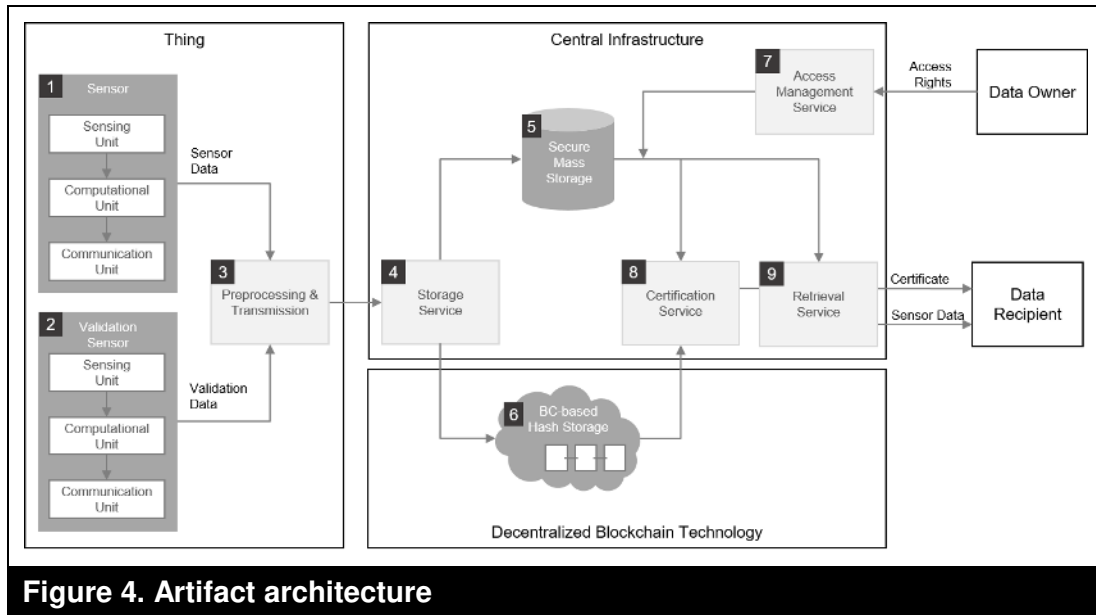
The third design principle postulates that data owners determine when and to what extent their data is communicated to others, which results in the implementation of two design features, namely, an access management service (DF7) and a data retrieval service (DF9). The access management service ensures that the raw data,

which is stored in an encrypted form in a centralized mass storage, can only be decrypted by the owner of the data. The data retrieval service is implemented in such a way that, in accordance with the access management settings, only selected parts of the whole raw data can be transferred to the data-requesting party. Hence, in the odometer example, the data owner has the possibility to choose between only sharing the last odometer value or providing the full history of odometer values, e.g., in the form of a daily, weekly, or monthly history.

The fourth design principle, requiring a linearly scalable system architecture, needs three more design features, namely, a storage service (DF4) that writes into the raw data storage (DF5) as well as into an independent verification storage (DF6). In practice, the storage service saves the encrypted raw data in a cloud storage and propagates the signed transaction with the hash to the blockchain network. By implementing these three features, a “hybrid architecture” that addresses a central challenge of public permissionless blockchain technology is realized. It is well known that certain public permissionless blockchain technologies have severe technical and economic scalability issues, so that dedicated approaches have to be applied (Beck et al., 2016; Notheisen et al., 2017; Risius & Spohrer, 2017). More specifically, hybrid architectures that build upon blockchain-based “on-chain” transactions and non-blockchain-based “off-chain” transactions are known to cope with large amounts of data while preserving the key characteristics of distributed blockchain systems (Zyskind, Nathan, & Pentland, 2015). In hybrid architectures, not all data is made available on a fully distributed blockchain. Instead, some data is stored centrally or shared only by a selected number of nodes. However, to enable trust and prevent manipulation, off-chain data has to be linked to on-chain transactions. In the case of IoT sensor data, sensor values can be stored in a central repository, and only the

digital fingerprint (hash) of one or multiple records is recorded on-chain. Thereby, the data stored in the blockchain can dramatically be reduced while still ensuring data integrity.

To summarize, Figure 4 presents a general architecture for an SDPS, including all of the design features introduced above. In practice, different instantiations of this architecture are possible. In some cases, for example, the collected data itself might not be privacy-relevant, and hence a selected retrieval thereof would not be necessary (DF7, DF9). In other cases, the validation data might be publicly available or even has to be gathered manually by inspections, which would replace the validation sensor (DF2).



The architecture highlights that a sensor is not necessarily a monolithic component and that the data, which is preprocessed and incorporated in a blockchain transaction as a hash (DF3), has already been processed through several steps as follows: it is recorded by a sensing unit, then processed by a computational unit into meaningful information, and finally communicated to a receiver outside of the sensor

through a communication unit. Hence, there are several attack vectors between the sensing unit of a sensor and the location where the blockchain transaction is actually signed. One important goal, therefore, is to build and sign the blockchain transaction (DF3) as close as possible to the sensing unit. In the future, one could imagine blockchain-enabled hardware that combine sensing and transaction management within one chip, similar to current hardware security modules. This would significantly reduce the attack vectors and ease the implementation of DP1. Storing only a hash in the blockchain supports several goals, in addition to the main objective of guaranteeing the immutability of the stored data. As opposed to storing the raw data in the blockchain, using only a hash additionally prohibits other participants from gaining useful, potentially privacy-related information, as the blockchain is public and accessible for everyone (DP3). Furthermore, the hash serves to reduce the amount of data that needs to be stored in the blockchain and therefore supports the scalability of the solution (DP4).

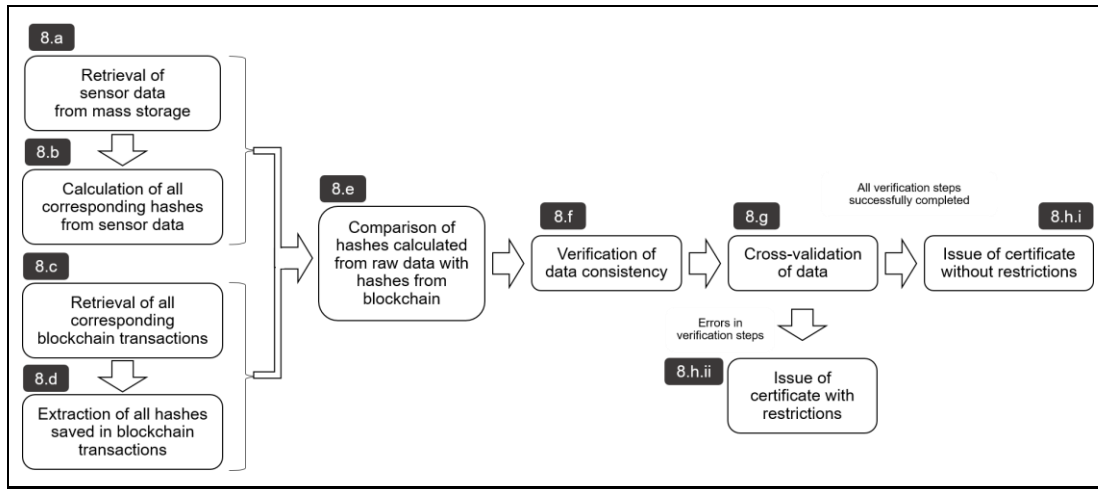


Figure 5. A detailed view of the certification process

The details of the certification mechanism and its individual steps (DF8) are outlined in Figure 5. The process is initiated by the owner of the data by granting access (DF7). The raw dataset is decrypted and sent to the unit responsible for the

certification (8.a). In the next step, the hashes of each raw data package (hashes can be calculated on the basis of single or multiple values) are calculated and stored (8.b). In parallel, for each raw data package, the corresponding transaction is looked up in the blockchain (8.c), and the saved hash is extracted (8.d). Then, the algorithm compares the hashes calculated from the raw data with those retrieved from the blockchain (8.e). A match proves that the data package in question was not changed since the signature of the corresponding blockchain transaction. Hence, the data was not manipulated on its way through the processing pipeline, and the data owner can be made accountable for the data. Any mismatches are noted and inserted as warnings in the final certificate. In the next step, the data consistency is verified (8.f). Here, the verification logic depends on given domain rules and constraints. In the case of mileage data, for example, verification can rely on the simple fact that the odometer value increases with every trip; a decrease in mileage is thus a clear indicator of an irregularity or manipulation. Typically, the more interdependent the sensor values that are recorded, the more sophisticated the tests are that can be applied. In the final step of the verification, the validation data can be leveraged (8.g). In the case of odometer fraud, the increase in the mileage of a trip should, for example, be larger or equal to the shortest distance between the GPS coordinates of the start and the end of the trip (data which is available in connected cars as of today). Finally, the certificate is issued, either without restrictions, if all verification steps were passed successfully (8.h.i), or with restrictions and a detailed report on the issues (8.h.ii).

5 Iterative Development of the Prototype

One of the core goals of design science research is to create utility for practitioners. To succeed in this task, practitioners have to understand how to apply

the abstract guidelines developed in the design science research process. As the implementation of such abstract guidelines is inherently ambiguous, scholars recommend describing the implementations of these guidelines, including the corresponding context, in detail and positioning the artifact in a natural setting, thus rendering these guidelines actionable (Baskerville, 2008; Chandra Kruse, Seidel, & Gregor, 2015; Chandra Kruse, Seidel, & Puro, 2016). Additionally, these descriptions enable researchers to establish the instantiation validity of the implementation by showing how the abstract guidelines can be linked to specific features of an artifact (Lukyanenko, Evermann, & Parsons, 2015). Hence, in the following, we present the iterative problem solving process used to design and develop our prototype CertifiCar.

The aim of our prototype is to prevent odometer fraud. Odometer fraud prevention is a relevant IoT use case in which the integrity of data is of high value and privacy is desirable. Odometer fraud, i.e., the fraudulent manipulation of a car's mileage records, is a huge problem in many countries, which is why numerous governments, for example, in Belgium, New Zealand, and the USA, have fostered the creation of systems that impede manipulation, with according legal policies (Car-Pass, 2018; Carfax, 2018; CarJam, 2018). Germany is one of the largest car markets without a centralized prevention system, and it is estimated that odometer fraud in Germany affects one third of all resold cars, leading to an annual damage of almost 6 billion euros (TÜV Rheinland, 2015). Usually, odometer fraud is committed to increase a car's value by reducing the mileage. The procedure is extremely simple and inexpensive and can be performed within minutes. Detailed step-by-step instructions are available on YouTube, and corresponding devices can be ordered online for less than 100 euros.

The existing systems that fight odometer fraud, such as Carfax (USA) and CarJam (NZ), have several substantial challenges. They are not able to detect odometer fraud reliably, have severe privacy issues, and cannot support cross-country transactions. More specifically, new records are only captured occasionally, and the interval between two records can span months or even years, giving rise to considerable fraud potential. In addition, there is no cross-validation. This makes it very difficult to detect odometer fraud. Moreover, continuous odometer fraud enabled by specific hardware manipulation devices within the car cannot be detected at all. Finally, sensitive data is stored in central databases accessible to the public, and data acquisition is limited to the country of the respective service provider. The privacy problems in the approaches of the existing systems prohibit their application in countries with strict privacy laws, such as Germany.

5.1. Iteration 1: End-to-End Processing and Initial Verification

An overview of the prototype architecture in its final state is displayed in Figure 6. In the first iteration, we implemented an initial version of the end-to-end data pipeline. This included the recording of the odometer data in the car (DF1), the processing of this data in the application (DF3, DF4), and the subsequent storing of the encrypted raw data and hashes in the private cloud storage (DF5) and the blockchain (DF6), respectively.

We chose the Ethereum blockchain because it offered the best development support and a vibrant ecosystem at the time of the development of the prototype in the beginning of 2017 (Buterin, 2013). As a proof of principle, we used the public Ethereum blockchain for a set of transactions. In addition to individual sample transactions on the Ethereum MainNet, we set up a private instance of the Ethereum blockchain exclusively for the prototype, which was operated and used. The system

has been in operation for over a year with only short interruptions. Additionally, a first version of the verification process (DF8) was implemented. This ensured that all data points were protected by a corresponding hash in the blockchain and were not manipulated (DF8.e). The verification process investigated if the mileage did not decrease at any point in time (DF8.f). To interact with the system seamlessly, a web-based user interface was added³. We chose to record the data points on the trip level to ensure reasonable transaction costs while guaranteeing a resolution high enough to detect fraud reliably.

