Decision Analytics and Decentralized Ledger Technologies for Determination and Preservation of Spare Part Value in Aircraft Maintenance

Inaugural Dissertation
Presented to the
School of Business and Economics
Freie Universität Berlin
in Candidacy for the Degree of
doctor rerum politicarum (Dr. rer. pol.)

by
Clemens Wickboldt, M.Sc.

Berlin, 2019

Referees

1. Univ.-Prof. Dr. Natalia Kliewer, Freie Universität Berlin
2. Univ.-Prof. Dr. Martin Gersch, Freie Universität Berlin
3. Univ.-Prof. Dr. Hannes Rothe, Freie Universität Berlin

Date of Disputation

July 9th, 2019

Publications

- Meise, C., Wickboldt, C., Kliewer, N. (2019). "Decentralized Workshop Event Documentation with Hyperledger Fabric". In: *International Conference on Information Systems*. Submitted.
- Wickboldt, C. (2019). "Benchmarking a Blockchain-based Certification Storage System". In: *Diskussionsbeiträge Fachbereich Wirtschaftswissenschaft*. Vol. 2019/5 : Information Systems. Freie Universität Berlin. URL: https://refubium.fu-berlin.de/handle/fub188/24484?locale-attribute=e.
- Wickboldt, C., Kliewer, N. (2019). "Blockchain for Workshop Event Certificates - A Proof of Concept in the Aviation Industry". In: *Proceedings of the 27th European Conference on Information Systems (ECIS)*. ISBN: 978-1-7336325-0-8. URL: https://aisel.aisnet.org/ecis2019_rp/35/.
- Wickboldt, C., Kliewer, N. (2018). "Blockchain zur dezentralen Dokumentation von Werkstattereignissen in der Luftfahrtindustrie". In: *HMD Praxis der Wirtschaftsinformatik*, pp. 1–14. ISSN: 1436-3011. URL: https://doi.org/10.1365/s40702-018-00452-y.
- Wickboldt, C., Kliewer, N. (2018). "Value Based Pricing meets Data Science: A Concept for Automated Spare Part Valuation". In: *Multikonferenz Wirtschaftsinformatik (MKWI) 2018*. ISBN: 978-3-935786-72-0. URL: http://mkwi2018.leuphana.de/wp-content/uploads/MKWI2018_Band1.pdf.
- Wickboldt, C., Kliewer, N. (2016). "A framework for value-based pricing in aircraft surplus markets". In: *Multikonferenz Wirtschaftsinformatik (MKWI) 2016*. ISBN: 978-3-86360-132-4. URL: http://www.mkwi2016.de/download/MKWI2016_Research-in-Progress-Poster-Band.pdf.

Declaration of Authorship

Except where reference is made in the text, this thesis contains no material published elsewhere or extracted in whole or in part from a thesis presented by me for another degree or diploma. No other person's work has been used without due acknowledgment in the main text of the thesis. This thesis has not been submitted for the award of any other degree or diploma in any other tertiary institution.

Berlin, July 11, 2019

Clemens Wickboldt

Abstract

Aircraft spare parts are used to quickly replace defective parts and ideally avoid expensive aircraft-on-ground situations. Understanding the Fair Market Value of surplus parts is of eminent importance for the competitive advantage of a company. Decisions such as purchase, sale, storage or scrapping are made on the basis of the determined value. Domain experts state that the value of a part depends significantly on its specific characteristics, condition and workshop event history. If the documentation of this history is incomplete, this can lead to a complete loss of value of the part, since, for example, safety-relevant parts may no longer be used without complete documentation.

For companies that want to be able to survive in the highly competitive Maintenance Repair and Overhaul market, the use of digital technologies for data-based decision making has become unavoidable. The drowning of data while at the same time thirsting for information affects all market participants who manage their spare parts using digital technologies. The competitive advantage over others is now to use this data efficiently and make decisions based on data rather than experience and instinct. On the other hand, processes still exist in this industry that require documentation in paper form. One such process is the documentation of workshop events for safety-relevant spare parts by means of certificates. Low mutual trust and the heterogeneity of regional requirements in a global market prohibit the establishment of a central instance for data management.

The determination of a Fair Market Value was carried out manually for a long time, with great personnel effort and low reliability. The design of an Automated Spare Part Valuation concept provides a basis for data owners to use the amount of data reliably. Similar implementations in industry and with integrated automated evaluation prove the usability. The problem of incomplete certificates of workshop events is addressed and solved by the conception, implementation and evaluation of a Blockchain-based Certification System. The characteristics of a blockchain, in particular its decentralization and persistence, meet the requirements that could not previously be met in an environment with a lack of trust and due to the danger of a single point of failure.

Zusammenfassung

Flugzeugersatzteile dienen dem schnellen Austausch von defekten Teilen und vermeiden im Idealfall teure Aircraft-on-Ground-Situationen. Das Verständnis für einen Fair Market Value der überschüssigen Teile ist von eminenter Bedeutung für den Wettbewerbsvorteil eines Unternehmens. Entscheidungen wie Kauf, Verkauf, Einlagerung oder Verschrottung werden auf Basis des ermittelten Werts getroffen. Domänenexperten geben an, dass der Wert eines Teils maßgeblich von seinen spezifischen Charakteristika, seinem Zustand und seiner Werkstattereignishistorie abhängt. Ist der Nachweis dieser Historie lückenhaft, so kann es zum vollständigen Wertverlust des Teils kommen, da etwa sicherheitsrelevante Teile ohne lückenlose Nachweise nicht weiter verwendet werden dürfen.

Für Unternehmen, die in der Lage sein wollen im starken Wettbewerb des Maintenance Repair and Overhaul Markts zu bestehen ist der Einsatz digitaler Technologien zur datenbasierten Entscheidungsfindung mittlerweile unumgänglich. Das Ertrinken an Daten bei gleichzeitigem Verdursten an Informationen trifft alle Marktteilnehmer, die ihre Ersatzteile mittels digitaler Technologien verwalten. Der Wettbewerbsvorteil gegenüber anderen besteht nun darin, diese Daten effizient zu nutzen und Entscheidungen weniger nach Erfahrung und Instinkt, sondern datenbasiert zu treffen. Andererseits existieren auch in dieser Branche immer noch Prozesse, die eine Dokumentation in Papierform erfordern. Ein solcher Prozess ist die Dokumentation von Werkstattereignissen für sicherheitsrelevante Ersatzteile durch Zertifikate. Ein geringes Vertrauen untereinander und die Heterogenität regionaler Anforderungen in einem globalen Markt verbieten die Etablierung einer zentralen Instanz zur Verwaltung der Daten.

Die Ermittlung eines Fair Market Value erfolgte lange Zeit manuell, unter großem personellen Aufwand und geringer Zuverlässigkeit. Die Konzeption eines Automated Spare Part Valuation Konzepts bildet eine Grundlage für Inhaber von Daten, um die Menge an Daten verlässlich zu nutzen. Ähnliche Umsetzungen in der Industrie und mit integrierter automatisierter Bewertung belegen die Einsatzfähigkeit. Das Problem der lückenhaften Zertifikate von Werkstattereignissen wird durch die Konzeptionierung, Implementierung und Evaluation eines Blockchain-based Certification System adressiert und gelöst. Die Eigenschaften einer Blockchain, insbesondere die Dezentralität und Persistenz, erfüllen die Anforderungen, die in einem Umfeld mit mangelndem Vertrauen und aufgrund der Gefahr eines Single Point of Failure, bisher nicht zu erfüllen waren.

Contents

List of Figures

List of Tables

List of Algorithms

Chapter 1

Introduction

"We are drowning in information but starved for knowledge." stated Naisbitt and Cracknell (1984) already 35 years ago. This statement hasn't changed much, as Neely and Jarrar (2004) put it very similar: "Managers, today complain of 'drowning in data while thirsting for information'.". The use of data analytics by organizations is no longer a voluntary measure that improves the efficiency of their own processes in order to gain an edge over the competition. Meanwhile, companies that do not use their data with analytics are falling behind (Davenport, 2006). Digitizing is understood as "the process of converting analogue signals into a digital form" while digitization is "a sociotechnical process of applying digitizing techniques to broader social and instutional contexts" (Tilson et al., 2010). According to this, the view in IT strategy has moved from being aligned to business strategy (Venkatraman, 1994) to a view that recognized digital technologies for shaping business strategy. Digital technologies, such as business intelligence and analytics have become increasingly important over the last twenty years, both in academia and in practice. They can be defined as "techniques, technologies, systems, practices, methodologies, and applications that analyze critical business data to help an enterprise better understand its business and market and make timely business decisions" (Chen et al., 2012). In business intelligence applications, data moves from external sources and operational databases through an Extract Transform Load (ETL) (Kimball and Ross, 2011) process via data warehouse servers and analytic engines to a front-end application facing the user (Chaudhuri et al., 2011). Meanwhile, Deloitte (2019) identifies blockchain technology as becoming as important as analytics is today. As the technology "has reached the tip of a global hype across a variety of industries" (Notheisen, Hawlitschek, et al., 2017), it is now in a transition from a buzzword to real-world applications.

Within a supply chain network, information flows between buyer and supplier impact both parties positively through reduced costs of operations, enhanced productivity and improved management of assets (Klein and Rai, 2009). Dyer and Singh (1998) state that "alliance partners" would only share proprietary knowledge if they can be certain that this knowledge is not shared with the competition. Hence, there is a trust issue which has to be overcome. Digital technologies which speed up global supply chain orchestration are becoming a major source of competitive advantage (Bharadwaj et al., 2013). One example of a global value chain is the supply of spare parts in the aviation industry. This value chain is characterized by a large number of different parts that are moved in an intransparent market without central trading platforms and at the same time are subject to legal regulations in order to guarantee flight safety.

1.1 Motivation: The Example of Spare Parts in the Aviation Industry

Regular and conscientious maintenance is an essential part of ensuring flight safety. According to the International Air Transport Association (IATA), overall safety has been improved by 54% between 2007 and 2017 (IATA, 2017). The world aircraft fleet had a size of 25,870 aircrafts in 2017. An aircraft was utilized on average 2,601 hours a year (or around 7 hours per day) by performing 1,222 cycles (takeoffs and landings). In the same year, the airline industry recorded a profit after taxes of $38 billion which is an increase by 11% compared to the year before. Airlines spent globally $76 billion on Maintenance Repair and Overhaul (MRO) in 2017 which represents 11% of total operation costs. Until 2027 an annual MRO market growth of 4.6% is forecasted which would lead to a market size of $118 billion. The market for MRO services is divided by Original Equipment Manufacturers (OEM), MRO service providers and operators (Wyman, 2018b).

Maintenance activities are divided into four groups: base, line, engine and component. Base maintenance includes regularly aircraft checks. According to Quantas (2016) these checks range from short A-checks to long-lasting D-checks. A-checks are carried out every 10 weeks, including changing filters, renewing hydraulic fluids and checking emergency equipment. The B-check is increasingly performed within the other checks in modern aircraft. C-Checks are performed every 18 to 24 months and take three weeks. D-Checks are carried out approximately every 6 years. The complete aircraft is disassembled, all cabin details are removed so that engineers can inspect the metal frame. The chassis is dismantled and overhauled. All systems are subject to inspection. This check takes about 6 weeks, the costs amount to several million dollars. Base maintenance accounts for 20% of all maintenance costs. Line maintenance keeps the aircraft ready for on-time take-off by operational routines on-site and accounts for 17% of global MRO spending. Engine maintenance has own maintenance requirements and accounts for 42% of total costs. Material costs for components make up 21% of all MRO spending (IATA, 2018a). Between 2018 and 2028, 20,346 new deliveries for passenger aircrafts and 8,433 retirements are expected (Wyman, 2018a). Retired aircrafts will contribute to the component aftermarket by teardowns (MRO-Network, 2015).