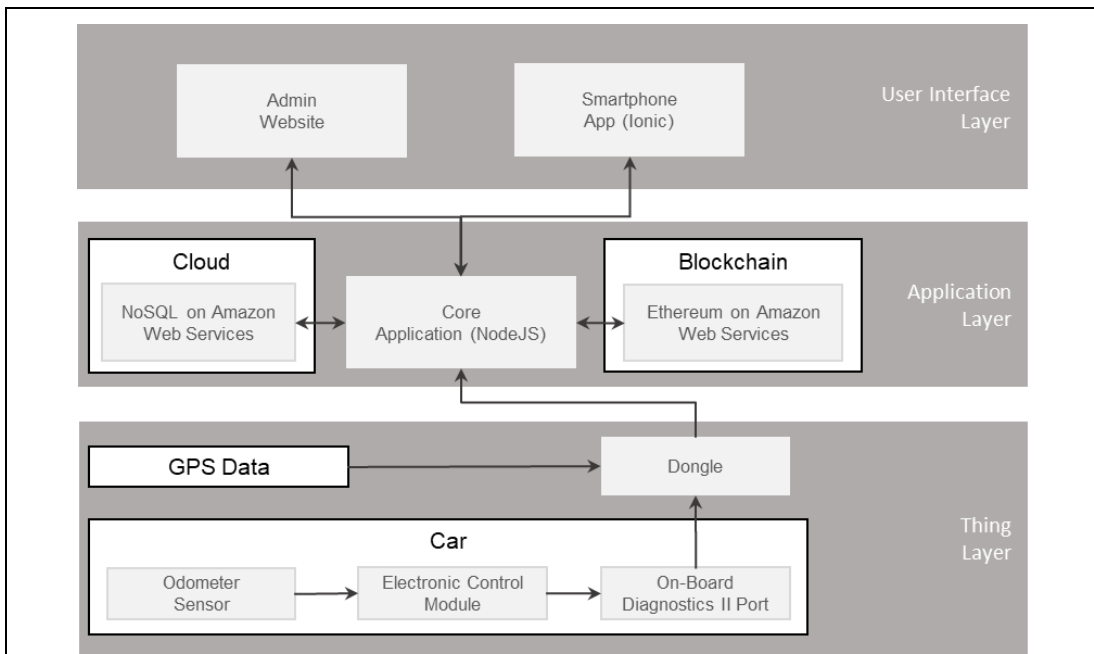


Figure 6. Prototype architecture

Every iteration of the creative and heuristic design as a search process should generate a representation of the artifact that is being demonstrated and evaluated (Hevner et al., 2004; Peffers et al., 2007). We tested this iteration with five cars that were driven daily for several hours for two weeks. This showed that the prototype was running without any major issues.

³ For details, please see Appendix, Figure A-1

A sample Ethereum transaction of the prototype at this stage, written to the public blockchain, is shown in Figure 7 and Figure 8. Figure 7 shows the view of the transaction in the JavaScript command line interface of the geth client, the official Go implementation of the Ethereum protocol. Please note that the hash of the sensor data is labeled “input”, while the value depicted as “hash” is the hash value of the overall blockchain transaction. In the example at hand, the corresponding transaction is the first transaction (“transactionIndex”) in the 2,806,957th block (“blockNumber”). To prove that the transaction was submitted to the Ethereum MainNet, Figure 8 shows a view from etherscan.io, where one can recognize the stored hash (“Input Data”).

```

> web3.eth.getTransactionFromBlock(2806957, 1)
{
  blockHash: "0x10a31a62aead015458c0afb00e341a155b1f68785d10dd249d7ed9bda3a8d231",
  blockNumber: 2806957,
  from: "0x962e55bf08e57382b047c9e47765cdf89125261a",
  gas: 90000,
  gasPrice: 20000000000,
  hash: "0x1f086129e9ee3242ce661690dabc7ed312dff6266c331b9bc4455ea19b366e13",
  input: "0x003700650065003600330031006500660066006300640034006300340035006100350061003300037006100660066003500340034006300300034006300330003400630035003800620033003100650031003500640066003200320061003300320030003700320033003900380038006300330066006200360039003200310034003900390064",
  nonce: 21,
  r: "0x227f73b1afaf96697abbd74e465b43d3ee63e00bfc8a3da51992cc42a51705f",
  s: "0x6cbfa5161d650f81159b28f8eb3bdeb0bb155783bbd8bfa5724c68046bfed9b",
  to: "0xdf65b9e6d1228aafa9c4bb0e4d914311eb544a54",
  transactionIndex: 1,
  v: "0x26",
  value: 1000000000000000
}
    
```

Figure 7. Sample Ethereum transaction retrieved with a local instance of the geth client

For this Ethereum transaction, Figure 9 shows the verification process at that point in time and how it links to the respective design feature (DF8) and its sub-processes. Note that at this stage of the prototype, the cross-validation (DF8.g, see Figure 5) was not yet implemented.

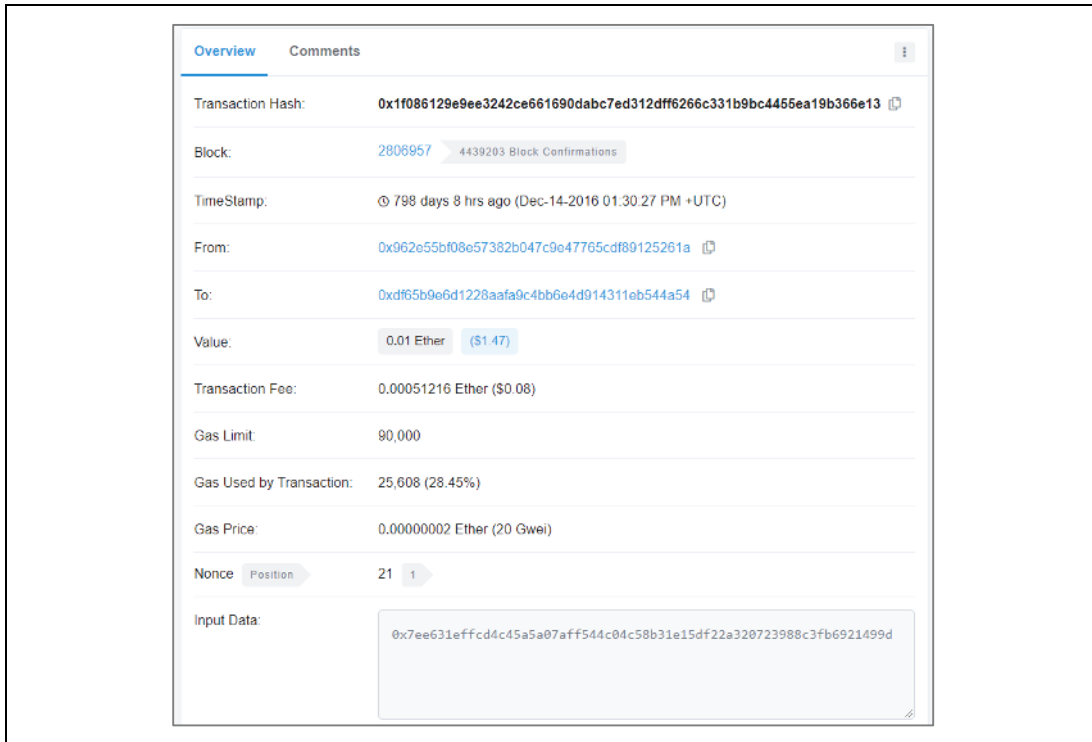


Figure 8. Transaction of Figure 7 on etherscan.io with the Ether price at the time of the transaction




Figure 9. First implementation of the verification process (hashes are presented in ASCII format to increase readability)

Finally, we conducted a series of workshops and semi-structured interviews with automotive and information technology (IT) industry experts. This revealed that a special case of so-called “continuous odometer fraud”, previously unknown to us, was impossible to detect with the existing system. In the case of continuous odometer fraud, the mileage of the car is continuously recorded at a lower-than-actual rate, i.e., only a certain percentage of the mileage actually driven is added to the odometer, for example, 80%. This is achieved by installing additional hardware, a so-called “CAN filter”, in the car. Such odometer filters are readily available on the Internet, for example, on eBay, for less than 50 USD. Our solution after this first iteration, however, was only focusing on odometer mileage reduction to uncover potential fraud; an increase at a lower rate could not be related to fraudulent behavior. We addressed this issue in our next iteration by adding a cross-validation feature (DF2, DF8.g).

5.2. Iteration 2: Cross-Validation and Scalability

To address the problem of continuous odometer fraud, we leveraged GPS data (start and end coordinates of a trip) in the second iteration (DF2). To use the GPS data to enhance fraud prevention, the verification process needed a substantial update. In addition to verifying the increase of the odometer value, we also checked that the trip distance calculated on the basis of the odometer mileage exceeded the distance between the GPS points from the start and the end of the trip (DF8.g).

Furthermore, we addressed the scalability of the solution in this iteration (DP4). The internal processes of the application were optimized and structured by several queues to enable the fault-tolerant processing of data from a larger fleet of cars⁴. For the evaluation of this iteration, 100 cars were deployed in a field test. These cars were

⁴ For details, please see Appendix, Figure A-2

supplied by one of the leading German car manufacturers, whom we contacted for the evaluation of the initial iteration of the prototype. Supplying the whole fleet with dongles would have been very costly and out of the scope of this study, which is why, as of this iteration, the data was routed over the internal backend of the car manufacturer, where it was sent directly by the connected cars used for the field test.

This version of the prototype was tested over twelve weeks with 100 cars that were used on a daily basis. We also conducted another series of workshops and interviews. The test revealed that the processing and verification of the enriched dataset, including the GPS values, worked as intended. By manipulating the sensor data from the administrator interface, the usage of odometer filters was simulated. The cross-validation procedure thus reliably detected the simulated continuous odometer fraud. Even minor manipulations (e.g., a continuous reduction in the mileage by 10%) could be consistently identified after 15 trips.

The evaluation also revealed stability problems with the underlying infrastructure, specifically in respect to the Ethereum integration. Issues such as clients losing connection to the blockchain network or cloud servers running out of storage could easily be fixed. Other problems were more severe. Especially delicate was the fact that the Ethereum client responded to the sending of a signed transaction to the blockchain network with a valid transaction hash, even if the transaction itself had not necessarily been successfully processed by the network. Hence, additional logic was necessary to assure that a transaction had successfully been processed by the blockchain network. These issues were addressed in a third iteration, leading to the final prototype.

5.3. Iteration 3: Stability and Usability

In the third iteration, we addressed the stability problems observed in the evaluation step of iteration two and implemented a smartphone app for end users to interact with the CertifiCar system. We improved the stability of the system with several measures. First, we started relevant processes via a daemon to ensure their uptime and introduced an additional queue⁵ to check that a transaction had been successfully inserted in the blockchain. Additionally, we set up an infrastructure monitoring tool (Nagios) to reduce the response time to system problems.

To improve the usability, we provided a smartphone app, which is shown in Figure 10. It includes an overview screen and a history of the driven distance, as well as a screen that allows the creation of a certificate that can be sent to a receiving party via email. The smartphone application, and in particular the process of creating a certificate, was tested by a focus group of 16 people. The feedback led to a simpler design, specifically with respect to data sharing and certification. The data owner has the option to share only the current odometer value, for example, with a potential buyer. Hence, no detailed car-usage data is revealed. However, data owners might want to share historic data to increase trust and ultimately the sales price. Therefore, they can also share the odometer history on a monthly, weekly, or daily basis.

Overall, the final iteration resulted in a prototype with increased stability and an intuitive interaction possibility through the smartphone application. The robustness and the reception by users was encouraging, resulting in the final clearance for a larger field test, which is now ongoing.