Surplus parts (e.g from teardowns) move within an Alternative Closed Loop Supply Chain in the sense of Blumberg (2004). If there is an internal demand, they are stored in a pool stock for (future) replacement. Parts with value but no internal demand will be put on a trade stock and sold to the surplus market. Either from this surplus market network or directly from the manufacturer, parts could be procured. If they have no value at all, parts are scrapped. Decision-making anywhere in this supply chain requires knowledge about a part's Fair Market Value (FMV) which can be found by manual analysis and ideally reflects the most likely trading price (Kelly, 2008). A FMV is determined as the basis for decisions to scrap, use, repair or trade in aircraft spare parts. Studies have shown that the simple disposal of parts is a major cost contributor within the whole supply chain (Blumberg, 2004). In order to generate economic and ecological benefits, valuable parts may not be disposed, but could be used or sold to surplus markets, if the value is recognizable and quantifiable.

The industry is characterized by non-transparent transactions in the handling, repair and trade of spare sparts. Due to mutual distrust between market participants, there is a lack of uniform platforms for the exchange of information. Resulting information asymmetries lead to inefficiencies in the use of aircraft spare parts. Although reliably determining a FMV represents a competitive advantage and a basis for decision making, it has been done manually until today. A spare parts value is mostly dependent on its history and condition. Especially for safety-related spare parts, a complete lifecycle history is mandatory to use the part. Safety-related spare parts are valued on average in the six digits. Having an incomplete 'back to birth' documentation of these parts leads to a loss of their value, since they are not allowed to be used anymore. Still, tracking workshop events couldn't yet be digitized, as the industry is highly competitive and characterized by regional regulatory authorities.

The supply chain is determined by two interdependent characteristics of spare parts, value and provenance. Provenance is defined as "the history of ownership of a valued object or work of art or literature" (Merriam-Webster, 2019). Surplus parts have to be valued in order to have a basis for decision making. Also, provenance of (safety-related) spare parts has to be ensured to keep a part's value. This thesis focusses on evaluating methods from data analytics to reliably determine a FMV and decentralized ledger technologies to digitize the documentation process for workshop events in order to retain this value. These approaches are presented under the understanding of "Technology as Instrument", meaning systems which "improve efficiency of human decision-making, but do not constrain or manipulate human decision makers" (Markus and Rowe, 2018).

1.2 Research Approach

The design-oriented approach of Peffers et al. (2007) is used to explain the research process in this thesis. Figure 1.1 shows the Design Science Research Methodology (DSRM) process model of Peffers et al. (2007). First, the specific research problem is identified. To motivate the research process, the value of a possible solution is justified. Target parameters of the solution are then defined. This is followed by the development of the actual artefact. The developed artefact is applied in a suitable context to show that the artefact can solve the problem. In order to prove the effectiveness of the artefact, an evaluation takes place. This can be achieved for example by "quantitative performance measures" or "client feedback" (Peffers et al., 2007). The procedure concludes with the communication to a scientific or professional audience.

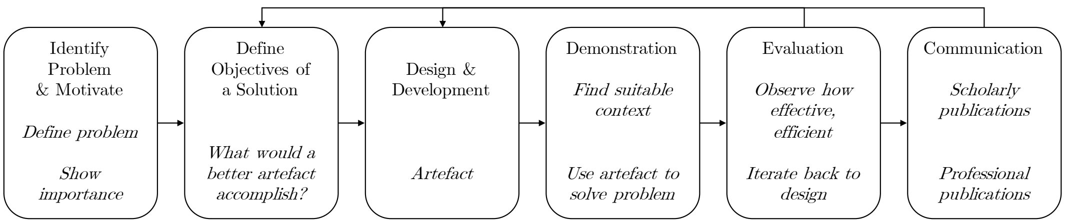

Figure 1.1: DSRM process model based on Peffers et al. (2007)

This work adresses two problems (P) that arise in the environment of valuating and tracking spare parts:

P1 A valuation of spare parts is only possible with great effort due to incomplete information about these parts.

P2 The history of workshop and trading events for safety-related spare parts cannot be documented transparently and persistently.

1.3 Thesis Outline

The development of a conceptual model for an Automated Spare Part Valuation (ASPV) is performed to adress problem 1 and the concept and implementation of a Blockchain-based Certification Storage System (BCSS) to adress problem 2.

1.3.1 Part 1: Automated Spare Part Valuation

Discussions with domain experts of the industry show that the surplus part market offers a great number of different products traded by only a few participants. First signs of Digital Economy are visible, such as platforms which list offers and prices. Still, transactions are frequently done manually via e-mail or telephone which leads to a poor market transparency. Further digitized modes of operations, such as end-to-end trading platforms or data-driven models for value determination and pricing are desperately needed to overcome the lack of transparency. The current manual value estimation does not only dismiss much information, it can also not keep up with the rapidly growing number of parts on the market. As pricing data is digitized today but lost in heterogenous data silos, the most pressing challenge is to understand how this data sources can be recombined as a basis for FMV estimation. Table 1.1 shows the application of the DSRM process to adress problem P1.

Activity	Approach
Problem identification and motivation	The reliable valuation of a spare part is the basis for supporting the decision how to use the part. Manual valuation leads to errors and bad decisions. The back to birth documentation for spare parts is performed analogue. This leads to loss of information, which in turn leads to loss of value of the parts.
Define the objectives for a solution	Objectives for a solution include automatically determining a FMV for spare parts and present it to the end user in a meaningful way. Also, it should prepare data as foundation for automated valuation of spare parts. A way for validation of the FMV has to be found.
Design and development	The artefact is a conceptual model and system architecture as a basis for implementation of an ASPV. Due to confidentiality of pricing data, publishing an instatiation of the conceptual model is not in scope.
Demonstration	To ensure that the artefact solves instances of the problem, the evaluated methods are discussed on scientific conferences and with domain experts in the aviation industry.
Evaluation	Feedback on efficacy of the artefact is given by domain experts which implemented a similar solution for a leading MRO service provider in the aviation industry.
Communication	Wickboldt and Kliewer (n.d.) communicate the embedding in pricing literature and propose a system architecture for a Fair Market Evaluator. In Wickboldt and Kliewer (2018b), an overview of methods from supply chain management and data analytics is given to propose a system for Automated Spare Part Valuation.

Table 1.1: Mapping research approach towards ASPV to DSRM activities of Peffers et al. (2007)

Discussions with domain experts have shown that the value of the spare part depends to a large extent on the special properties of the individual part in question. This includes both the part condition and the history of the past workshop events that led to this condition. This insight lead to start a second DSRM process and adressing problem 2.

1.3.2 Part 2: Blockchain-based Certification Storage System

Ensuring the authenticity of the provenance and condition of physical goods is a particular challenge in value chains that are both global and shaped by regional rules. The establishment of a central instance to manage data to track the location and condition of goods is virtually impossible in a highly competitive environment and little trust. This is why the aviation industry but also other global spanning transportation industries like the maritime industry (Stahlbock et al., 2018) are still producing paper documents to a considerable extent. Recent developments in blockchain technologies support better supply chain provenance (Armstrong, 2016). Due to their decentralized characteristics, blockchains can be used to evaluate provenance without one party claiming ownership of all supply chain data (Kim and Laskowski, 2018). Table 1.2 shows the application to adress problem P2.

Activity	Approach
Problem identification and motivation	The information of back to birth documentation for safety-related spare parts is stored on paper documents which can get lost or destroyed. This leads to loss of information, which in turn leads to loss of value of the parts.
Define the objectives for a solution	Analogue records of workshop events on paper hinder complete documentation and by that risk a loss of a part's value. Efforts to digitize the documentation process need to be taken in order to avoid gaps in the workshop event documentation. What is needed is a solution for fast, orderly, persistent, forgery-proof, access-restricted documentation solution that is aligned with the rules of the business process.
Design and development	A conceptual model for adressing the requirements is developed.
Demonstration	A proof of concept of the conceptual model in the aviation industry is developed.
Evaluation	Following a Technical Risk & Efficacy evaluation strategy (Venable et al., 2016), the proof of concept is evaluated in an artificial environment before incrementally moving it into a naturalistic environment.
Communication	A conceptual model is published in Wickboldt and Kliewer (2018a), a proof of concept in Wickboldt and Kliewer (2019) and performance measurements in Wickboldt (2019). The proof of concept is enhanced and benchmarked in Meise et al. (2019).

Table 1.2: Mapping research approach towards BCSS to DSRM activities of Peffers et al. (2007)

In order to arrive at these contributions, this thesis is organized as follows. Chapter 2 and 3 address problem P1. Problem P2 is addressed in chapters 4, 5, 6 and 7. Chapter 8 concludes and gives an outlook towards future research. Table 1.3 summarizes the mapping of publications[1] and thesis chapters.

Chapter	Title	Authors	Year	Publisher
2	A framework for value-based pricing in aircraft surplus markets	Wickboldt, Kliewer	2016	Multikonferenz Wirtschaftsinformatik
3	Value Based Pricing meets Data Science: A Concept for Automated Spare Part Valuation	Wickboldt, Kliewer	2018	Multikonferenz Wirtschaftsinformatik
4	Blockchain zur dezentralen Dokumentation von Werkstattereignissen in der Luftfahrtindustrie	Wickboldt, Kliewer	2018	HMD Praxis der Wirtschaftsinformatik
5	Blockchain for Workshop Event Certificates – A Proof of Concept in the Aviation Industry	Wickboldt, Kliewer	2019	European Conference on Information Systems
6	Benchmarking a Blockchain-based Certification Storage System	Wickboldt	2019	Diskussionsbeiträge Fachbereich Wirtschaftswissenschaft
7	Decentralized Workshop Event Documentation with Hyperledger Fabric	Meise, Wickboldt, Kliewer	2019	International Conference on Information Systems

Table 1.3: Mapping of publications to chapters of the thesis

[1]The publication Wickboldt (2019) was created in sole authorship. All other papers were created in co-authorship. All authors contributed equally to the papers.

References

Armstrong, S. (2016). *Eight things you need to know about the future of retail.* Electronic Article. URL: https://www.wired.co.uk/article/wired-retail-2015 (visited on 03/25/2019).

Bharadwaj, A., El Sawy, O., Pavlou, P., and Venkatraman, N. (2013). "Digital business strategy: toward a next generation of insights". In: *MIS quarterly*, pp. 471–482. ISSN: 0276-7783.

Blumberg, D. F. (2004). *Introduction to Management of Reverse Logistics and Closed Loop Supply Chain Processes.* CRC Press. ISBN: 9780203502938. DOI: 10.1201/b12427.

Chaudhuri, S., Dayal, U., and Narasayya, V. (2011). "An overview of business intelligence technology". In: *Communications of the ACM* 54.8, pp. 88–98. ISSN: 0001-0782.

Chen, H., Chiang, R. H., and Storey, V. C. (2012). "Business intelligence and analytics: From big data to big impact". In: *MIS quarterly* 36.4. ISSN: 0276-7783.

Davenport, T. H. (2006). "Competing on analytics". In: *harvard business review* 84.1, p. 98. ISSN: 0017-8012.

Deloitte (2019). *Tech Trends 2019: Beyond the digital frontier.* Electronic Article. URL: https://www2.deloitte.com/content/dam/Deloitte/br/Documents/technology/DI_TechTrends2019.pdf (visited on 03/15/2019).

Dyer, J. H. and Singh, H. (1998). "The relational view: Cooperative strategy and sources of interorganizational competitive advantage". In: *Academy of management review* 23.4, pp. 660–679. ISSN: 0363-7425.

IATA (2017). *IATA Annual Review 2017.* Electronic Article. URL: https://www.iata.org/publications/Documents/iata-annual-review-2017.pdf (visited on 03/20/2019).

— (2018a). *AIRLINE MAINTENANCE COST EXECUTIVE COMMENTARY: An Exclusive Benchmark Analysis (FY2017 data) by IATA's Maintenance Cost Task Force.* Electronic Article. URL: https://www.iata.org/whatwedo/workgroups/Documents/MCTF/MCTF-FY2017-Report-Public.pdf (visited on 03/15/2019).

Kelly, D. B. (2008). "Forecasting aircraft values: An appraiser's perspective". In: *Air-finance annual.*

Kim, H. M. and Laskowski, M. (2018). "Toward an ontology-driven blockchain design for supply-chain provenance". In: *Intelligent Systems in Accounting, Finance and Management* 25.1, pp. 18–27. ISSN: 1055-615X.

Kimball, R. and Ross, M. (2011). *The data warehouse toolkit: the complete guide to dimensional modeling.* John Wiley & Sons. ISBN: 1118082141.

Klein, R. and Rai, A. (2009). "Interfirm strategic information flows in logistics supply chain relationships". In: *Mis quarterly*, pp. 735–762. ISSN: 0276-7783.

Markus, M. L. and Rowe, F. (2018). “Is IT Changing the World? Conceptions of Causality for Information Systems Theorizing”. In: *Management Information Systems Quarterly* 42.4, pp. 1255–1280.

Meise, C., Wickboldt, C., and Kliewer, N. (2019). “Decentralized Workshop Event Documentation with Hyperledger Fabric (submitted)”. In: *Fortieth International Conference on Information Systems*.

Merriam-Webster (2019). *Provenance*. Electronic Article. URL: `https://www.merriam-webster.com/dictionary/provenance` (visited on 03/25/2019).

MRO-Network (2015). *Teardown Numbers Attract New Entrants*. Electronic Article. URL: `https://www.mro-network.com/maintenance-repair-overhaul/teardown-numbers-attract-new-entrants` (visited on 03/20/2019).

Naisbitt, J. and Cracknell, J. (1984). *Megatrends: Ten new directions transforming our lives*. Report. Warner Books New York.

Neely, A. and Jarrar, Y. (2004). “Extracting value from data–the performance planning value chain”. In: *Business Process Management Journal* 10.5, pp. 506–509.

Notheisen, B., Hawlitschek, F., and Weinhardt, C. (2017). “Breaking down the blockchain hype–towards a blockchain market engineering approach”. In: *European Conference on Information Systems*.

Peffers, K., Tuunanen, T., Rothenberger, M. A., and Chatterjee, S. (2007). “A design science research methodology for information systems research”. In: *Journal of management information systems* 24.3, pp. 45–77.

Quantas (2016). *The A, C and D of aircraft maintenance*. Electronic Article. URL: `https://www.qantasnewsroom.com.au/roo-tales/the-a-c-and-d-of-aircraft-maintenance/` (visited on 03/10/2019).

Stahlbock, R., Heilig, L., and Voß, S. (2018). “Blockchain in der maritimen Logistik”. In: *HMD Praxis der Wirtschaftsinformatik* 55.6, pp. 1185–1203. ISSN: 1436-3011.

Tilson, D., Lyytinen, K., and Sørensen, C. (2010). “Research commentary—Digital infrastructures: The missing IS research agenda”. In: *Information systems research* 21.4, pp. 748–759. ISSN: 1047-7047.

Venable, J., Pries-Heje, J., and Baskerville, R. (2016). “FEDS: a framework for evaluation in design science research”. In: *European Journal of Information Systems* 25.1, pp. 77–89. ISSN: 0960-085X.

Venkatraman, N. (1994). “IT-enabled business transformation: from automation to business scope redefinition”. In: *Sloan management review* 35.2, pp. 73–88. ISSN: 0019-848X.

Wickboldt, C. (2019). “Benchmarking a Blockchain-based Certification Storage System”. In: *Diskussionsbeiträge Fachbereich Wirtschaftswissenschaft*. Vol. 2019/5 : Information Systems. Freie Universität Berlin. URL: `https://refubium.fu-berlin.de/handle/fub188/24484?locale-attribute=en`.

Wickboldt, C. and Kliewer, N. “A framework for value-based pricing in aircraft surplus markets”. In: *Multikonferenz Wirtschaftsinformatik (MKWI) 2016*. ISBN:

978-3-86360-132-4. URL: http://www.mkwi2016.de/download/MKWI2016_Research-in-Progress-Poster-Band.pdf.

Wickboldt, C. and Kliewer, N. (2018a). "Blockchain zur dezentralen Dokumentation von Werkstattereignissen in der Luftfahrtindustrie". In: *HMD Praxis der Wirtschaftsinformatik*, pp. 1–14. ISSN: 1436-3011. URL: https://doi.org/10.1365/s40702-018-00452-y.

— (2018b). "Value Based Pricing meets Data Science: A Concept for Automated Spare Part Valuation". In: *Multikonferenz Wirtschaftsinformatik (MKWI) 2018*. ISBN: 978-3-935786-72-0. URL: http://mkwi2018.leuphana.de/wp-content/uploads/MKWI2018_Band1.pdf.

— (2019). "Blockchain for Workshop Event Certificates - A Proof of Concept in the Aviation Industry". In: *Proceedings of the 27th European Conference on Information Systems (ECIS)*. ISBN: 978-1-7336325-0-8. URL: https://aisel.aisnet.org/ecis2019_rp/35/.

Wyman, O. (2018a). *GLOBAL FLEET & MRO MARKET FORECAST COMMENTARY*. Electronic Article. URL: https://www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2018/January/2018-2028_Global_Fleet_MRO_Market_Forecast_Commentary_Public_Final_web.pdf (visited on 03/15/2019).

— (2018b). *MRO SURVEY 2018: TACKLING INDUSTRY DISRUPTION*. Electronic Article. URL: https://www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2018/april/MRO-Survey-2018-web.pdf (visited on 03/15/2019).

Chapter 2

A framework for value-based pricing in aircraft surplus markets

Abstract

Aircraft Maintenance, Repair and Overhaul (MRO) companies are facing pricing related decision problems on a daily basis. Although trading platforms are listing offers for some parts, prices are rarely published. Knowing a Fair Market Value (FMV) in this market represents a significant competitive advantage, be it to assess an aircraft for a tear-down or support internal part related decision-making processes. The goal of this work is to develop a framework which is able to aggregate relevant data from different sources, automatically determine a FMV and present context information around it in a transparent and user friendly manner. The Fair Market Evaluator is embedded in a comprehensive software system which enables the user to automatically evaluate single parts or whole part packages. Incomplete historical data sets, aggregation of information from different data sources and integration into existing IT infrastructure are some of the challenges we will address in our ongoing research work.

Contents

2.1 Introduction

In order to maintain an airline's fleet, parts are needed in a timely manner. New parts can be procured directly from a supplier, used ones from the aircraft surplus market. If there is no direct use, parts will be stored. Parts with decreasing internal demand will be sold again on the surplus market, which has been growing over the last couple of years. One reason for this growth is an increasing number of parts from aircraft tear-downs. At the same time, this market still lacks transparency, e.g. transactions are usually done via e-mail or telephone. The market for aircraft surplus parts is characterized by a great number of different products but just a few market participants.

There are several ways to determine a price, including auction based market places and prediction models based on historical data. The compartmentalized structure of traders and MRO companies in the aircraft industry may be one reason why there is still no end-to-end auction based market place for surplus parts. As soon as context based FMVs are needed for internal decision making, external determined prices fail to include company-specific information and don't provide a sufficient foundation for decision-making. Predicting a FMV based on historical data however may work well for market players with high amounts of structured and high quality historical data. The goal of this work is to develop a framework which is able to automatically determine a FMV and present context information around it in a transparent and user friendly manner.

An example use case would be: A supplier offers an airplane for a tear-down. The value of this plane would be the sum of all the parts in it. The user would have to determine a Fair Market Value for every single part of this plane and sum it all up to get an idea of what would be a reasonable price for it. Alternatively, the user could let the Fair Market Evaluator calculate values for each item and in that way provide a solid foundation for decision making. If there is not enough data available for a reliable Fair Market Value calculation, the user still gets all the aggregated information as internal historical prices and external market prices, which supports a manual Fair Market Value determination.

2.2 Fair Market Value-based pricing

In order to find an approach to determinate a Fair Market Value, we take a look at existing pricing methods, define the term Fair Market Value and see if we can combine both.

It has been shown that price has the greatest impact on a company's earnings before interest and taxes (Hinterhuber, 2004). The way a product's or service's price is determined can be categorized into three different groups (Hinterhuber, 2008):

1. Cost-based pricing uses data from cost accounting and for example considers original purchase price and internal costs (e.g. repairing) to determine a sale price. It doesn't take the competition or market into account and is therefore considered as the weakest approach.

2. Competition based-pricing observes price levels of the competition and uses market prices for orientation in price setting.

3. Value-based pricing uses a predefined value as the basis for determining the price.

As we are looking for a value which is able to serve as foundation for decision making and price setting, value-based pricing seems the way to go. In the literature, value is defined as the difference between the customer's anticipated benefit of a good or service and the supplier's cost of producing it (Johansson et al., 2015) or as benefits received in exchange for price (Anderson and Narus, 1998). This work focuses on a business to business secondary market in the aircraft industry, which mainly differentiates between the following value terms:

- Base Value is considered as the economic value and assumes balanced supply and demand (Kelly, 2008), completely informed market participants and is considered as a hypothetic value (Ackert, 2012).

- (Current) Market Value is defined value after manual analysis (Ackert, 2012) or the "most likely trading price" and is synonymous with Fair Market Value (Kelly, 2008).

Still, many companies are not able to benefit from value-based pricing. One obstacle is that the value has to be determined before it can be used as an argument for pricing (Hinterhuber, 2008). Value-based pricing and value estimation methods are often related to product introductions in primary markets. State of the art literature about value-based pricing takes the customer point of view for value determination. Common methods are surveys or conjoint analysis (Anderson et al., 1992). This is applicable to this scenario in a limited way only. Our sample data contains more than 500.000 different parts[1] . Conducting surveys for this large quantity of different products seems to be a disproportionate effort. Also, as soon as context-based values are needed, external determined prices fail to include company-specific information and don't provide a sufficient foundation for a company's decision-making. Therefore, a new and innovative approach is needed.