⁵ For details, please see Appendix, Figure A-2

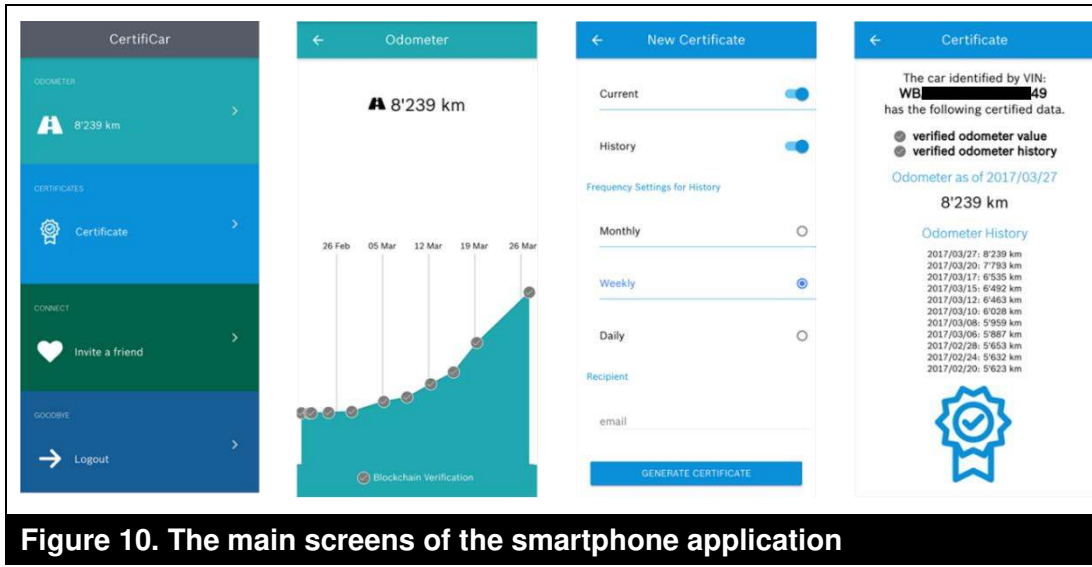


Figure 10. The main screens of the smartphone application

5.4. *Prototype Evaluation*

We have continuously evaluated the implementation against practical results from an accompanying field test that consisted eventually of 100 cars and expert feedback from workshops and semi-structured interviews. Overall, we held six workshops between December 2016 and September 2017, each comprising two to six experts and three to four researchers (in total, 22 evaluators participated) and lasting between three to five hours. Additionally, we interviewed sixteen experts between January and August 2017 for 45 minutes to 1 hour each. We prompted the participants for specific feedback and related that to the corresponding design decisions to adapt the design principles and features. Among the experts in the workshops and interviews were engineers from a German car manufacturer, a data protection law expert from a German car manufacturer, specialists from a German technical certification provider, an online car sales platform CEO, and engineers from a German car supplier. An overview of the evaluation is shown in Table 2.

Table 2: Overview and summary of the prototype evaluation			
	Iteration 1	Iteration 2	Iteration 3
Core Developments and Improvements	<ul style="list-style-type: none"> • Initial end-to-end prototype • First implementation of verification process to detect odometer fraud 	<ul style="list-style-type: none"> • GPS-based cross-validation to address continuous odometer fraud • Queue management for scalability and reliability 	<ul style="list-style-type: none"> • Transaction queue to assure reliable blockchain transaction processing • Smartphone app for end users
Evaluation	<ul style="list-style-type: none"> • Focus on fraud detection, scalability, and reliability • Field test with 5 cars • 5 interviews and 2 workshops (lasting 3.5 and 4.5 hours and with 2 and 5 participants, respectively) 	<ul style="list-style-type: none"> • Focus on continuous odometer fraud detection, scalability, and reliability • Field test with 100 cars and simulated continuous odometer fraud • 7 interviews and 2 workshops (4 hours each, 3 people each) 	<ul style="list-style-type: none"> • Focus on smartphone application, particularly the process of certificate creation, and system reliability • Field test with 100 cars and focus group of 16 people • 4 interviews and 2 workshops (lasting 4 and 3 hours and with 6 and 3 people, respectively)
Core Results	<ul style="list-style-type: none"> • Initial verification procedure detects odometer reductions but not continuous odometer fraud • Limited scalability and fault tolerance 	<ul style="list-style-type: none"> • Cross-validation procedure reliably detects continuous odometer fraud • Successful blockchain transaction processing is not guaranteed, occasional loss of transactions 	<ul style="list-style-type: none"> • Stable prototype • Well-accepted smartphone app • Clearance for larger field test

6 Ex-Post Evaluation of the Design

The scope of the evaluation in DSR reaches beyond the question of whether an artifact works and fulfils the design requirements proposed. Additionally, DSR should thoroughly explore how and why an artifact works (Pries-Heje et al., 2008). Therefore, we conducted an additional ex-post evaluation (Beck et al., 2013; Pries-Heje et al.,

2008) to address these questions and to investigate to what extent the proposed guidelines are actionable and help to create a solution for the underlying problem class. In addition to the case of odometer fraud that we have investigated in detail for the development of our artifact, we included two other use cases for this ex-post evaluation. Thereby, we want to go beyond a single prototype evaluation and gear the evaluation more towards the overall problem class. As suggested by Beck et al. (2013), to reach a higher level of abstraction, we collected additional slices of data and discussed the viability of our proposed guidelines with a purposive group of domain experts. Therefore, we selected two additional cases in which IoT sensor data needs to be protected and that are discussed extensively as fruitful blockchain use cases, namely, supply chain management (Pilkington, 2016; Tian, 2016; Underwood, 2016) and energy microgrids (Imbault, Swiatek, De Beaufort, & Plana, 2017; Mengelkamp et al., 2018; Münsing, Mather, & Moura, 2017). The first case relates to cold chains, where it must be ensured that the temperature along the supply chain

Table 3: Ex-post evaluation interview participants			
Participant name	Role	Industry	Case
BC Dev, Manufacturing	Blockchain developer	Manufacturing and engineering	Supply chain
PM BC, Manufacturing	Project manager blockchain	Manufacturing and engineering	Automotive
BC Sol Arch, Energy	Blockchain solution architect	Energy	Energy
PM BC, Energy	Project manager blockchain	Energy	Energy
BC Dev, Software	Blockchain developer	Software consulting	Supply chain
Sol Arch, Automotive	Solution architect	Automotive	Automotive
PM Innovation, Automotive	Project manager innovation	Automotive	Supply chain
PM Innovation, Manufacturing	Project manager innovation	Manufacturing and engineering	Energy
Certification Expert, Inspection	Certification expert	Inspection and product certification	Automotive

stays within a certain range (Modum, 2018). In the second case, we considered an energy microgrid with participating consumers and prosumers, where it is essential to protect the readings of smart meters for a well-functioning peer-to-peer market (Exergy, 2017b, 2017a).

As we had already developed a real-world instantiation of an artifact for the odometer fraud case, we were already in contact with several experts from the automotive and IT certification industries. These relations helped us to recruit a purposive sample of interview participants (Miles & Huberman, 1994; Robinson, 2014) with expertise in the IoT domain, including dedicated experts on subjects such as IoT sensor systems, blockchain technology, odometer fraud, supply chain management, and energy microgrids (see Table 3). We conducted a total of nine interviews (three per use case: odometer, cold chain, and microgrid), each of which lasted 45 to 70 minutes (four face-to-face and five via phone). The conversations were semi-structured, fully recorded, transcribed, and analyzed. We opted for the format of semi-structured interviews to decrease the risk of biasing participants to concrete answers and to allow a more free way of expression, especially as the interviewees often had more expertise in the specific subject matter than the interviewer (Myers & Newman, 2007; Wengraf, 2001). Below, we provide evidence from the transcripts of the nine interviews regarding the efficacy of the proposed design principles and corresponding features to address the design requirements defining our problem class.

DP1: Sensor data is certified on the basis of source to sink protection.

A majority of the participants deemed DP1 to be of “utmost importance” (*PM BC, Manufacturing; BC Sol Arch, Energy; Certification Expert, Inspection*) or even “the most important” (*BC Dev, Manufacturing; PM Innovation, Automotive*), independent of the use case. One participant mentioned that DP1’s “implementation applies to all

use cases” and that DP1 is a “necessary basis to guarantee the validity of sensor data” (*BC Dev, Manufacturing*). However, interviewees agreed (*BC Dev, Manufacturing; BC Sol Arch, Energy; PM BC, Energy; BC Dev, Software Consulting; PM Innovation, Automotive*) that “in practice, it is difficult to comply 100% with [the DP]” (*PM BC, Manufacturing*), especially “in the fragmented ecosystem of the IoT, where sensors are built by one company, deployed by another, and a third runs a service on top of that infrastructure” (*PM BC, Manufacturing*). With regard to future developments, it was articulated that the implementation of DP1 could become easier, for example “if sensors can communicate directly with the blockchain” (*PM Innovation, Automotive*) or at least “sign transactions” (*BC Dev, Manufacturing*). One participant additionally noted that “[for a scalable solution] a sensor that is able to sign transactions would be sufficient” (*BC Dev, Manufacturing*), although a sensor that is able to directly communicate (bidirectional) with the blockchain “would open fascinating new possibilities, as it could directly interact with smart contracts and, instead of a one-way communication, a dialogue could be realized” (*BC Dev, Manufacturing*), which would allow the sensor to also receive coins and instructions from the blockchain.

Many interviewees expect blockchain-enabled hardware (*BC Dev, Manufacturing; PM BC, Manufacturing; PM Innovation, Automotive*), for example “sensors similar to hardware security modules” (*PM BC, Manufacturing*) that ease the implementation of DP1 to be available in the future. However, one participant noted that “currently, the software specifications of blockchains [e.g., of signature algorithms] are still evolving [and therefore] the development of sensor ASICs still has to wait” (*BC Dev, Manufacturing*). The advantage of application-specific integrated circuits (ASICs) would rather lie in a “more energy efficient processing than in

increased speed” (*BC Dev, Manufacturing*). While the implementation specifics are expected to change, the “basic concept of blockchains as a record of an immutable shared truth [is not]” and, therefore, the usage of blockchain transactions as in DF3 to fulfil DP1 “should continue to make sense” (*BC Sol Arch, Energy*).

While generally source to sink protection through a blockchain transaction was appreciated as a sound measure to hinder data tampering, several participants agreed (*BC Dev, Manufacturing; Certification Expert, Inspection*) that there “will probably never be a way to ensure a completely tamper-proof solution” (*PM Innovation, Automotive*). For example, “one could simply manipulate the surrounding of the sensor – in the case of a cold chain, for example, by putting a cooling element or ice on top of the temperature sensor” (*BC Dev, Manufacturing*). One participant concluded that “while it makes sense to aim for a tamper-proof solution, it is sufficient to ensure tamper resistance that is strong enough to make it economically unprofitable to commit fraud, similar to proof of work [a blockchain mining mechanism]” (*PM Innovation, Automotive*), which is in line with DR1. Overall, DP1 and its implementation (corresponding design features) were strongly supported by all nine interviewees, and the participants provided fruitful insights into how blockchain technology might evolve to enable DP1.

DP2: Sensor data is certified on the basis of cross-validation.

Most participants acknowledged that an implementation of DP2 would be needed, as either DP1 could not uniquely prevent all data tampering or it could not be implemented to the full extent. As such, one participant noted that “it is good that the dependence on DP1 is reduced by the introduction of DP2” (*BC Sol Arch, Energy*), and another stated that “[some kind of] cross-validation is always necessary because already the reading of the sensor could be influenced [in a manipulative way]” (*Sol*

Arch, Automotive). Relating to future developments, a participant noted that “increasing the security by implementing DP2 is probably faster and more economically viable than perfecting the implementation of DP1, possibly with future technology” (*PM Innovation, Automotive*).