[1]Real world data is being provided by a large MRO company and can be used to test the Fair Market Evaluator.

We are working on a framework which is able to determine a part's value under consideration of the company's specific situation. It contains information about the following:

- Original launch price,
- historical purchase price within the company,
- historical offer price within the company,
- historical selling price within the company,
- historical internal and manual evaluation and
- current external market price.

We suggest a combination of a set of business intelligence methods to transparently determine and present a Fair Market Value. Based on the results of Voß and Lessmann (2013), who had determined prices for the used car market, an application of prediction methods seems to be promising to evaluate secondary market items. There is a wide variety of methods[2] , which may perform differently in this specific scenario and will be evaluated in the final paper.

The Fair Market Evaluator framework delimits from state of the art value-based pricing methods as shown in table 2.1. Value-based pricing and value estimation methods are related to product introductions in primary consumer markets[3]. This work focuses on a business to business secondary market in the aircraft industry. Value-based pricing focuses on an individual consideration of a small amount of products, which is an obvious scope for pricing related customer surveys and conjoint analysis. This work aims at automatic evaluation of thousands of different products and still present the valuation in a transparent way.

Dimension	State of the art	Fair Market Evaluator
Target	Primary Markets	Secondary Market
Audience	B2C	B2B
Scope	Small amount of products	Thousands of products
Automation	Individual consideration	Automated evaluation

Table 2.1: Delimitation from state of the art value-based pricing methods

An automated estimation of the FMV as a basis for evaluating purchasing and selling prices represents a promising approach to address the transparency issue. It

[2]See Hastie et al. (2009) for an overview.
[3]See for example Hinterhuber (2004, p. 767).

could not only be used for supporting the surplus part utilization process but also in any department which would work with value-based pricing methods.

2.3 Fair Market Evaluator embedded in comprehensive software system

After defining the approach to determine the Fair Market Value, we introduce a software framework which aggregates information from different sources, prepares data for the Fair Market Evaluator model and presents the results to the end user.

In the past, FMV determination has been done manually, mainly based on experts' estimates, including the risk of not using all available information to determine the FMV objectively. Until now and for various organizational reasons, not all historical pricing data is collected and stored. Because of incomplete datasets the FMV can be recognized directly in only a few cases. Prediction models naturally assume a complete set of historical data and can't deal with missing datasets. We will address this challenge by relying on statistical methods for dealing with missing data.

A similar approach is a web marketing system with automatic pricing, called *digiprice*, which was proposed by Abe and Kamba (2000). Their goal is to maximize the profit within an on-line marketing site. Table 2.2 shows the differences between our proposed Fair Market Evaluator and digiprice.

Dimension	digiprice	Fair Market Evaluator
Market	Digital content (p. 775)	Aircraft surplus parts
Audience	"consumers" (p. 776)	B2B
Scope	Marketing and pricing (p. 775)	Value determination
Goal	Find price which maximizes the overall profit (p. 775)	Find value which reflects the part's value within the company
Considered variables	Historic prices, historic sales, product description provided by seller (p. 776)	Introductory price, current market price, historic purchase price, repair costs, historic offer, historic sell price, historic manual evaluation
Algorithms	Stochastic and linear approximation of optimal price (p. 779 and p. 781)	Forecasting per multiple linear regression for now.

Table 2.2: Differentiation of Fair Market Evaluator from alternative framework

As digiprice is looking to price digital content for consumers, we're determining a value for aircraft surplus parts in a B2B market. Also, we're not looking for a sales price only but for a value which is the foundation for decision making and a basis for deducing prices. At last, we consider multiple dimensions within a linear regression model whereas digiprice considers three dimensions and uses stochastic and linear approximation methods.

The Fair Market Evaluator itself is embedded in a comprehensive software system to support pricing, trading and logistic processes, see figure 2.1 for an approach.

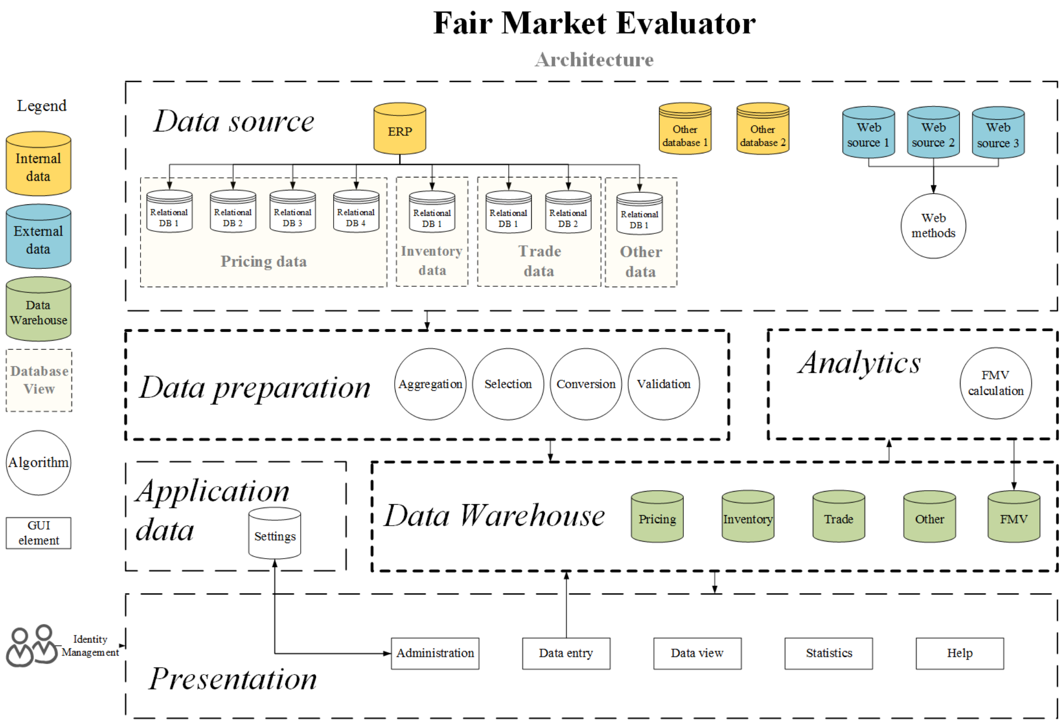

Figure 2.1: Approach for a Fair Market Evaluator system architecture

The system architecture of the Fair Market Evaluator is divided into six layers. The first layer contains data gathering from different sources. Internal data lies in a corporate ERP with relational databases in the back-end, which are aggregated in logical database views for convenient access. Additionally, there are other (non-)relational databases, outside the ERP. External data is collected via web interfaces from data providers, e.g. trading platforms. In the data preparation layer, data is aggregated, relevant data is selected, converted to one consistent format and validated in the end. The data flows in one direction to avoid contamination of the production data. The Data Warehouse layer contains one data base per data category. In this case, there are pricing, inventory, trade and master data. The analytics layer calculations are supplied by Data Warehouse data. Calculated FMVs are stored in a database in the Data Warehouse layer. The presentation layer provides the user interface. The user is able to enter data into the data warehouse as well as change existing data. In this way, the user is able to override a modelsuggested Fair Market

Value which will enrich the model's quality with information about the prediction's accurancy. The presentation layer also provides statistics which are only available because of the prepared data stored in the Data Warehouse. Administrative settings are stored in a database within the application data layer. Authentication is handled by a corporate identity management system.

2.4 Outlook

An automated FMV calculation may deliver valid results if the underlying data is complete and of good quality. An automated value determination addresses the market characteristic of many different products. Data aggregation prior to the Fair Market Value calculation could be one way to build a generic surplus part trading platform and overcome the market's transparency issue. In the end, an automated value determination prevents bad part utilization decisions and could lead to a more efficient and sustainable use of surplus material.

Subsequent to introducing the Fair Market Evaluator framework, an evaluation of it's utility and benefit will follow. A possible approach to determine the prediction assurance ex-post is to measure the amounts of FMVs which had been overwritten by the user. Evaluation of the framework in practice and empirical insights in cooperation with a MRO company are one part of the final paper.

References

Abe, N. and Kamba, T. (2000). "A Web marketing system with automatic pricing". In: *Computer Networks-the International Journal of Computer and Telecommunications Networking* 33.1-6, pp. 775–788. ISSN: 1389-1286.

Ackert, S. (2012). "Basics of Aircraft Market Analysis". In: *Aircraft Monitor.*

Anderson, J., Jain, D., and Chintagunta, P. (1992). "Customer value assessment in business markets: A state-of-practice study". In: *Journal of Business-to-Business Marketing* 1, pp. 3–29.

Anderson, J. and Narus, J. (1998). "Business marketing: understand what customers value". In: *Harvard business review* 76, pp. 53–67.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). "The Elements of Statistical Learning: Data Mining, Inference, and Prediction". In: *Springer Series in Statistics.*

Hinterhuber, A. (2004). "Towards value-based pricing—An integrative framework for decision making". In: *Industrial Marketing Management* 33.8, pp. 765–778. ISSN: 00198501. DOI: 10.1016/j.indmarman.2003.10.006.

— (2008). "Customer value-based pricing strategies: why companies resist". In: *Journal of Business Strategy* 29.4, pp. 41–50. ISSN: 0275-6668. DOI: 10.1108/02756660810887079.

Johansson, M., Keränen, J., Hinterhuber, A., Liozu, S., and Andersson, L. (2015). "Value assessment and pricing capabilities? How to profit from value". In: *Journal of Revenue & Pricing Management.*

Kelly, D. B. (2008). "Forecasting aircraft values: An appraiser's perspective". In: *Air-finance annual.*

Voß, S. and Lessmann, S. (2013). "Resale Price Prediction in the Used Car Market". In: *Tristan Symposium VIII.*

Chapter 3

Value Based Pricing meets Data Science: A Concept for Automated Spare Part Valuation

Abstract

Turning data into value is an exciting challenge for Data Science in times of an exponential growing amount of data. Maintenance, Repair and Overhaul companies are facing pricing related decision problems on a daily basis. The industry sits on vast amount of data. Due to lacks of transparency in the surplus part market and missing concepts to efficiently use internal data, existing information is not used exhaustively to improve data-based part utilization decisions. An early-stage concept for automated spare part valuation which classifies pricing data before applying appropriate valuation methods is presented and hereby combines methods from multiple disciplines. Information from heterogeneous sources is aggregated, transformed and then supports machine learning methods to automatically determine a Fair Market Value for surplus spare parts. Handling incomplete historical data sets as well as validating the calculated Fair Market Value are some of the challenges which become visible.

Keywords Automated Value Determination, ETL, Data Classification, Spare Parts, Fair Market Value

Contents

3.1 Introduction

"We are drowning in information but starved for knowledge." stated author John Naisbitt in his book Megatrends in the year 1984 (Naisbitt and Cracknell, 1984). Until today, decision makers still complain in a very similar way about "drowning in data while thirsting for information" (Neely and Jarrar, 2004). As this is a general problem, it's especially true for market participants with a vast number of different products, for example the aircraft spare parts industry. The UK's Royal Airforce alone is managing around 685.000 line items (Eaves and Kingsman, 2004), the largest commercial airplane Airbus A380 consists of around 2.5 million individual parts[1]. The aircraft parts industry is an especially interesting example since aircraft maintenance has to face high safety regulations but also operate in a very short timeframe to avoid very expensive aircraft downtime. Maintenance Repair and Overhaul (MRO) companies, especially in the aircraft industry, are constantly confronted with managing the logistics of spare parts. An inexpensive repair and maintenance in a timely manner requires an efficiently organized material supply. It cannot always be resorted to new parts, either for cost reasons or because certain parts are not available on the market. Additionally, constant pulling and replacing of parts leads to a high number of used parts in stock which not always have a direct use, in other words they are surplus. The market for aircraft surplus spare parts is characterized by a great number of different products but just a few market participants. Transactions are done via e-mail or telephone. Although trading platforms are listing offers for some parts, prices are rarely published. The compartmentalized structure of traders and MRO companies in the aircraft industry may be one reason why there is still no end-to-end auction based market place for surplus parts. Valuating spare parts on the market or in stock in this environment is done today in a time consuming and error-prone manual manner by industry experts.

There are several ways to determine a price as foundation for the part utilization decision, including auction based market places and prediction models based on historical data. As soon as context based value information is needed for internal decision making, external determined prices fail to include company-specific information and don't provide a sufficient basis for decision-making. Predicting a value based on historical data however may work well for market players with a huge amount of structured and high quality historical data. In this paper, we propose a concept for combining methods from multiple disciplines with the goal of automated part valuation. The aircraft industry is used as an example. This work addresses problems of the MRO industry in general. Existing pricing approaches which had been already discussed in Wickboldt and Kliewer (n.d.) are shortly described in Section 3.2.1. In cooperation with a large aircraft MRO provider, we gained insights

[1]Airbus Press Office 2017: Facts & Figures, June 2017

in existing spare part utilization processes which are described in Section 3.2.2. In Section 3.3 Pareto Principle and Product Lifecycle are used to classify pricing data which is a basis for an extension of the developed system architecture for automated spare part valuation (ASPV) of Wickboldt and Kliewer (n.d.). Section 3.3 closes with approaches for handling the validation of the resulting values. Section 3.4 summarizes this work and gives an outlook for what will be discussed during following research.

3.2 State of the Art of Pricing and Surplus Part Supply Chain

This section gives an overview about existing pricing approaches and the surplus spare part supply chain to understand the need for a part's value as basis for the utilization process.

3.2.1 Pricing

It has been shown that price has the greatest impact on a company's earnings before interest and taxes (Hinterhuber, 2004). Existing literature provides approaches for pricing. Before proceeding to automatically finding a surplus part's value, a short overview about the underlying terminologies is given. According to Anderson et al. (2000), '*Price* in business market is what a customer firm pays a supplier for its product offering'. Furthermore, Anderson and Narus (1998) state, '*Value* in business markets is the worth in monetary terms of the technical, economical service and social benefits a customer company receives in exchange for the price it pays for a market offering'. Pricing approaches can be categorized into three different groups (Hinterhuber, 2008):

Cost Based Pricing uses data from cost accounting and for example considers original purchase price and internal costs (e.g. repairing) to determine a selling price. It doesn't take the competition or market into account and is therefore considered as the weakest approach.

Competition Based Pricing observes price levels of the competition and uses market prices for orientation in price setting but dismisses company-internal information such as inventory or repair costs of defective spare parts.

Value Based Pricing uses a predefined value as the basis for determining the price. As only value based pricing takes company internal information into account, this approach is now described in more detail to discuss its potential to serve as a foundation for decision making and price setting.

This work focuses on a business to business secondary market for surplus parts which mainly differentiates between the following more granular value terms. *Base*

Value is considered as the economic value and assumes balanced supply and demand (Kelly, 2008) as well as completely informed market participants and thus is considered as a hypothetic value (Ackert, 2012). *(Current) Market Value* is defined as determined value after manual analysis (Ackert, 2012) or the 'most likely trading price' and is used synonymously with Fair Market Value (Kelly, 2008). Mercer and Brown (1999) state that the *Fair Market Value (FMV)* is '(..) the price at which the property would change hands between a willing buyer and a willing seller when the former is not under any compulsion to buy and the latter is not under any compulsion to sell, both parties having reasonable knowledge of the relevant facts'.

Still, many companies are not able to benefit from Value Based Pricing. One obstacle is that the value has to be determined before it can be used as an argument for pricing (Hinterhuber, 2008). Value Based Pricing and value estimation methods are often related to product introductions in primary markets. State-of-the-art literature about Value Based Pricing takes the customer point of view for value determination. Common methods are surveys or conjoint analysis (Anderson et al., 1992). This is applicable to this scenario in a limited way only. Conducting surveys for this large quantity of different products seems to be a disproportionate effort. Also, as soon as context based values are needed, external determined prices fail to include company-specific information and don't provide a sufficient foundation for a company's decision-making. Therefore, an approach is needed which includes company internal information and is feasible to a large number of parts.

3.2.2 Understanding the Spare Part Supply Chain

As pointed out in Blumberg (2004), products at the end of the traditional supply chain may be surplus but still have value. As shown in Figure 3.1, surplus parts are stored, depending on their value.

Expert interviews with a large MRO provider showed that parts with value but no internal demand will be put on a trade stock and sold to the surplus market via direct sale or auction. From the same surplus market, parts could be bought (B). If there is internal demand they are stored in a pool stock for (future) replacement (R). The pool stock could also be filled by parts pulled out of the operations, such as an aircraft or production facility (P). If they have no value at all, they will be scrapped. Because of the fact that parts circulate in a loop but not return to the manufacturer, this process is called Alternative Closed Loop Supply Chain Process.

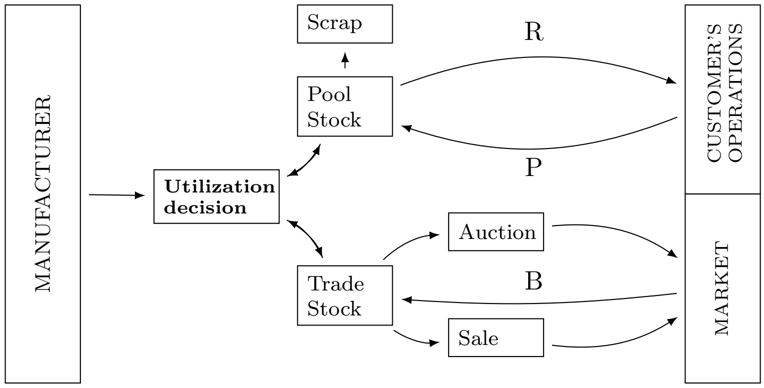

Figure 3.1: Alternative Closed Loop Supply Chain Process from a MRO point of view

Studies have shown that simple disposal of parts which have failed or appeared to have failed is a major cost contributor within the whole supply chain (Blumberg, 2004). In this situation, reverse logistics are important from an economic[2] and ecological[3] point of view. Blumberg (2004) also points out, that extra revenues are realized if there is an ability to liquidate parts with value in secondary markets. MRO companies therefore need to estimate a value as characteristic parameter and foundation for price derivation in order to support internal decision-making processes.

3.3 Framework for Automated Spare Part Valuation

From characteristics of the surplus market, an overview of relevant pricing literature and value definitions, the following research goals result:

- Design a framework in order to automatically determine a FMV for surplus spare parts and present it to the end user in a meaningful way (*ASPV framework* in following).
- Find a way to prepare data as foundation for automated valuation of surplus spare parts.
- Validate the calculated FMV.

The resulting ASPV framework should be able to determine a part's FMV under consideration of the company's specific situation. The calculated value contains

[2]Valuable parts will not be disposed but could be used or sold.
[3]Probably fewer parts will be disposed if the value is recognizable.

information as displayed in Table 3.1. The variables are categorized by time (historic and current) and by point of view (company-internal and -external).

Time	*Internal*	*External*
Historic	Purchase price, offer price, selling price	Original launch price
Current	–	Market price

Table 3.1: Variables for FMV calculation

We suggest a combination of a set of data analytics methods to transparently determine and present a FMV. Based on the results from Voß and Lessmann (2013), proposing prediction methods for determining prices for the used car market, an application of prediction methods seems to be promising in order to evaluate secondary market items. There is a wide variety of machine learning methods, from simple Linear Regression Models to Neural Networks which may perform differently in this specific scenario[4].

The ASPV framework differs from state-of-the-art Value Based Pricing methods as shown in Table 3.2. State-of-the-art Value Based Pricing and value estimation methods are related to product introductions in primary consumer markets. The proposed ASPV framework focuses on business to business secondary markets in the MRO industry. State-of-the-art Value-Based Pricing focuses on an individual consideration of a small number of products which is an obvious scope for pricing related customer surveys and conjoint analysis. On the contrary, this work aims to automated valuation of thousands of different products from multiple variables (see Table 3.1), present the valuation in a transparent way and in a second phase increase objectivity in contrast to the status quo results from manual value determination.

Dimension	*State-of-the-art Value Based Pricing*	*ASPV framework*
Target	Primary markets	Secondary market
Audience	B2C	B2B
Scope	Small number of products	Thousands of products
Automation	Individual consideration	Automated valuation

Table 3.2: Differentiation from state-of-the-art Value Based Pricing methods

3.3.1 Classification of Pricing Information

As Value Based Pricing is effort intensive, it should not be used on all products. To address the effort issue, we propose a rule-based segmentation of the products based

[4]An extensive overview can be found in Hastie et al. (2009)

on their need for an approximation on the one side or for exact value determination on the other side. To reduce complexity in the classification stage we prefer this rule-based approach over machine learning.

Grosfeld-Nir et al. (2007) take a closer look at the Pareto principle which led to a classification of products into A, B and C groups. The 'A' group (the 'vital few'), consisting of approximately 20% of the attributes (surplus spare parts in this case), accounts for 80% of the phenomenon (value in this case); the 'B' group, i.e. the next 30% of the items, accounts for 10% of the phenomenon, and the 'C' group (the 'trivial many'), which contains 50% of the items, accounts for also 10% of the phenomenon. Given that, Ronen et al. (2005)[5] recommend the following procedure:

- **Classification** – Find attributes for sorting products into A, B and C groups,
- **Differentiation** – set differentiation policy for each class,
- **Allocation** – allocate (pricing) effort according to classification and differentiation.

Based on the Pareto principle and illustrated in Table 3.3, category A parts which make about 20% of the parts but are responsible for 80% of the inventory's value would be valuated via the following ASPV framework. Category B and C parts which only contribute a combined 20% of the value but 80% of quantity could be just valued via market data. If competition based pricing is not feasible because of a lack of data, there is no alternative option than using any internal price information as basis for the FMV.

Category	*Amount*	*Value*	*Method*
A	20%	80%	Valuation via ASPV framework
B	30%	10%	Competition Based Pricing
C	50%	10%	Competition Based Pricing

Table 3.3: Classification prior to choosing valuation method

Another perspective on product importance is by classification into stages of a product lifecycle. Blythe (2008) gives an overview of a product's lifecycle progress in four phases, beginning with an introduction phase where the amount of sold products grows slowly. The following growth phase is characterized by even more rising sales. Due to market saturation or introduction of improved products, sales reach a plateau and later on decline until obsolescence. Östlin et al. (2009) extend this concept and state that the number of new products and the demand for remanufactured products or spare parts is positively correlated but also delayed in time. As spare parts are

[5]Quoted from Grosfeld-Nir et al. (2007)

removed some time later from operations, the number of parts on the surplus market peaks after the maturity phase of new parts (see Figure 3.2).