Participants also noted that DP2 is “rather use case specific” (*PM BC, Manufacturing*), in contrast to DP1, and speculated that “in some cases it might be difficult to find appropriate data for cross-validation” (*BC Dev, Software Consulting*). Regarding the cold chain case, an interviewee suggested that “weather data could be combined with cooling power consumption data of the truck to detect anomalies” (*BC Dev, Software Consulting*). With respect to the microgrid case, they proposed to use data of “a transformer station supplying several houses with electricity” and “weather data in combination with power data from the installed solar panel” (*BC Sol Arch, Energy*) for cross-validation. In the case of odometer fraud, the “cross-validation could be expanded considerably with service and maintenance data”, for example, by “validating that the exchange of brake disks occurs after roughly 50,000 kilometers” (*PM Innovation, Manufacturing*). In essence, all participants supported DP2 and highlighted its context dependency as well as the interlinked nature of DP1 and DP2.

DP3: Data owners determine when and to what extent their certified data is communicated to others.

Generally, the participants stated that the privacy-preserving mechanisms introduced through DP3 are very strong. According to one interviewee, “the propagation of information is organized well in the system and occurs in a very safe way” (*BC Sol Arch, Energy*). Three participants mentioned the upcoming General Data Protection Regulation (GDPR) of the EU (European Commission, 2018) and noted that the most important parts thereof are covered in DP3 and its features (*PM*

BC, Manufacturing; PM BC, Energy; PM Innovation, Manufacturing). One participant stressed additionally that “there is also an obligation to inform the data owner about how her data will be used by the receiving party” (*PM BC, Energy*), and another stressed that “there should be a possibility to revoke the sharing of data any time after the data has been sent for the first time” (*PM Innovation, Manufacturing*).

Regarding the importance of privacy, it was noted that it is highly dependent on the gathered data in the specific case and, importantly, on the perception of the data owner towards sharing this data (*BC Dev, Manufacturing; BC Dev, Software Consulting*). For example, people are “used to sharing their electricity consumption data with their energy supplier” (*PM BC, Energy*), and in a cold chain, “a driver might not perceive the sharing of temperature data as very sensitive” (*BC Dev, Software Consulting*). Therefore, a participant argued, “it might actually be a challenge to convince users that data privacy is valuable in their case” and raised the question of “how do you want to raise awareness for that?” (*PM BC, Manufacturing*). This statement is in line with the comment of another participant that “at the moment, privacy is typically driven by regulatory decisions [in Europe] and not by customer demand”, concluding that “currently, it is often not essential for flourishing businesses [to provide privacy-preserving solutions], but it will probably become a core feature in the future” (*PM Innovation, Automotive*). In line with this last comment, one participant noted that new technology enables gathering and transmitting data at a more granular level, possibly changing users’ perceptions as follows: “If you start sharing your electricity consumption on a minute basis, instead of delivering a quarterly or annual meter reading, you might get more uncomfortable” (*PM BC, Energy*). In summary, the participants appreciated DP3 and the corresponding design features. They also

emphasized that privacy is becoming increasingly important as the technological performance and the ability to collect detailed data increases.

DP4: Data is certified on the basis of a linearly scalable system architecture.

Several participants noted that DP4 is, together with DP1, essential for any solution trying to solve the problem of data protection and certification (*BC Dev, Manufacturing; PM Innovation, Automotive*). One participant with a strong business background said that this “needs to be fulfilled right away” (*PM Innovation, Automotive*). In general, the participants noted that the scalability provided by the proposed principles and features is indeed sufficient for real-world applications like, for example, the processing of the majority of all cars in the EU. One participant noted that the “scalability properties of blockchain-based solutions strongly depend on the use case at hand and the specific implementation”, continuing that “the often-heard statement that anything involving blockchain technology does not scale and costs a lot is simply not true [...] as, for example, CertifiCar and the OpenTimestamps project reveal” (*PM Innovation, Manufacturing*).

The hybrid approach of using both decentralized and traditional infrastructures was deemed appropriate by all interviewed blockchain experts, independent of the cases discussed. “Currently, such a solution can only be built on the basis of a hybrid approach” (*PM BC, Energy*), noted one participant, while another added that “taking into account the current state of the blockchain ecosystem, this approach definitely makes sense” (*BC Dev, Manufacturing*). However, considering future developments, many participants speculated (*PM BC, Manufacturing; BC Dev, Software Consulting*) that “it might be possible to build the entire system on a decentralized infrastructure in a far future” (*BC Dev, Manufacturing*), as already hinted before, and it was also noted that already “many players are working, for example, towards decentralized

storage solutions with throughput and scalability for enterprise environments” (*BC Dev, Software Consulting*).

Regarding the question of whether a scalable protection system is better built without blockchain technology, i.e., disregarding DF6, many interviewees agreed (*PM BC, Manufacturing; Sol Arch, Automotive*) that “technically, this would be possible” (*BC Dev, Software Consulting*). However, different considerations in favor of the usage of blockchain technology were made. One participant noted that “using a blockchain to store the hashes makes sense whenever the certification happens in an environment with a multitude of parties with [partially] conflicting interests” (*PM BC, Manufacturing*). For example, in the case of odometer fraud, “the owner of the car, a potential buyer of the car, the car manufacturer, associated and independent workshops, and even different departments within a car manufacturer have different interests regarding odometer fraud” (*PM BC, Manufacturing*). Therefore, establishing a central database for all participants that is operated by just one of the involved parties is a major challenge. Uninvolved third parties can take over the responsibility to run such a system. It was also noted that “new business models based on other sensor data that is shared in a multi-party system” (*Certification Expert, Inspection*) will increase in importance. In principle, it “might be possible to find a traditional database provider [for this role]” (*BC Dev, Software Consulting*); however, it could be costly and potentially difficult to reach an agreement between all parties involved. “A blockchain provides a viable alternative in such a case, with no need to trust a third party” (*BC Dev, Software Consulting*).

In addition, the participants noted that the “overhead of the blockchain is small – really expensive are [hardware] sensors and connectivity” (*BC Dev, Manufacturing*). The blockchain “can even reduce costs”, as its security is less dependent on third-

party certification, which is costly and time-consuming (*BC Dev, Manufacturing*). This is especially important for smaller companies, which might not have the resources and processes to deploy highly secure databases. An expert in the research department of a multinational company stated that “the business side clearly does not see the need for a blockchain-based solution yet”, as they think that “a secure and trustworthy database can also be provided by the company itself and its brand name” (*Sol Arch, Automotive*). In line with that, several participants noted that when a blockchain is used, the trust question is transferred to “technology” or “engineering”, while in traditional systems, it is addressed with “brand names” and “company processes” (*BC Dev, Manufacturing; BC Sol Arch, Energy*).

An additional interesting point was made regarding the standardization potential of a solution relying on blockchain technology. An expert from the energy sector noted that individual energy suppliers “might be more willing to accept a solution as an industry standard if its cornerstone is based on blockchain technology, and this decreases the dependence on another company” (*PM BC, Energy*). In contrast, “if a solution’s core is in control of another energy supplier or technology provider, the adoption as a standard would be very difficult” (*PM BC, Energy*).

In essence, the interviewees highlight the importance of DP4 and agree that the proposed features are indeed appropriate to address this design principle. Furthermore, they provide several reasons why a blockchain-based SDPS might be superior to a traditional solution in particular situations. First and foremost, they highlight the potential of blockchain technology in cases where sensor data protection has to be assured in ecosystems with multiple parties with conflicting interests.

In summary, the nine interviews provided additional evidence of the usefulness of our proposed design. The participants reinforced the core considerations and major

design decisions of the SDPS design. In addition, the interviews revealed new insights, for example, with respect to the evolution of blockchain technology and its specific business potential. The results also correspond to the findings from the development and evaluation of our prototype. However, by building upon additional slices of data (Beck et al., 2013), they go beyond a “one instance evaluation” of the design.

7 Discussion

7.1. SDPS Design Theory

After the ex-post evaluation, we integrate our findings and formulate a design theory as summarized in Table 4. Thereby, we follow the seminal work of Gregor and Jones (2007), who laid out six fundamental components of a design theory. Finally, we discuss our findings in light of their theoretical and practical implications.

According to Gregor and Jones (2007), the first component of a design theory is its purpose and scope. The aim of our artifact is to develop a system that protects IoT sensor data generation, processing, and exchange in a privacy-preserving and efficient manner. With respect to the boundaries of the design, we want to highlight that the development of the guidelines was clearly focused on the processing of IoT sensor data and the corresponding challenges, such as big data, multistage data processing pipelines, and distributed data processing across organizational boundaries or multi-party ecosystems. This problem class covers a wide range of relevant issues, which is in stark contrast to existing studies on SDPSs (e.g., Ayoade et al., 2018; Liang et al., 2017; Machado & Fröhlich, 2018) that focus on specific solutions to very specific problems. The generalizability within our wide problem class constitutes an important foundation for our theoretical contribution.

1	Purpose and scope	The aim is to develop a system that protects IoT sensor data generation, processing, and exchange in a privacy-preserving and efficient manner.
2	Constructs	<ul style="list-style-type: none"> • Tamper resistance • Privacy • Scalability • Economic feasibility • Certification
3	Principles of form and function	Design principles (DP1-4) to support the protection of IoT sensor data and corresponding design features (DF1-9) are presented.
4	Artifact mutability	SDPSs have to be mutable, specifically with respect to the amount of data they can handle. DR2 and DR3 articulate this fundamental thought, and DP4 subsequently poses a linearly scalable system. SDPS can be used with benefit by different organizations. However, they need to be adapted particularly with respect to cross-validation. The cross-validation data and the certification procedure are highly dependent on the context.
5	Testable propositions	<ul style="list-style-type: none"> • P1: The artifact enables tamper-resistant IoT sensor data generation, processing, and exchange • P2: The artifact enables privacy-preserving IoT sensor data generation, processing, and exchange • P3: The artifact is capable of processing large amounts of IoT sensor data • P4: The positive effects of the artifact are not negated by artifact development and operation costs
6	Justificatory knowledge	Design requirements are based on the literature on IoT, security, and privacy. Design principles are derived from theory on information asymmetry, privacy, and IS success. Design features build upon blockchain literature.

The second component that Gregor and Jones (2007) depict is constructs, which represent core entities of interest in the design. The core constructs we propose are tamper resistance, privacy, scalability, and economic feasibility, which are reflected in our design requirements. These constructs capture the impact of an SDPS and may therefore serve as dependent variables in efforts to investigate SDPS success. In addition, the theory on information asymmetry (Akerlof, 1970) suggests that certification is a core concept and means to overcome information deficits and avoid opportunistic behavior, such as intentional data manipulation. We build upon these

insights and base our design on certification. Therefore, certification is a fundamental, independent construct of our work.

Regarding the third component of a design theory, we present principles of form and function that may serve as a blueprint for the construction of IoT sensor data protection systems. To this end, we identify the SDPS design requirements (DR1-4), derive design principles (DP1-4) to support the protection of the IoT sensor data and depict corresponding design features (DF1-9) (see Figure 3). The requirements, principles, and features constitute actionable guidelines, which highlights a core difference between our work and the extant research. Thereby, we reflect the various calls in the IS literature to support the development of implementable tools to increase security and privacy, especially in the IoT (Bélanger & Crossler, 2011; Lee et al., 2018; Medaglia & Serbanati, 2010; Pavlou, 2011).

To account for the special nature of IS artifacts, Gregor and Jones (2007) call for explicitly addressing the mutable nature of these artifacts as a fourth component. In the case of SDPSs, we reflected the importance of mutability specifically with respect to the amount of data they can handle. DR2 and DR3 articulate this fundamental thought, and DP4 subsequently poses a linearly scalable system. However, the design that we derived is not universally applicable, nor is it “one-size-fits-all”. While SDPSs can be used with benefit by different organizations, they need to be adapted particularly with respect to cross-validation. The cross-validation data and the certification procedure are highly dependent on the context, as the development of the instantiation that we presented clearly indicates.

The fifth component of a design theory comprises testable propositions. These propositions might be presented as “if a system or method that follows certain principles is instantiated, then it will work, or it will be better in some way than other

systems or methods”. Following this argumentation, we can deduce propositions from the presented design requirements. The design requirements disentangle the “it will work, or it will be better” into specific, contextualized needs that must be addressed by the artifact. Propositions postulate that these needs have been successfully addressed and serve as a basis for assessing the impact of the artifact. Applying this rationale to DR1-4, we deduce the following four propositions: the artifact enables tamper-resistant IoT sensor data generation, processing, and exchange (P1). The artifact enables privacy-preserving IoT sensor data generation, processing, and exchange (P2). The artifact is capable of processing large amounts of IoT sensor data (P3). The positive effects of the artifact are not negated by the artifact development and operation costs (P4). These propositions might be helpful in developing test cases for future instantiations.