To refine the allocation of parts, the consideration of the part's product lifecycle is helpful. Products wander through a product lifecycle which is characterized by a dependence from product value to demand and distribution on the market. The value of the parts is very dependent on the product lifecycle and therefore the demand of the underlying operations. Nevertheless, the product lifecycle of the new part can be an estimate of the demand for the spare part and therefore an indicator for the part's importance for valuation.

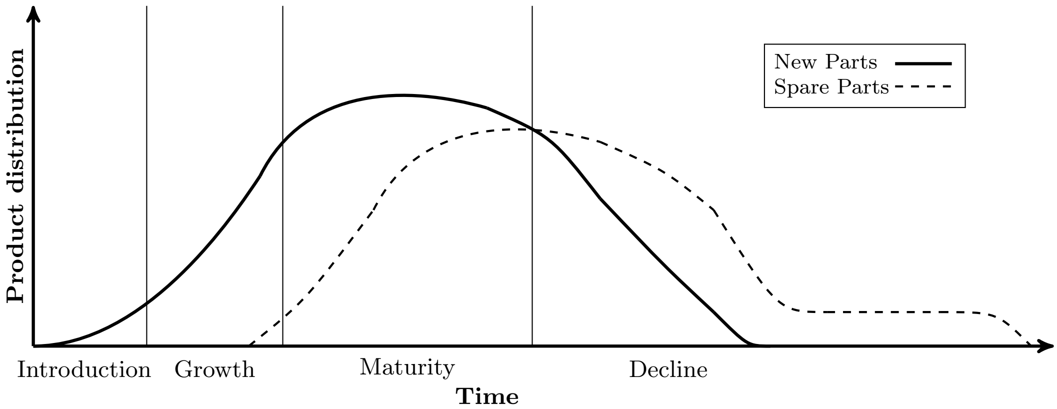

Figure 3.2: Product lifecycle based on Östlin et al. (2009)

As extension to the procedure of Ronen et al. (2005) and in the context of this scenario, the following procedure is therefore recommended:

- **Classification** – Find attributes for sorting products into A, B and C groups,
- **Differentiation** – set differentiation policy for each class,
- **Refine** – differentiate further by putting in perspective of product lifecycle,
- **Apply** – use ASPV framework for A parts and competition based pricing for B and C products,
- **Revise** – repeat process regularly based on improved data.

3.3.2 ASPV Framework

Even with a smaller number of parts for valuation, the process from raw data to information is a walk through the jungle of data analytics methods. An extension to the existing Extract Transformation and Load (ETL) process of Kimball and

Ross (2011) is proposed to address the particularities of pricing data and provide a system for method evaluation. The system architecture of the ASPV framework is visualized in Figure 3.3.

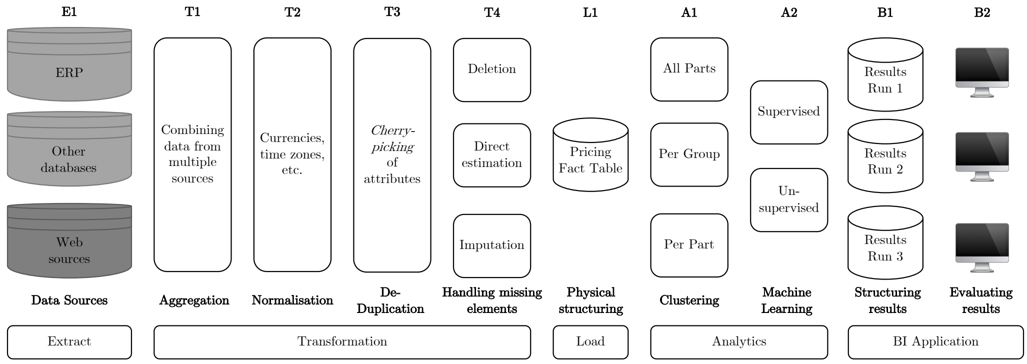

Figure 3.3: ASPV framework

The first step is the extraction of raw data from available data sources which include internal ERP systems, other internal databases and web sources from trading platforms (E1). Every data source contains complementary pricing information, that is why data pre-processing is of high importance prior to applying machine learning methods. Extracted data is now aggregated (T1) and subsequently normalized concerning currencies, time zones etc. (T2). The resulting datasets containing pricing information are now uniquely distinguishable by the combination of part number and date.

	partNo	partDate	partGroup	purchasePrice	offerPrice	sellingPrice	originalLaunchPrice	marketPrice
1	(string)	(datetime)	(string)					(double)
2	(string)	(datetime)	(string)				(double)	
3	(string)	(datetime)	(string)	(double)				
4	(string)	(datetime)	(string)		(double)			
5	(string)	(datetime)	(string)			(double)		

1	(string)	(datetime)	(string)	(double)	(double)	(double)	(double)	(double)

Figure 3.4: Deduplication of aggregated data

As there are usually multiple pricing points per part, duplicative datasets are combined (T3) as shown in the example in Figure 3.4. Multiple incomplete datasets are combined to one complete dataset. In case of multiple available data, aggregation or selection methods need to be selected according to the business process environment. Still missing elements need to be handled (T4). Two simple and widely used concepts are listwise and pairwise deletion. In case of listwise deletion a complete record is deleted when at least one attribute is missing. This might be the quickest and most effortless approach but also loses much data. Pairwise deletion on the other hand doesn't always ignore incomplete datasets but considers records with connected attributes. Consider the following example: the correlation between attribute A and attribute B is calculated. Pairwise deletion takes all records (including incomplete) into account where attribute A and B still exist. A detailed overview about handling missing data methods is given in Allison (2002). The result of the preceding transformation is a clean pricing fact table (L1) which is the basis for applying analytics methods. Now parts are combined in clusters of different size (A1). Clustering depends on the underlying characteristics of the parts, e.g. number and size of part groups. Machine Learning methods can then be applied on different aggregation levels (A2). Resulting is a FMV per part number and part date (B1). Finding a suitable machine learning method depends on the underlying data structure. The framework's performance can be evaluated by repeatedly running the process in multiple combinations from E1 to A2 and comparing the results (B1) in B2.

3.3.3 The Challenge of Validation

After collecting the results, the question of validation still stands out. In which combination of methods does the ASPV framework determine a good or correct FMV as target variable? Which FMV is the best or what is the target number of the FMV? We propose a two-step process for using the ASPV to find a FMV: First, automate the valuation process, second increase objectivity of the automated valuation to improve the status quo.

Supervised machine learning methods could be used to automate the valuation process. Predicted FMVs are validated with a subjective opinion of a domain expert. By that the prediction model will be trained to deliver the same quality as the domain expert.

To realize a substantial contribution to existing approaches the ASPV has to be more objective than the domain expert. The FMV should be in context to the input data which are the basis for FMV calculation. In reality, that might lead to choosing a value between reference values such as purchase and selling price. As this seems straightforward for validation, the question may arise, why taking all the effort and complicated models to calculate a FMV to, in the end, simply taking a mean value as reference. Achieving higher objectivity could be realized by deriving a model which leads to a FMV-based decision which contributes most to the overall business profitability. A FMV-based decision which achieves a more profitable business leads to the conclusion that this FMV was correctly determined regarding the definition given by Mercer and Brown (1999). This determination could be done by comparing time periods before using the FMV and periods while using the FMV for decision making. The challenges coming with this approach are:

1. Separating the FMV from its exogenous influence factors such as
   a) macroeconomic factors like overall economic growth,
   b) general profitability within the industry sector,
   c) general market conditions,
2. and then again, handling emerging missing values which might prevent this separation.

The third approach is not only the most complicated to realize but also the most objective.

3.4 Conclusion and Outlook

To find methods to determine a surplus spare part's value, existing pricing literature was analyzed. As Cost Based and Competitive Based Pricing are not applicable, Value Based Pricing is the right approach to find a Fair Market Value. The concept of the ASPV framework is an expansion of Value Based Pricing for manual value determination in primary markets to automated valuation in secondary markets.

The vast amount of different parts makes an automated part valuation inevitable. This is impeded by missing data and a non-transparent market structure. A generic ASPV framework for reliable value determination of spare parts enables companies to recognize a part's value right away and leads to a more sustainable use of parts. Due to the information gain, it would lead to a more efficient market overall. In the end, an automated value determination prevents bad part utilization decisions and could lead to a more efficient and sustainable use of surplus material. The concept for automated spare part valuation is a promising alternative for value determination and pricing in secondary markets and thus may serve as a foundation for building a generic surplus part trading platform to overcome market transparency issues if the obstacles of validation are overcome.

An early-stage blueprint of a framework for automated spare part valuation which serves as a guide for segmenting pricing data and aligning software layers from data aggregation to transformation, analytics and evaluation has been provided (research question 1). Future research has to deal with refining the alignment of the methods for segmentation of pricing data (research question 2) by benchmarking the approach against competitive concepts such as mean estimation, a k-nearest neighbors clustering, Bayesian Personalized Ranking or approaches without segmentation. Also the methods for predicting a FMV (research question 3) must be evaluated with real world data. Only experiments with real world data are able to give us an insight on whether a solution based on machine learning is even applicable to find a FMV.

References

Ackert, S. (2012). "Basics of Aircraft Market Analysis". In: *Aircraft Monitor*.

Allison, P. D. (2002). "Missing data. Series: Quantitative applications in the social sciences". In: *Thousand Oaks. CA*.

Anderson, J., Jain, D., and Chintagunta, P. (1992). "Customer value assessment in business markets: A state-of-practice study". In: *Journal of Business-to-Business Marketing* 1, pp. 3–29.

Anderson, J. and Narus, J. (1998). "Business marketing: understand what customers value". In: *Harvard business review* 76, pp. 53–67.

Anderson, J., Thomson, J. B., and Wynstra, F. (2000). "Combining value and price to make purchase decisions in business markets". In: *International Journal of Research in Marketing* 17.4, pp. 307–329. ISSN: 0167-8116.

Blumberg, D. F. (2004). *Introduction to Management of Reverse Logistics and Closed Loop Supply Chain Processes*. CRC Press. ISBN: 9780203502938. DOI: 10.1201/b12427.

Blythe, J. (2008). *Essentials of marketing*. Pearson Education.

Eaves, A. H. C. and Kingsman, B. G. (2004). "Forecasting for the ordering and stock-holding of spare parts". In: *Journal of the Operational Research Society* 55.4, pp. 431–437. DOI: 10.1057/palgrave.jors.2601697.

Grosfeld-Nir, A., Ronen, B., and Kozlovsky, N. (2007). "The Pareto managerial principle: when does it apply?" In: *International Journal of Production Research* 45.10, pp. 2317–2325.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). "The Elements of Statistical Learning: Data Mining, Inference, and Prediction". In: *Springer Series in Statistics*.

Hinterhuber, A. (2004). "Towards value-based pricing—An integrative framework for decision making". In: *Industrial Marketing Management* 33.8, pp. 765–778. ISSN: 00198501. DOI: 10.1016/j.indmarman.2003.10.006.