Finally, Gregor and Jones (2007) encourage scholars to provide the justificatory knowledge of their design. We base our design requirements on insights from the literature on IoT, security, and privacy (see Section 4.1). The design principles are mainly derived from theory on information asymmetry, privacy, and IS success (see Section 4.2). Ultimately, the design features build primarily upon the blockchain literature (see Section 4.3). This theoretical grounding enabled us, in close interplay with insights from practice, to derive a set of purposive guidelines for the design of SDPSs in the form of DRs, DPs, and DFs. Gregor and Jones (2007) emphasize the importance of explanatory theory as a “linking mechanism for a number, or all, of the other aspects of the design theory” (p. 327). We reflect this role of explanatory theory by explicitly deriving design principles that serve as a link between design requirements and design features. This thorough conceptualization of the problem is a key distinction from previous literature (e.g., Ayoade et al., 2018; Liang et al., 2017;

Machado & Fröhlich, 2018), and it facilitates the generalizability of our findings, which enables our theoretical contribution.

7.2. Design Implications

Our research has important design implications for SDPSs that address IoT-related security and privacy challenges (Ayoade et al., 2018; Crossler & Posey, 2017; Liang et al., 2017), specifically with respect to the value proposition of blockchain technology. Blockchain-based SDPSs inherit core characteristics of blockchain technology (Notheisen et al., 2017) and therefore are particularly useful in certain scenarios (see Table 5). While SDPSs are used to protect simple data pipelines, for example, to secure data transfer from sensors to one single intra-organizational system, they are also leveraged in the case of multi-stage data pipelines that cross organizational boundaries and involve a potentially large ecosystem of players, as our prototype case reveals. In the latter case, blockchain-based SDPSs are particularly valuable because they can protect sensor data even in large ecosystems with conflicting interests through the use of a shared, immutable ledger. In addition, a blockchain-based SDPS is a decentralized system. Hence, the involved parties are peers, and no single party controls the overall system (Beck et al., 2018). As our ex-post evaluation reveals, such a system is often perceived as “neutral” and might be accepted as an industry standard much faster than a centralized system. Finally, important security and protection technology, such as public-key cryptography, is already built into blockchain technology (Buterin, 2013; Noyen, Volland, Wörner, & Fleisch, 2014). Additionally, the infrastructure to use these protocols is readily provided by a decentralized set of actors (e.g., miners), who are typically incentivized through the economics of cryptocurrencies. Essentially, blockchain technology offers a ready-to-use set of well-defined security protocols. For smaller companies, in

particular, that do not have cryptography specialists or corresponding technology available, blockchain-based SDPSs offer the opportunity to leverage state-of-the art security technology that is usually license-free and often designed for rapid adoption.

Table 5: Blockchain-based SDPS usage implications		
Blockchain characteristic	Related advantages	SDPS usage implications
Shared, immutable ledger	<ul style="list-style-type: none"> • Blockchain integrates the advantages of distributed databases and crypto technology • Well-managed data redundancy across different parties • Secure data processing that fosters data integrity 	“using a blockchain to store the hashes makes sense whenever the certification happens in an environment with a <i>multitude of parties with [partially] conflicting interests</i> ” (PM BC, Manufacturing)
Decentralized system	<ul style="list-style-type: none"> • No central authority • All parties are peers with the same rights • No single party controls the overall system 	“[members of an ecosystem] might be more willing to <i>accept a solution as an industry standard</i> if its cornerstone is based on blockchain technology and this <i>decreases the dependence on another [single] company</i> ” (PM BC, Energy)
Ready-to-use set of well-defined security protocols and infrastructure	<ul style="list-style-type: none"> • Private and public key cryptography stack built into blockchain • Infrastructure readily provided by a decentralized set of actors incentivized through economics of cryptocurrencies • Security does not rely on third-party certification, which is costly and time-consuming • Even smaller companies with no dedicated cyber-security or cryptography specialists can leverage state-of-the art security technology 	“overhead of the blockchain is small – really expensive are [hardware] sensors and connectivity” (BC Dev, Manufacturing), “the <i>blockchain can reduce costs</i> ” (BC Dev, Manufacturing)

However, as our design theory reveals, blockchain-based SDPSs have to be carefully designed. Blockchain technology is not a universal solution that addresses the derived design requirements out of the box. The fundamental design implications must be considered to address the derived design requirements (see Table 6). With

respect to DP1 (sensor data certified on the basis of source to sink protection), it is important to note that, as of today, sensors cannot communicate directly with the blockchain. Therefore, the data must be protected as early as possible in the processing chain by building and signing blockchain transactions as close as possible to the sensing unit. In the future, blockchain-enabled sensors could drastically simplify this and might allow for signing within the sensing unit. In addition, DP2 (sensor data certified on the basis of cross-validation) has to be carefully addressed. More specifically, system designers have to realize that blockchain technology generally cannot assure “tamper-proof” processes, and the additional cross-validation of the sensor data is necessary to enable effective tamper resistance. Thereby, a nondetection risk of fraud remains. With respect to DP3 (data owners determine when and to what extent their data is communicated to others) it should be noted that a blockchain is not a universal remedy that can guarantee privacy (Conti, Kumar, Lal, & Ruj, 2018; Fabian et al., 2016; Goldfeder, Kalodner, Reisman, & Narayanan, 2018; Kumar, Fischer, Tople, & Saxena, 2017). In the context of sensor data sharing specifically, privacy mechanisms have to be implemented on top of the blockchain in the form of an access management service. In addition, by relying on a hybrid blockchain approach, there must be assurances that the sensor data itself is not stored in a public permissionless blockchain and that data integrity can be maintained. Finally, regarding DP4 (data certified on the basis of a linearly scalable system architecture), specific blockchain architectures have to be implemented. With the current state of technology, hybrid blockchain architectures (Ayoade et al., 2018; Zyskind et al., 2015) are necessary to enable scaling. Therefore, viable systems store sensor values in a central repository, and only the digital fingerprint (hash) of the sensor values is recorded on the blockchain.

Table 6: Design implications for blockchain-based SDPS	
DP1 Sensor data is certified on the basis of source to sink protection	
Prototype design & eval.	<ul style="list-style-type: none"> • <i>Data must be protected as early as possible in the processing chain</i> • In the prototype, we collected odometer data and preprocessed it as soon as possible in a way that data manipulation from that point on was prevented, and we built and signed the blockchain transaction as close as possible to the odometer sensing unit • However, in the prototype, we could only do this rather late in the processing chain, as a blockchain cannot be directly integrated into the odometer sensor
Ex-post evaluation	<ul style="list-style-type: none"> • “[source to sink protection] is a necessary basis to guarantee the validity of sensor data” (<i>BC Dev, Manufacturing</i>) • “in practice, it is difficult to comply 100% with [source to sink protection]”, especially “in the fragmented ecosystem of the IoT” (<i>PM BC, Manufacturing</i>) • Implementation of DP1 could become easier, for example “if sensors can communicate directly with the blockchain” (<i>PM Innovation, Automotive</i>) or at least “sign transactions” (<i>BC Dev, Manufacturing</i>)
DP2 Sensor data is certified on the basis of cross-validation	
Prototype design & eval.	<ul style="list-style-type: none"> • <i>Blockchain technology cannot assure “tamper-proof” processes per se, so additional cross-validation is necessary to enable effective tamper resistance, and a nondetection risk of fraud remains</i> • Initial prototype verification procedure detects odometer reductions but not continuous odometer fraud • Prototype cross-validation procedure finally reliably detects continuous odometer fraud
Ex-post evaluation	<ul style="list-style-type: none"> • “[blockchain] will probably never be a way to ensure a completely tamper-proof solution” (<i>PM Innovation, Automotive</i>) • “[some kind of] cross-validation is always necessary because already the reading of the sensor could be influenced [in a manipulative way]” (<i>Sol Arch, Automotive</i>) • DP2 is “rather use case specific” (<i>PM BC, Manufacturing</i>)
DP3 Data owners determine when and to what extent their data is communicated to others	
Prototype design & eval.	<ul style="list-style-type: none"> • <i>Blockchain technology cannot assure data privacy per se, so privacy must be implemented on top of the blockchain in the form of an access management service</i> • Feedback of 16 prototype users that fine-grained sharing mechanisms have to be implemented • Clearance of app for large field test that included user feedback & legal compliance check
Ex-post evaluation	<ul style="list-style-type: none"> • “there should be a possibility to revoke the sharing of data any time” (<i>PM Innovation, Manufacturing</i>) • “the propagation of information is organized well [in the proposed design] and occurs in a very safe way” (<i>BC Sol Arch, Energy</i>)
DP4 Data is certified on the basis of a linearly scalable system architecture	
Prototype design & eval.	<ul style="list-style-type: none"> • <i>Hybrid blockchain architecture necessary to enable scaling</i> • Odometer sensor values are stored in a central repository, and only the digital fingerprint (hash) of the records is recorded on-chain • System for 100 cars was deployed on the basis of two low-performance standard Amazon EC2 instances, and there were no performance issues during the evaluation
Ex-post evaluation	<ul style="list-style-type: none"> • “scalability properties of blockchain-based solutions strongly depend on the use case at hand and the specific implementation” (<i>PM Innovation, Manufacturing</i>) • “the often-heard statement that anything involving blockchain technology does not scale and costs a lot is simply not true” (<i>PM Innovation, Manufacturing</i>) • “Currently, such a solution can only be built on the basis of a hybrid approach” (<i>PM BC, Energy</i>)

7.3. Theoretical and Practical Contributions

In summary, the proposed SDPS design theory is the key theoretical contribution of our work. We synthesize our design into a conceptual solution that addresses a whole problem class. Notably, the codification and abstraction of our design, including the design requirements, design principles, and design features, enables generalizability beyond a particular problem. The provision of actionable guidelines based on such a thorough conceptualization is, to the best of our knowledge, a novel contribution, which was specifically called for (e.g., Bélanger & Crossler, 2011). Thereby, we add to the literature on IoT and IoT-related security and privacy challenges, as well as to the literature on blockchain technology.

More specifically, our investigation of the problem class confirms and conceptualizes earlier evidence from the literature (Aggarwal et al., 2013; Lowry et al., 2017) that the distributed, multilayered nature of IoT systems, as well as IoT ecosystems with multiple parties and potentially diverging interests, introduces very specific and particularly serious challenges. The derived design requirements can serve as a basis for future research, for example, investigating how their fulfillment affects the adoption of IoT IS. Furthermore, we base the design principles, in particular, on the theory of information asymmetry, which has been used before as a fruitful basis in the design of IS that enables the reliable exchange of data (e.g., Notheisen et al., 2017). In contrast to the existing SDPS-related literature, we specifically focus on certification as a well-known means of overcoming information asymmetries. As such, we leverage deep insights from the existing body of knowledge on information asymmetries (Bond, 1982; Genesove, 1993; Spence, 1976), and certification in particular (Akerlof, 1970; Albersmeier et al., 2009), which we strongly believe represents a useful basis for other design research in the realm of SDPSs.

Finally, we discuss the design features and the design implications of our research on the usage of blockchain technology in detail. Notably, we shed light on both the advantages as well as the potential problems of using a blockchain for SDPSs. We elaborate how the proposed design can address the widely discussed shortcomings of blockchains, such as scalability and privacy. We do this by building upon the existing research on hybrid blockchain architectures (Ayoade et al., 2018; Zyskind et al., 2015) and thereby encourage design researchers to specifically reflect the latest developments in this domain.