— (2008). "Customer value-based pricing strategies: why companies resist". In: *Journal of Business Strategy* 29.4, pp. 41–50. ISSN: 0275-6668. DOI: 10.1108/02756660810887079.

Kelly, D. B. (2008). "Forecasting aircraft values: An appraiser's perspective". In: *Air-finance annual*.

Kimball, R. and Ross, M. (2011). *The data warehouse toolkit: the complete guide to dimensional modeling*. John Wiley & Sons. ISBN: 1118082141.

Mercer, Z. C. and Brown, T. S. (1999). "Fair Market Value vs. The Real World". In: *Business Valuation Review* 18.1, pp. 16–25.

Naisbitt, J. and Cracknell, J. (1984). *Megatrends: Ten new directions transforming our lives*. Report. Warner Books New York.

Neely, A. and Jarrar, Y. (2004). “Extracting value from data–the performance planning value chain”. In: *Business Process Management Journal* 10.5, pp. 506–509.

Östlin, J., Sundin, E., and Björkman, M. (2009). “Product life-cycle implications for remanufacturing strategies”. In: *Journal of cleaner production* 17.11, pp. 999–1009.

Ronen, B., Spiegler, I., and Kozlovsky, N. (2005). “The use, misuse and abuse of the Pareto principle”. In: *Working paper*.

Voß, S. and Lessmann, S. (2013). “Resale Price Prediction in the Used Car Market”. In: *Tristan Symposium VIII*.

Wickboldt, C. and Kliewer, N. “A framework for value-based pricing in aircraft surplus markets”. In: *Multikonferenz Wirtschaftsinformatik (MKWI) 2016*. ISBN: 978-3-86360-132-4. URL: http://www.mkwi2016.de/download/MKWI2016_Research-in-Progress-Poster-Band.pdf.

Kapitel 4

Blockchain zur dezentralen Dokumentation von Werkstattereignissen in der Luftfahrtindustrie

Zusammenfassung

Zur Wartung einer Flotte, etwa für den Flug-, See-, oder Schienenverkehr, sind Maintenance Repair und Overhaul (MRO, deutsch: Wartung, Reparatur und Betrieb) Unternehmen auf eine sichere und schnelle Ersatzteilversorgung angewiesen. Um Ausfallzeiten zu minimieren, werden defekte Teile kurzfristig ausgetauscht. Anschließend werden die ausgetauschten Teile repariert oder überholt. Überschüssige Teile werden in einem Sekundärmarkt gehandelt. Dem Teilekreislauf werden neue Teile zugeführt und letztlich obsolete Teile zur Verschrottung entnommen. In diesem Beitrag liegt der Fokus auf der Luftfahrtindustrie.

Eine besondere Bedeutung kommt den sicherheitsrelevanten Ersatzteilen zu, die wiederum eine lückenlose Lebenszyklusdokumentation aller Werkstattereignisse aufweisen müssen, um weiter verwendbar zu bleiben. Die aktuelle Dokumentationspraxis ist durch manuelle Erfassung und analoge Speicherung geprägt, was zu einer hohen Unsicherheit durch manipulierbare, fehleranfällige und unvollständige Dokumentation und somit unnötigen Verschrottung von wertvollen Ersatzteilen führt. Ein funktionierender Sekundärmarkt ist vor diesem Hintergrund nur sehr eingeschränkt möglich.

Eine Möglichkeit, diese Intransparenz zu adressieren, ist die Dokumentation von Werkstattereignissen in einer zutrittsbeschränkten (engl. permissioned) Blockchain. In diesem Beitrag wird ein theoretisches Transaktionsflussmodell basierend auf dem Hyperledger Fabric Framework vorgeschlagen, das die genannten besonderen Anforderungen der Geschäftsprozesse der MRO-Anbieter in der Luftfahrtindustrie berücksichtigt. Die praktische Relevanz der Problemstellung und die Anwendbarkeit der Lösung wird durch eine enge Zusammenarbeit mit einem der führenden MRO-Anbieter (anonym) sichergestellt.

Schlüsselwörter Blockchain · Hyperledger · Smart Contract · Ersatzteil · Luftfahrtindustrie · Werkstattereignis

Blockchain for Decentralized Documentation of Workshop Events in the Aviation Industry

Abstract

For the maintenance of a fleet, such as for air, sea or rail traffic, maintenance repair and overhaul (MRO) companies depend on a secure and fast spare parts supply. To minimize downtime, defective parts are replaced at short notice. Subsequently, the replaced parts are repaired or overhauled. Surplus parts are traded in a secondary market. The parts cycle is fed new parts and finally obsolete parts are removed for scrapping. This article focuses on the aviation industry.

Of particular importance are safety-relevant spare parts, which must have a complete back-to-birth documentation of all workshop events in order to remain usable. The current documentation practice is characterized by manual capture and analog storage, resulting in high uncertainty due to manipulatable, error-prone and incomplete documentation and thus unnecessary scrapping of valuable spare parts. A functioning secondary market is only possible to a very limited extent with this background.

One possibility to address this lack of transparency is the documentation of workshop events in a permissioned blockchain. This paper proposes a theoretical transaction flow architecture based on the Hyperledger Fabric Framework that addresses the specific needs of the business processes of MRO companies in the aviation industry. The practical relevance of the problem and the applicability of the solution is ensured by close cooperation with one of the leading MRO providers (anonymous).

Keywords Blockchain · Hyperledger · Smart Contract · Spare Part · Aerospace Industry · Maintenance Event

Contents

4.1 Einleitung

Der Markt für Flugzeugersatzteile zeichnet sich durch eine hohe Anzahl heterogener Teile, einem hohen Umschlag verschiedener Teile pro Tag, hohen Sicherheitsanforderungen an den Umgang mit den Teilen und gleichzeitig einer geringen Zahl von Marktteilnehmern aus. Alleine die britische Royal Airforce betreut rund 685.000 Artikel mit einem Wert von über 2Mrd. GBP (Eaves und Kingsman, 2004). Das größte Verkehrsflugzeug Airbus A380 besteht beispielsweise aus rund 2,5Mio. Einzelteilen[1]. Die Fluggesellschaften übernehmen die Wartung der Flugzeuge oft nicht selbst, sondern schließen Wartungsverträge mit sogenannten Maintenance Repair and Overhaul (MRO)-Anbietern ab. Die Aufgabe der MRO-Anbieter liegt wiederum darin, eine schnelle, sichere und effiziente Wartung sicherzustellen. Eine Besonderheit stellen dabei sicherheitsrelevante Teile dar. Um diese Teile verwenden zu können, ist eine lückenlose Dokumentation bis zum Herstelldatum notwendig[2] . Die derzeit analoge Dokumentation von Werkstattereignissen mit Zertifikaten stößt an ihre Grenzen, wenn einzelne Zertifikate im Verlauf des Teilelebenszyklus verloren gehen. Treten Unregelmäßigkeiten bei der Dokumentation dieser Teile auf, werden diese sicherheitshalber verschrottet, was für das Unternehmen mit hohen und eigentlich vermeidbaren Verlusten einhergeht. Ein Sekundärmarkt ist aufgrund dieser Unsicherheiten nur sehr eingeschränkt vorhanden. Experten aus der Flugzeug-MRO-Branche bestätigen, dass die analoge Dokumentation maßgeblich für diese Unsicherheiten verantwortlich ist. Die Digitalisierung der Prozesse wird bereits in der Branche diskutiert. Aufgrund dessen, dass die Luftfahrtindustrie zwar global agiert, jedoch regional reguliert wird, existiert keine globale Autorität, die eine zentrale digitale Plattform zur Dokumentation der Werkstattereignisse und Teilehistorie durchsetzen könnte. Um diese Lücke zu schließen, werden in diesem Beitrag die in Tab. 4.1 dargestellten Anforderungen an eine digitale Lösung adressiert.

Dieser Beitrag untersucht, inwiefern Fortschritte in verteilten Systemen und Datenbanken helfen können, das Problem der fehlenden zentralen Autorität zu umgehen und eine vertrauenswürdige und sichere Plattform zur Dokumentation bereitzustellen. Dazu wird insbesondere auf die jüngsten Entwicklungen im Bereich Blockchain und sogenannter *Smart Contracts* eingegangen. Smart Contracts bieten aufgrund der Tatsache, dass sie dezentral und trust-free funktionieren die Möglichkeit, etablierte Vermittler obsolet werden zu lassen und befähigen dadurch zwei Parteien zum direkten Austausch (Peer-to-Peer oder P2P) (Glaser, 2017).

Hevner et al., 2004 halten fest, dass das effektive Design eines neuen Artefakts sowohl Wissen über Anforderungen und Beschränkungen der Anwendungsdomäne als auch technische Eigenschaften der Lösungsdomäne erfordert. Daher wird im Sinne

[1] Airbus Press Office 2017: Facts & Figures, Juni 2017.

[2] Diese Anforderungen ergeben sich durch die jeweiligen Luftfahrtbehörden (z.B. EASA in Europa oder FAA in den USA).

Code	Kategorie	Beschreibung
A1	Geschwindigkeit	Eine hohe Geschwindigkeit, um den Prozess der Teileverwendung nicht zu behindern
A2	Zusammenführung	Die Möglichkeit Ereignisse aus verschiedenen Quellen geordnet zusammenzuführen
A3	Persistenz	Eine hohe Geschwindigkeit, um den Prozess der Teileverwendung nicht zu behindern
A4	Vertrauensfrei (*trust-free*)	Eine fälschungssichere Lösung, die gegebenenfalls ohne Vertrauen funktioniert
A5	Zutrittsbeschränkt (*permissioned*)	Eine beschränkte Sichtbarkeit, das heißt eine mit Zutrittsberechtigungen versehene Lösung
A6	Abbildung Geschäftsprozess	Eine Lösung, die in der Lage ist, Regeln des Geschäftsprozesses abzubilden

Tabelle 4.1: Anforderungen zur dezentralen und digitalen Dokumentation von Werkatttereignissen

der allgemein zitierten Phrase „blockchain is an innovative technology in search of use cases“ (Glaser, 2017), also einer innovativen Technologie auf der Suche nach Anwendungsfällen, ein Modell für eine dezentrale Dokumentation von Werkstattereignissen auf Basis der Blockchain, zugeschnitten auf die Besonderheiten des Prozesses für sicherheitsrelevante Flugzeugersatzteile, vorgeschlagen. In diesem Zusammenhang wird unter anderem die Frage adressiert, ob eine durch dezentrale Smart Contracts ermöglichte P2P-Dokumentation dafür hilfreich ist. Das Transaktionslussmodell lehnt sich an existierende Ansätze zur dezentralen Nachverfolgung von Gütern (Abeyratne und Monfared, 2016; Christidis und Devetsikiotis, 2016) an. Das Forschungsprojekt zu diesem Beitrag verwendet Informationen von Domänenexperten aus der Flugzeug-MRO-Branche, um so die Problemstellungen in einem realen Anwendungsfall zu adressieren und somit die praktische Relevanz der Themenstellung sicherzustellen. Das Modell stützt sich auf erste Erkenntnisse zur Anwendbarkeit von Smart Contracts in einer Blockchain (Glaser, 2017), erweitert diese zu einem nachvollziehbaren und in der Praxis evaluierbaren Artefakt im Sinne von Hevner et al. (2004) und kommt der Aufforderung nach Forschung in der Schnittstelle zwischen Technologie, Märkten und Geschäftsmodellen (Nofer et al., 2017) nach, einen praxisnahen Beitrag zur noch kleinen Wissensbasis des Themenbereichs Blockchain zu liefern und damit das Problem der Informationsasymmetrien zu lösen.