With regard to practical contributions, we first of all provide a blueprint that guides the development of SDPSs. Furthermore, we address emerging blockchain concerns that more and more practitioners share, namely, blockchains have no scalability, they induce high costs, and they cannot assure privacy. Our design – and more specifically the prototype – reveals that these concerns can be addressed with existing technology. This might inspire practitioners to overcome their concerns and start leveraging blockchain technology for their enterprises. In addition, in line with the existing research (Beck et al., 2016; Christidis & Devetsikiotis, 2016), our evaluation reveals where the use of blockchains might be particularly helpful in practice. Ecosystems with a multitude of parties with potentially conflicting interests often rely on an intermediary to ensure reliable data exchange and trust. In these cases, blockchain technology might serve as such an intermediary. Additionally, blockchain-based solutions might facilitate the establishment of industry standards. Finally, in light of ever-increasing regulation, blockchain-based solutions might serve as a cost-efficient complement to third-party certification. Smaller companies, in particular, might benefit from the ready-to-use security protocols and corresponding infrastructure that the blockchain provides. In the realm of IoT, however, physical

devices must be blockchain-enabled. As of today, the data pipeline too often remains unprotected directly after the sensing unit of such devices.

8 Conclusion

The study at hand uses a design science research approach to propose a design theory for a sensor data protection system (SDPS). More specifically, we derive design requirements, design principles, and design features for a blockchain-based SDPS. In addition, we design and develop an instantiation of an SDPS (CertifiCar) on the basis of three iterative cycles. Our prototype prevents the fraudulent manipulation of car mileage data. Finally, we provide an ex-post evaluation of our design theory considering two additional use cases in the realms of pharmaceutical supply chains (Modum, 2018) and energy microgrids (Mengelkamp et al., 2018). The findings of our evaluation suggest that the proposed design ensures the tamper-resistant gathering, processing, and exchange of IoT sensor data in a privacy-preserving, scalable, and efficient manner.

The results of this study should be assessed in light of its limitations. We derive design principles on the basis of specific theoretical lenses. Building upon an alternative selection of theoretical lenses, we might have identified different or additional design requirements and principles (see Meth et al., 2015). However, the chosen theories are well accepted and undisputed and represent a reliable and stable basis for analysis. In addition, our evaluation confirms that our design principles are concise and independent of current technology and upcoming technology developments, as well as applicable to the chosen problem class across different use cases. A second limitation refers to the design features that are grounded in the capabilities of today's blockchain technology. Blockchain technology is in an early

stage of development (Beck et al., 2017), and, in particular, new on-chain/off-chain approaches are still emerging (Ayoade et al., 2018; Machado & Fröhlich, 2018; Zyskind et al., 2015). Therefore, the proposed design features might change with future, potentially disruptive blockchain breakthroughs. However, we want to highlight the fact that we build upon the latest blockchain research at the forefront of technology, and our features reflect latest on-chain/off-chain architecture approaches that provide a viable tradeoff between security and scalability (Ayoade et al., 2018; Zyskind et al., 2015). A third limitation is related to the evaluation of our design theory. We developed and evaluated CertifiCar and investigated two additional use cases to reflect our design. While a quantitative and broader evaluation is desirable and encouraged, we want to emphasize that at this point in time, corresponding systems and domain experts are not widely available.

Beyond the aforementioned opportunities, there are many other possible extensions to our work. We contribute to an emerging literature stream that aims to advance the theoretical understanding of blockchain technology. We hope that our study serves as a fruitful basis for further research on how blockchain technology facilitates new modes of ecosystem collaboration, for example, by establishing security, privacy, and trust. More specifically, we encourage scholars to investigate and compare the various blockchain-based data protection approaches that are currently emerging with respect to their business potential (Risius & Spohrer, 2017). Finally, while there are several industry initiatives, such as the Trusted IoT Alliance, and many companies are currently developing promising use cases, we see an absence of design and theory to bridge the gap between technology and business. Blockchain technology is rapidly evolving, but its business potential still remains vague. It is not only researchers who have been too optimistic about the potential of

blockchain technology (Beck et al., 2017). In practice, blockchain technology is still overhyped, and discussions are either very technology-focused or business-driven without reflecting the actual capabilities and restrictions of the current technology. In line with Bélanger and Crossler's (2011) call for more actionable solutions, we encourage design science researchers to fill the articulated gap and link (business) problem classes to blockchain technology and corresponding applications.

References

- Abbasi, A., Sarker, S., & Chiang, R. H. L. (2016). Big data research in information systems: Toward an inclusive research agenda. *Journal of the Association for Information Systems*, 17(2), i–xxxii.
- Aggarwal, C. C., Ashish, N., & Sheth, A. (2013). The internet of things: A survey from the data-centric perspective. In C. C. Aggarwal (Ed.), *Managing and mining sensor data* (pp. 383–428). Berlin, Germany: Springer Science+Business Media.
- Akerlof, G. A. (1970). The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3), 488–500.
- Albersmeier, F., Schulze, H., Jahn, G., & Spiller, A. (2009). The reliability of third-party certification in the food chain: From checklists to risk-oriented auditing. *Food Control*, 20(10), 927–935.
- Alqassem, I., & Svetinovic, D. (2014). A taxonomy of security and privacy requirements for the Internet of Things (IoT). In *Proceedings of the 2014 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)* (pp. 1244–1248). Bandar Sunway, Malaysia.
- Anderson, C., Baskerville, R. L., & Kaul, M. (2017). Information Security Control Theory: Achieving a Sustainable Reconciliation Between Sharing and Protecting the Privacy of Information. *Journal of Management Information Systems*, 34(4), 1082–1112.
- Atzori, L., Iera, A., & Morabito, G. (2010). The Internet of Things: A survey. *Computer Networks*, 54(15), 2787–2805.
- Avital, M., Beck, R., King, J., Rossi, M., & Teigland, R. (2016). Jumping on the Blockchain Bandwagon: Lessons of the Past and Outlook to the Future. In *Proceedings of the 37th International Conference on Information Systems (ICIS)*. Dublin, Ireland.
- Ayoade, G., Karande, V., Khan, L., & Hamlen, K. (2018). Decentralized IoT Data Management Using BlockChain and Trusted Execution Environment. In *Proceedings of the 2018 IEEE International Conference on Information Reuse and Integration (IRI)* (pp. 15–22). Salt Lake City, UT.
- Baskerville, R. (2008). What design science is not. *European Journal of Information Systems*, 17(5), 441–443.
- Baskerville, R., Kaul, M., & Storey, V. (2015). Genres of Inquiry in Design-Science Research: Justification and Evaluation of Knowledge Production. *MIS Quarterly*, 39(3), 541–564.
- Baskerville, R., & Pries-Heje, J. (2010). Explanatory Design Theory. *Business & Information Systems Engineering*, 2(5), 271–282.
- Baskerville, R., & Siponen, M. (2002). An information security meta-policy for emergent organizations. *Logistics Information Management*, 15(5/6), 337–346.
- Beck, R., Avital, M., Rossi, M., & Thatcher, J. B. (2017). Blockchain Technology in Business and Information Systems Research. *Business & Information Systems Engineering*, 59(6), 381–384.
- Beck, R., & Müller-Bloch, C. (2017). Blockchain as Radical Innovation : A Framework for Engaging with Distributed Ledgers. In *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)* (pp. 5390–5399). Waikoloa, USA.
- Beck, R., Müller-Bloch, C., & Ling, L. J. (2018). Governance in the Blockchain Economy: A Framework and Research Agenda. *Journal of the Association for*

- Information Systems*, in press.
- Beck, R., Stenum Czepluch, J., Lollike, N., & Malone, S. (2016). Blockchain - The Gateway to trust-free cryptographic Transactions. In *Proceedings of the 24th European Conference on Information Systems (ECIS)*. Istanbul, Turkey.
- Beck, R., Weber, S., & Gregory, R. W. (2013). Theory-generating design science research. *Information Systems Frontiers*, 15(4), 637–651.
- Bélanger, F., & Crossler, R. E. (2011). Privacy in the digital age: a review of information privacy research in information systems. *MIS Quarterly*, 35(4), 1017–1042.
- Bogner, A., Chanson, M., & Meeuw, A. (2016). A decentralised sharing app running a smart contract on the ethereum blockchain. In *Proceedings of the 6th International Conference on the Internet of Things*. Stuttgart, DE: ACM.
- Bond, E. W. (1982). A direct test of the "Lemons" model: The market for used pickup trucks. *The American Economic Review*, 72(4), 836–840.
- Bonvin, N. (2012). *Linear Scalability of Distributed Applications*. École Polytechnique Fédérale de Lausanne, Thèse No. 5278.
- Brynjolfsson, E., & McAfee, A. (2012). *Race against the machine: How the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy*. Lexington, MA: Digital Frontier Press.
- Bulgurcu, B., Cavusoglu, H., & Benbasat, I. (2010). Information security policy compliance: an empirical study of rationality-based beliefs and information security awareness. *MIS Quarterly*, 34(3), 523–548.
- Buterin, V. (2013). Ethereum White Paper. Retrieved September 28, 2017, from <https://github.com/ethereum/wiki/wiki/White-Paper>
- Car-Pass. (2018). Car-Pass is your guarantee of an accurate odometer. Retrieved January 5, 2018, from <https://www.car-pass.be/en/about-car-pass>
- Carfax. (2018). Buying Used American Cars? Check the Carfax Report. Retrieved January 13, 2018, from <https://www.carfax.eu/de>
- CarJam. (2018). CarJam. Vehicle Facts, History, Money Owing and more. Retrieved January 17, 2018, from <https://www.carjam.co.nz/>
- Chandra Kruse, L., Seidel, S., & Gregor, S. (2015). Prescriptive knowledge in IS research: Conceptualizing design principles in terms of materiality, action, and boundary conditions. In *Proceedings of the 48th Hawaii International Conference on System Sciences (HICSS)* (pp. 4039–4048). Kauai, USA.
- Chandra Kruse, L., Seidel, S., & Purao, S. (2016). Making Use of Design Principles. In *Proceedings of the International Conference on Design Science Research in Information Systems and Technology (DESRIST)* (pp. 37–51). St. John's, Canada.
- Chanson, M., Gjoen, J., Risius, M., & Wortmann, F. (2018). Initial Coin Offerings (ICOs): The role of Social Media for Organizational Legitimacy and Underpricing. In *Proceedings of the 39th International Conference on Information Systems (ICIS)*. San Francisco, CA.
- Chatterjee, S., Sarker, S., & Valacich, J. S. (2015). The behavioral roots of information systems security: Exploring key factors related to unethical IT use. *Journal of Management Information Systems*, 31(4), 49–87.
- Chellappa, R. K., & Pavlou, P. A. (2002). Perceived information security, financial liability and consumer trust in electronic commerce transactions. *Logistics Information Management*, 15(5/6), 358–368.
- Chen, H., Chiang, R. H., & Storey, V. C. (2012). Business Intelligence and Analytics:

- From Big Data to Big Impact. *MIS Quarterly*, 36(4), 1165–1188.
- Chen, Y., & Zahedi, F. M. (2016). Individuals' Internet Security Perceptions and Behaviors: Polycontextual Contrasts Between the United States and China. *Mis Quarterly*, 40(1), 205–222.
- Christidis, K., & Devetsikiotis, M. (2016). Blockchains and Smart Contracts for the Internet of Things. *IEEE Access*, 4, 2292–2303.
- Clarke, R. (2016). Big data, big risks. *Information Systems Journal*, 26(1), 77–90.
- Conti, M., Kumar, S., Lal, C., & Ruj, S. (2018). A survey on security and privacy issues of bitcoin. *IEEE Communications Surveys & Tutorials*, in press.
- Cram, W. A., Proudfoot, J. G., & D'Arcy, J. (2017). Organizational information security policies: a review and research framework. *European Journal of Information Systems*, 26(6), 605–641.
- Crossler, R. E., & Posey, C. (2017). Robbing Peter to pay Paul: Surrendering privacy for security's sake in an identity ecosystem. *Journal of the Association for Information Systems*, 18(7), 487.
- Curtis, S. (2015). Visa uses Bitcoin's blockchain technology to cut paperwork out of car leasing. Retrieved March 18, 2018, from <https://www.telegraph.co.uk/technology/news/11961296/Visa-uses-bitcoins-blockchain-technology-to-cut-paperwork-out-of-car-rental.html>
- Davenport, T. H. (2013). Analytics 3.0. *Harvard Business Review*, 91(12), 64–72.
- Delone, W., & McLean, E. R. (2003). The DeLone and McLean model of information systems success: a ten-year update. *Journal of Management Information Systems*, 19(4), 9–30.
- Dinev, T., Hart, P., & Mullen, M. R. (2008). Internet privacy concerns and beliefs about government surveillance - An empirical investigation. *The Journal of Strategic Information Systems*, 17(3), 214–233.
- Dong, W., Liao, S., & Zhang, Z. (2018). Leveraging Financial Social Media Data for Corporate Fraud Detection. *Journal of Management Information Systems*, 35(2), 461–487.
- Egelund-Müller, B., Elsmann, M., Henglein, F., & Ross, O. (2017). Automated Execution of Financial Contracts on Blockchains. *Business & Information Systems Engineering*, 59(6), 457–467.
- European Commission. (2018). Data protection in the EU. Retrieved February 20, 2018, from https://ec.europa.eu/info/law/law-topic/data-protection_en
- Exergy. (2017a). Electric Power Technical Whitepaper. Retrieved March 28, 2018, from <http://exergy.energy/wp-content/uploads/2017/11/Exergy-WhitePaper-v5.pdf>
- Exergy. (2017b). Exergy Business Whitepaper. Retrieved March 28, 2018, from <https://exergy.energy/wp-content/uploads/2017/12/Exergy-BIZWhitepaper-v5.pdf>
- Fabian, B., Ermakova, T., & Sander, U. (2016). Anonymity in Bitcoin? The users' perspective. In *Proceedings of the 37th International Conference on Information Systems (ICIS)*. Dublin, Ireland.
- Fernandes, E., Rahmati, A., Eykholt, K., & Prakash, A. (2017). Internet of things security research: A rehash of old ideas or new intellectual challenges? *IEEE Security & Privacy*, 15(4), 79–84.
- Genesove, D. (1993). Adverse selection in the wholesale used car market. *Journal of Political Economy*, 101(4), 644–665.
- Gervais, A., Karame, G. O., Wüst, K., Glykantzis, V., Ritzdorf, H., & Capkun, S. (2016). On the security and performance of proof of work blockchains. In

- Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS)*. Vienna, Austria.
- Glaser, F. (2017). Pervasive decentralisation of digital infrastructures: a framework for blockchain enabled system and use case analysis. In *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)*. Waikoloa, USA.
- Goes, P. B. (2014). Editor's comments: big data and IS research. *MIS Quarterly*, 38(3), iii--viii.
- Goldfeder, S., Kalodner, H., Reisman, D., & Narayanan, A. (2018). When the cookie meets the blockchain: Privacy risks of web payments via cryptocurrencies. *Proceedings on Privacy Enhancing Technologies*, 2018(4), 179–199.
- Gomber, P., Kauffman, R. J., Parker, C., & Weber, B. W. (2018). On the Fintech Revolution: Interpreting the Forces of Innovation, Disruption, and Transformation in Financial Services. *Journal of Management Information Systems*, 35(1), 220–265.
- Gregor, S. (2006). The nature of theory in information systems. *MIS Quarterly*, 30(3), 611–642.
- Gregor, S., & Hevner, A. (2013). Positioning and Presenting Design Science Research for Maximum Impact. *MIS Quarterly*, 37(2), 337–355.
- Gregor, S., & Jones, D. (2007). The Anatomy of a Design Theory. *Journal of the Association for Information Systems*, 8(5), 312–335.
- Heikka, J., Baskerville, R., & Siponen, M. (2006). A design theory for secure information systems design methods. *Journal of the Association for Information Systems*, 7(11), 31.
- Hevner, A. (2007). A Three Cycle View of Design Science Research. *Scandinavian Journal of Information Systems*, 19(192), 87–92. Retrieved from <http://aisel.aisnet.org/sjis>
- Hevner, A., & Chatterjee, S. (2010). *Design research in information systems: theory and practice*. Berlin, Germany: Springer Science+Business Media.
- Hevner, A., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105.
- Hogan, C. E., & Wilkins, M. S. (2008). Evidence on the audit risk model: Do auditors increase audit fees in the presence of internal control deficiencies? *Contemporary Accounting Research*, 25(1), 219–242.
- Hyvärinen, H., Risius, M., & Friis, G. (2017). A Blockchain-Based Approach Towards Overcoming Financial Fraud in Public Sector Services. *Business & Information Systems Engineering*, 59(6), 441–456.
- Iansiti, M., & Lakhani, K. R. (2014). Digital ubiquity: How connections, sensors, and data are revolutionizing business. *Harvard Business Review*, 92(11), 90–99.
- Imbault, F., Swiatek, M., De Beaufort, R., & Plana, R. (2017). The green blockchain: Managing decentralized energy production and consumption. In *Proceedings of the 2017 IEEE International Conference on Environment and Electrical Engineering and 2017 IEEE Industrial and Commercial Power Systems Europe (EEEIC/I&CPS Europe)*, (pp. 1–5). Milan, Italy.
- Kolias, C., Kambourakis, G., Stavrou, A., & Voas, J. (2017). DDoS in the IoT: Mirai and other botnets. *Computer*, 50(7), 80–84.
- Kolias, C., Stavrou, A., Voas, J., Bojanova, I., & Kuhn, R. (2016). Learning internet-of-things security “hands-on.” *IEEE Security & Privacy*, 14(1), 37–46.
- Kuechler, W., & Vaishnavi, V. (2008). On theory development in design science research: anatomy of a research project. *European Journal of Information Systems*, 17(5), 489–504.

- Kuechler, W., & Vaishnavi, V. (2012). A framework for theory development in design science research: multiple perspectives. *Journal of the Association for Information Systems*, 13(6), 395–423.
- Kumar, A., Fischer, C., Tople, S., & Saxena, P. (2017). A traceability analysis of Monero's blockchain. In *Proceedings of the 22nd European Symposium on Research in Computer Security (ESORICS)* (pp. 153–173). Oslo, Norway.
- Lee, J. K., Cho, D., & Lim, G. G. (2018). Design and Validation of the Bright Internet. *Journal of the Association for Information Systems*, 19(2), 63–85.
- Liang, X., Zhao, J., Shetty, S., & Li, D. (2017). Towards data assurance and resilience in IoT using blockchain. In *Proceedings of the 2017 IEEE Military Communications Conference (MILCOM)* (pp. 261–266). Baltimore, MD.
- Lindman, J., Rossi, M., & Tuunainen, V. K. (2017). Opportunities and risks of Blockchain Technologies in payments – a research agenda. In *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)* (pp. 1533–1542). Waikoloa, Hi.
- Loebbecke, C., & Picot, A. (2015). Reflections on societal and business model transformation arising from digitization and big data analytics: A research agenda. *The Journal of Strategic Information Systems*, 24(3), 149–157.
- Lowry, P. B., Dinev, T., & Willison, R. (2017). Why security and privacy research lies at the centre of the information systems (IS) artefact: proposing a bold research agenda. *European Journal of Information Systems*, 26(6), 546–563.
- Lukyanenko, R., Evermann, J., & Parsons, J. (2015). Guidelines for Establishing Instantiation Validity in IT Artifacts: A Survey of IS Research. In *Proceedings of the International Conference on Design Science Research in Information Systems and Technology (DESRIST)* (pp. 430–438). Dublin, Ireland. <https://doi.org/10.1007/978-3-319-18714-3>
- Machado, C., & Fröhlich, A. A. M. (2018). IoT Data Integrity Verification for Cyber-Physical Systems Using Blockchain. In *Proceedings of the 2018 IEEE International Symposium on Real-Time Distributed Computing (ISORC)* (pp. 83–90).
- Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, 15(4), 336–355.
- March, S. T., & Smith, G. F. (1995). Design and natural science research on information technology. *Decision Support Systems*, 15(4), 251–266.
- Margulies, J. (2015). Garage door openers: An internet of things case study. *IEEE Security & Privacy*, 13(4), 80–83.
- Margulis, S. T. (2011). Three theories of privacy: An overview. In S. Trepte & L. Reinecke (Eds.), *Privacy online: Perspectives on privacy and self-disclosure in the social web* (pp. 9–17). Berlin, Germany: Springer Science+Business Media.
- McAfee, A., Brynjolfsson, E., Davenport, T. H., Patil, D. J., & Barton, D. (2012). Big data: the management revolution. *Harvard Business Review*, 90(10), 60–68.
- Medaglia, C. M., & Serbanati, A. (2010). An overview of privacy and security issues in the internet of things. In D. Giusto, A. Iera, G. Morabito, & L. Atzori (Eds.), *The Internet of Things* (pp. 389–395). New York: Springer.
- Meeuw, A., Schopfer, S., Ryder, B., & Wortmann, F. (2018). LokalPower: Enabling Local Energy Markets with User-Driven Engagement. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*. Montreal, CA.
- Mengelkamp, E., Gärttner, J., Rock, K., Kessler, S., Orsini, L., & Weinhardt, C. (2018).

- Designing microgrid energy markets: A case study: The Brooklyn Microgrid. *Applied Energy*, 210, 870–880.
- Meth, H., Mueller, B., & Maedche, A. (2015). Designing a requirement mining system. *Journal of the Association for Information Systems*, 16(9), 799–837.
- Miles, M. B., & Huberman, A. M. (1994). *Qualitative data analysis: An expanded sourcebook*. Thousand Oaks, CA: SAGE Publications.
- Miorandi, D., Sicari, S., De Pellegrini, F., & Chlamtac, I. (2012). Internet of things: Vision, applications and research challenges. *Ad Hoc Networks*, 10(7), 1497–1516.
- Modum. (2018). Data integrity for supply chain operations powered by blockchain. Retrieved March 22, 2018, from <https://modum.io/>
- Moody, G. D., Siponen, M., & Pahlila, S. (2018). Toward a Unified Model of Information Security Policy Compliance. *MIS Quarterly*, 42(1), 285–311.
- Moura, J., & Serrão, C. (2016). Security and privacy issues of big data. In N. Zaman, M. E. Seliaman, M. F. Hassan, & F. P. G. Marquez (Eds.), *Handbook of Research on Trends and Future Directions in Big Data and Web Intelligence* (pp. 20–51). Hershey, PA: IGI Global.
- Moyano, J. P., & Ross, O. (2017). KYC optimization using distributed ledger technology. *Business & Information Systems Engineering*, 59(6), 411–423.
- Münsing, E., Mather, J., & Moura, S. (2017). Blockchains for decentralized optimization of energy resources in microgrid networks. In *Proceedings of the 2017 IEEE Conference on Control Technology and Applications (CCTA)* (pp. 2164–2171). Mauna Lani, HI.
- Myers, M. D., & Newman, M. (2007). The qualitative interview in IS research: Examining the craft. *Information and Organization*, 17(1), 2–26.
- Nærland, K., Müller-Bloch, C., Beck, R., & Palmund, S. (2017). Blockchain to Rule the Waves-Nascent Design Principles for Reducing Risk and Uncertainty in Decentralized Environments. In *Proceedings of the 38th International Conference on Information Systems (ICIS)*. Seoul, South Korea.
- Nakamoto, S. (2008). Bitcoin: A Peer-to-Peer Electronic Cash System (White Paper). Retrieved September 29, 2017, from <https://bitcoin.org/bitcoin.pdf>
- Nambisan, S., Lyytinen, K., Majchrzak, A., & Song, M. (2017). Digital innovation management: Reinventing innovation management research in a digital world. *MIS Quarterly*, 41(1), 223–238.
- Negroponte, N. (1995). *Being Digital*. New York City, NY: Alfred A. Knopf.
- Newell, S., & Marabelli, M. (2015). Strategic opportunities (and challenges) of algorithmic decision-making: A call for action on the long-term societal effects of datification. *The Journal of Strategic Information Systems*, 24(1), 3–14.
- Niemimaa, E., & Niemimaa, M. (2017). Information systems security policy implementation in practice: from best practices to situated practices. *European Journal of Information Systems*, 26(1), 1–20.
- Notheisen, B., Cholewa, J. B., & Shanmugam, A. P. (2017). Trading Real-World Assets on Blockchain. *Business & Information Systems Engineering*, 59(6), 425–440.
- Noyen, K., Volland, D., Wörner, D., & Fleisch, E. (2014). When Money Learns to Fly: Towards Sensing as a Service Applications Using Bitcoin. Retrieved from <http://arxiv.org/abs/1409.5841>
- Nunamaker Jr, J. F., Chen, M., & Purdin, T. D. M. (1990). Systems development in information systems research. *Journal of Management Information Systems*, 7(3), 89–106.

- Oetzel, M. C., & Spiekermann, S. (2014). A systematic methodology for privacy impact assessments: a design science approach. *European Journal of Information Systems*, 23(2), 126–150.
- Pavlou, P. A. (2011). State of the information privacy literature: Where are we now and where should we go? *MIS Quarterly*, 35(4), 977–988.
- Peppers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A Design Science Research Methodology for Information Systems Research. *Journal of Management Information Systems*, 24(3), 45–77.
- Pilkington, M. (2016). 11 Blockchain technology: principles and applications. In F. X. Ollerios & M. Zhegu (Eds.), *Research Handbook on Digital Transformations* (pp. 225–253). Cheltenham, UK: Edward Elgar Publishing.
- Porter, M. E., & Heppelmann, J. E. (2015). How smart, connected products are transforming companies. *Harvard Business Review*, 93(10), 96–114.
- Pries-Heje, J., Baskerville, R., & Venable, J. R. (2008). Strategies for Design Science Research Evaluation. In *Proceedings of the 16th European Conference on Information Systems (ECIS)* (pp. 255–266). Galway, Ireland.
- Rai, A. (2017). Editor's comments: diversity of Design Science Research. *MIS Quarterly*, 41(1), iii--xviii.
- Risius, M., & Spohrer, K. (2017). A Blockchain Research Framework. *Business & Information Systems Engineering*, 59(6), 385–409.
- Robinson, O. C. (2014). Sampling in interview-based qualitative research: A theoretical and practical guide. *Qualitative Research in Psychology*, 11(1), 25–41.
- Roman, R., Zhou, J., & Lopez, J. (2013). On the features and challenges of security and privacy in distributed internet of things. *Computer Networks*, 57(10), 2266–2279.
- Ronen, E., Shamir, A., Weingarten, A.-O., & O'Flynn, C. (2017). IoT goes nuclear: Creating a ZigBee chain reaction. In *Proceedings of the 2017 IEEE Symposium on Security and Privacy (SP)* (pp. 195–212). San Jose, CA.
- Schlossnagle, T. (2006). *Scalable internet architectures*. Indianapolis, IN: Sams Publishing.
- Schroeck, M., Shockley, R., Smart, J., Romero-Morales, D., & Tufano, P. (2012). Analytics: The real-world use of big data. Retrieved February 19, 2018, from <https://www-935.ibm.com/services/us/gbs/thoughtleadership/ibv-big-data-at-work.html>
- Sicari, S., Cappiello, C., De Pellegrini, F., Miorandi, D., & Coen-Porisini, A. (2016). A security-and quality-aware system architecture for Internet of Things. *Information Systems Frontiers*, 18(4), 665–677.
- Sicari, S., Rizzardi, A., Grieco, L. A., & Coen-Porisini, A. (2015). Security, privacy and trust in Internet of Things: The road ahead. *Computer Networks*, 76(2015), 146–164.
- Simon, H. A. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press.
- Singh, J., Pasquier, T., Bacon, J., Ko, H., & Eysers, D. (2016). Twenty security considerations for cloud-supported Internet of Things. *IEEE Internet of Things Journal*, 3(3), 269–284.
- Siponen, M., & Iivari, J. (2006). Six design theories for IS security policies and guidelines. *Journal of the Association for Information Systems*, 7(7), 445–472.
- Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: an interdisciplinary review. *MIS Quarterly*, 35(4), 989–1016.
- Spence, M. (1976). Informational aspects of market structure: An introduction. *The*

- Quarterly Journal of Economics*, 90(4), 591–597.
- Stone, E. F., Gueutal, H. G., Gardner, D. G., & McClure, S. (1983). A field experiment comparing information-privacy values, beliefs, and attitudes across several types of organizations. *Journal of Applied Psychology*, 68(3), 459–468.
- Takeda, H., Veerkamp, P., & Yoshikawa, H. (1990). Modeling design process. *AI Magazine*, 11(4), 37–48.
- Tian, F. (2016). An agri-food supply chain traceability system for China based on RFID & blockchain technology. In *Proceedings of the 13th International Conference on Service Systems and Service Management (ICSSSM)*. Kunming, China.
- Trappe, W., Howard, R., & Moore, R. S. (2015). Low-energy security: Limits and opportunities in the internet of things. *IEEE Security & Privacy*, 13(1), 14–21.
- Tschorsch, F., & Scheuermann, B. (2016). Bitcoin and Beyond : A Technical Survey on Decentralized Digital Currencies. *IEEE Communications Surveys & Tutorials*, 18(3), 2084–2123.
- TÜV Rheinland. (2015). Das Problem Tachomanipulation. Retrieved September 29, 2017, from https://www.arvato.com/content/dam/arvato/%0Adocuments/financial-solutions/PK_%0ATachomanipulation_TÜV_Rheinland.pdf
- Underwood, S. (2016). Blockchain beyond bitcoin. *Communications of the ACM*, 59(11), 15–17.
- Vaishnavi, V. K., & Kuechler, W. (2015). *Design science research methods and patterns: innovating information and communication technology*. Boca Raton, FL: CRC Press.
- Venable, J., Pries-Heje, J., & Baskerville, R. (2016). FEDS: a framework for evaluation in design science research. *European Journal of Information Systems*, 25(1), 77–89.
- Walls, J. G., Widmeyer, G. R., & El Sawy, O. A. (1992). Building an Information System Design Theory for Vigilant EIS. *Information Systems Research*, 3(1), 36–59.
- Weber, R. H. (2010). Internet of Things - New security and privacy challenges. *Computer Law & Security Review*, 26(1), 23–30.
- Webster, J., & Watson, R. T. (2002). Analyzing the past to prepare for the future: Writing a literature review. *MIS Quarterly*, xiii--xxiii.
- Wengraf, T. (2001). *Qualitative research interviewing: Biographic narrative and semi-structured methods*. Thousand Oaks, CA: SAGE Publications.
- Westin, A. F. (1967). *Privacy and freedom*. New York City, NY: Atheneum.
- Whittington, R., & Pany, K. (2015). *Principles of Auditing & Other Assurance Services*. New York City, NY: McGraw-Hill Education.
- Williams, L. G., & Smith, C. U. (2004). Web Application Scalability: A Model-Based Approach. In *Proceedings of the International Computer Measurement Group Conference (CMG)* (pp. 215–226). Las Vegas, USA.
- Winter, R. (2008). Design science research in Europe. *European Journal of Information Systems*, 17(5), 470–475.
- Xu, H., Dinev, T., Smith, J., & Hart, P. (2011). Information privacy concerns: Linking individual perceptions with institutional privacy assurances. *Journal of the Association for Information Systems*, 12(12), 798–824.
- Zyskind, G., Nathan, O., & Pentland, A. (2015). Decentralizing privacy: Using blockchain to protect personal data. *Proceedings of the 2015 IEEE Security and Privacy Workshops (SPW)*, 180–184.

Appendix

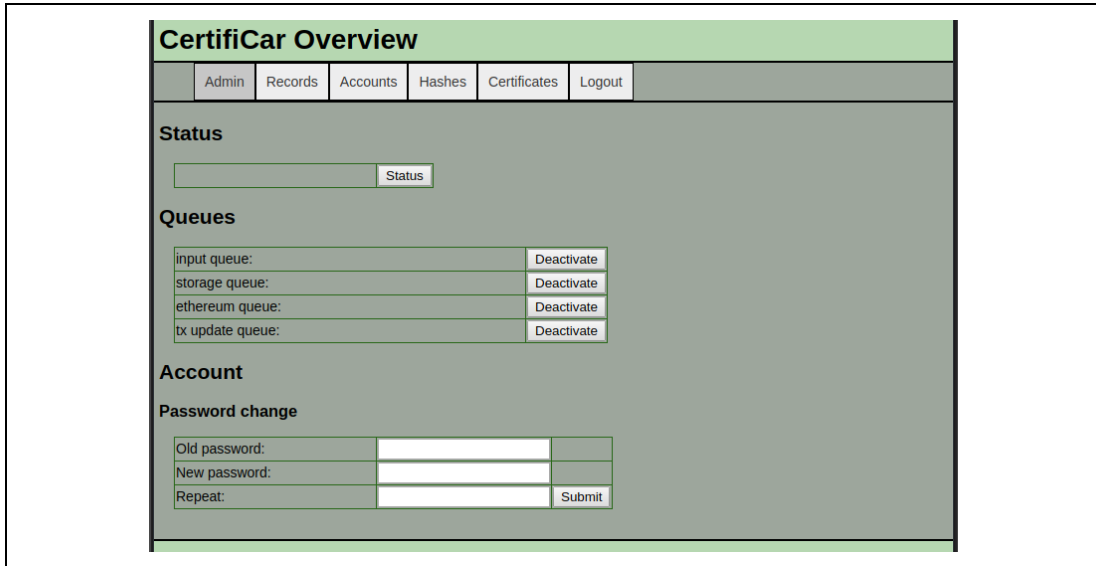


Figure A-1. First implementation of the web-based user interface

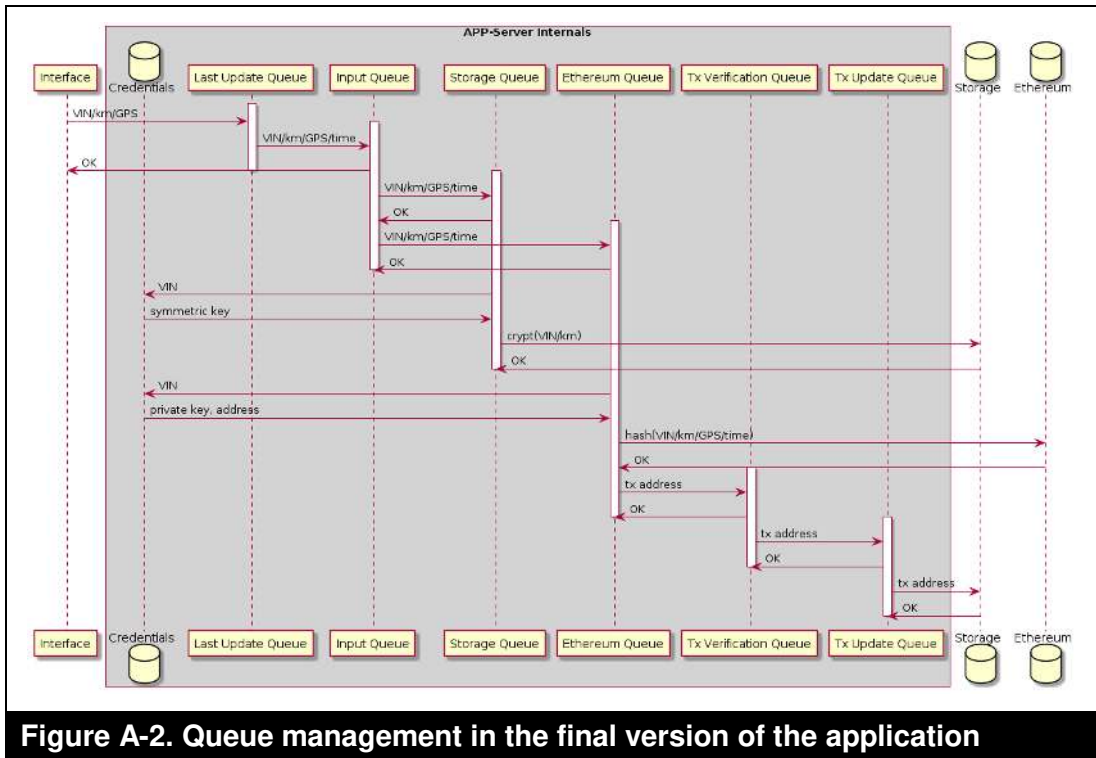


Figure A-2. Queue management in the final version of the application