Der nächste Abschnitt des Beitrags gibt einen Überblick über die technologischen Grundlagen zu Blockchain und Smart Contracts, den darin genutzten Konsensverfahren und die Plattform Hyperledger, ein Framework zur Nutzung bestehender Verfahren. In Abschn. 4.3 folgt die Einführung in die Spezifika des Dokumentationsprozesses für sicherheitsrelevante Flugzeugersatzteile und die Vorstellung eines Modells zur Abbildung der Transaktionslüsse unter der Verwendung von Hyperledger Technologie. Abschn. 4.4 fasst den Beitrag zusammen, diskutiert den vorgestellten Ansatz sowie seine praktischen Implikationen und gibt einen Ausblick auf Validierungsmöglichkeiten.

4.2 Technologische Grundlagen

Ein Smart Contract fungiert als vertrauenswürdige verteilte Anwendung und hat die Aufgabe, sichere (Geschäfts-) Beziehungen innerhalb eines verteilten Systems abzubilden. Unabhängig von der Blockchain-Technologie wurde die Idee für digitale Verträge in Szabo (1997) vorgestellt. Die Digitalisierung der Verträge ermöglicht eine automatische Ausführbarkeit dieser, zum Beispiel eine automatisierte Bestellung bei niedrigem Lagerbestand zu dem im Smart Contract vereinbarten Preis. Ein aktueller Blockchain-basierter Ansatz, der Smart Contracts in den Vordergrund stellt, ist Ethereum (vgl. Buterin (2014)). Dieser Ansatz ähnelt dem bekannten *State Machine Approach*, einer allgemeinen Methode zur Implementierung fehlertoleranter Dienste in verteilten Systemen (Schneider, 1990). Eine *State Machine* (deutsch: Zustandsmaschine) besteht sowohl aus Zustandsvariablen, die den Zustand festhalten als auch aus Kommandos, die den Zustand ändern. Eine Komponente (jedwede Anwendung, ob Smart Contract oder nicht) gilt als fehlerhaft, wenn ihr Verhalten nicht mehr mit ihrer Spezifikation übereinstimmt. Die Verteilung auf mehrere Systeme hat den Zweck, dass im Fall eines Fehlers nur der jeweilige Teil des Systems betroffen ist und auf diese Weise die Fehlertoleranz des Gesamtsystems steigt. Das Verfahren baut unter anderem auf Methoden auf, die Konsens[3] in verteilten Systemen erreichen.

Bei traditionellen verteilten State Machines läuft gleichzeitig nur eine Anwendung auf dem verteilten System. Dagegen ermöglichen Blockchains die gleichzeitige Ausführung mehrerer unterschiedlicher Anwendungen, die dynamisch von jedem Stakeholder erstellt und verändert werden können. Der Code einer solchen Anwendung ist durch das Fehlen eines zentralen Intermediärs nicht vertrauenswürdig und damit potenziell schädlich. Durch die Tatsache, dass Transaktionen dezentral stattfinden, gilt es bei der Zusammenführung zu einem einheitlichen Stand, einen Konsens zu finden, um die in Abschn. 4.1 dieses Beitrags definierte Anforderung A2 (Zusammenführung) zu erfüllen. Diese Konsensfindung wird in einigen Verfahren (z.B. Proof of Work) als Mining und der konsenssuchende Knoten als Miner bezeichnet. Die einzelnen Elemente und Charakteristiken von Blockchain-Systemen werden nachfolgend erläutert.

[3] Eine detaillierte Übersicht über Konsensverfahren in Blockchains findet sich in Zheng et al. (2017).

4.2.1 Blockchain

Blockchain-Systeme sind ein innovativer Ansatz zur Speicherung von Daten und zeichnen sich nach Zheng et al. (2017) unter anderem durch die folgenden Charakteristika aus.

- *Dezentralisierung*: In der Blockchain werden vertrauensbildende Intermediäre durch einen Konsensmechanismus ersetzt, so dass eine Blockchain trust-free ist. Die Datenstände liegen verteilt auf den Knoten des Netzwerks.
- *Persistenz*: Ist der Konsens einmal gefunden, ist es nicht möglich, den Datenstand rückwirkend zu verändern. Blocks mit Transaktionen, die dem gefunden Konsens widersprechen, werden nicht akzeptiert.
- *Prüfbarkeit*: Transaktionen bauen auf dem aktuellen Informationsstand auf und aktualisieren diesen anschließend. Auf diese Weise ist die Verkettung der Informationen nachvollziehbar.

Blockchain-Systeme lassen sich nach Zheng et al. (2017) hinsichtlich der Zutrittsbeschränkungen einordnen. In einer öffentlichen Blockchain gibt es keine Zutrittsbeschränkungen, jeder kann an der Konsensfindung teilnehmen. Eine private Blockchain agiert in einem zentralisierten Netzwerk, da sie vollständig von Knoten einer einzigen Organisation kontrolliert wird. Eine Konsortium-Blockchain, in der die Konsensfindung über Knoten ausgewählter Organisationen stattfindet, ist teilweise dezentralisiert.

Einen solch begrenzten Kreis von teilnehmenden Organisationen gibt es auch bei der Verwendung und dem Handel von sicherheitsrelevanten Flugzeugersatzteilen. Gegenwärtig entwickelt Hyperledger (Cachin, 2016; Group, 2018) Blockchain-Frameworks für Geschäftskonsortien. Ethereum hat ebenfalls Werkzeuge zur Entwicklung von Konsortium-Blockchains bereitgestellt (Ethereum, 2017). Umgesetzte Fallbeispiele für das Tracking von Gütern über eine Blockchain sind Everledger und Provenance. Everledger[4] ist ein Projekt zur Nachverfolgung von Diamanten entlang der Wertschöpfungskette auf Basis von Hyperledger. Damit soll nachvollziehbar sichergestellt werden, dass keine sogenannten Blut-Diamanten in den Handel gelangen. Provenance[5] ist ein Unternehmen, das Anwendungen zur Nachverfolgung von Gütern, insbesondere in der Lebensmittelindustrie, auf Basis von Ethereum-Blockchains entwickelt.

Ist ein Konsens gefunden, wird dieser über eine sogenannte aktive Replikation verteilt. Dabei werden die Transaktionen, für die ein Konsens gefunden wurde, nach

[4] `https://www.everledger.io`
[5] `https://www.provenance.org/whitepaper`

diesem geordnet und die komplette Kette der Transaktionen (Ledger) gleichermaßen an alle Teilnehmer des Netzwerks (Peers) propagiert. Jeder Peer führt diese Transaktion sequentiell aus. Damit der Stand der Peers konsistent bleibt, muss die Ausführung deterministisch sein, das heißt bei derselben Anwendung und demselben Zustand wird von jedem Peer dieselbe Zustandsaktualisierung erzeugt (Renesse und Guerraoui, 2010). Androulaki et al. (2018) nennen dies Order-Execute-Architektur (deutsch: Ordnen-Ausführen). Durch diese Architektur entstehen Nachteile, die in einer Lösung für die Spezifika der Dokumentation sicherheitsrelevanter Ersatzteile nicht annehmbar sind. Dazu zählen:

- Die sequenzielle Ausführung aller Transaktionen durch alle Peers schränkt die Performance ein, dies widerspricht Anforderung A1 (Geschwindigkeit).
- Die Tatsache, dass jede Transaktion auf jedem Peer ausgeführt wird, verhindert eine granulare Zugriffsberechtigung der Peers auf einzelne Transaktionen, dies widerspricht Anforderung A5 (Zutritt).
- Das Vertrauensmodell der Transaktionsvalidierung wird durch das Konsensverfahren bestimmt und kann nicht an die Anforderungen des Smart Contract angepasst werden, dies könnte zu Problemen hinsichtlich Anforderung A6 (Abbildung Geschäftsprozess) führen.
- Die Tatsache, dass Transaktionen deterministisch sind kann ebenfalls in Einschränkungen bei der Umsetzung von Anforderung A6 resultieren.

Diese Nachteile werden durch die nachfolgend beschriebene Plattform adressiert.

4.2.2 Hyperledger Fabric

Hyperledger Fabric ist ein modulares und erweiterbares Open-Source-System zum Bereitstellen und Betreiben von permissoned Blockchains mit der Möglichkeit der Verwendung verschiedener Konsensfindungsverfahren. Es ist eines der von der Linux Foundation betriebenen Hyperledger-Projekte[6] (Androulaki et al., 2018). Die Innovation von Fabric liegt in der Anpassung der Systemarchitektur von Order-Execute zu Execute-Order-Validate (Ausführen-Ordnen-Validieren). Der Transaktionsfluss wird in drei Schritte aufgeteilt, die auf unterschiedlichen Peers im Netzwerk ausgeführt werden können. Die Geschäftsprozesslogik wird in Form von Smart Contracts über den *Chaincode* als Software implementiert.

[6] `https://www.hyperledger.org/projects`

Es wird zwischen drei verschiedenen Peer-Typen unterschieden:

- Committer: Startet Transaktionen, führt den Ledger und den Status.
- Endorser: Erhält Transaktionsvorschläge, gibt diese auf Basis des Chaincode frei oder lehnt sie ab.
- Orderer: Nimmt freigegebene Transaktionen in die Blockchain auf und ordnet sie.

Zuerst wird die Transaktion vom Committer ausgeführt und vom Endorser auf Basis des Chaincode simuliert. Entspricht die Ausführung der Transaktion theoretisch dem festgelegten Geschäftsprozess, wird die Transaktion vom Endorser freigegeben (Anforderung A6). Anschließend und unabhängig von der Art der Transaktion erfolgt die Ordnung durch den Orderer auf Basis eines Konsensfindungsverfahrens. Nach der Ordnung wird die Transaktion validiert. Die Besonderheit liegt dabei in der Tatsache, dass die Transaktion zuerst ausgeführt und erst anschließend geordnet und validiert wird. Mittels der Ausführung der Transaktion durch eine Teilmenge der Knoten ist eine parallele Ausführung von verschiedenen Transaktionen möglich. Dadurch, dass die Ordnung und Validierung erst anschließend erfolgt, können Regeln von Untermengen von Peers beachtet und damit beziehungsspezifisch angewandt werden (Anforderung A6). Durch die Verteilung der Aufgaben sowie der Ausführung der Schritte auf Teilen des Netzes und der sich daraus ergebenen Parallelität wird zudem eine höhere Ausführungsgeschwindigkeit erreicht (Anforderung A1). Ein *Membership Service Provider* ist verantwortlich für die Zuordnung von Peers zu kryptografischen Identitäten und regelt damit den Zugang zum Netzwerk (Anforderung A5).

4.2.3 Abgrenzung der Technologien

Zusammenfassend zeigt Tab. 4.2 die Abgrenzung der vorgestellten Ansätze. Systeme, die derzeit zur Nachverfolgung von Gütern genutzt werden, z.B. ERP-Systeme, basieren auf zentral bereitgestellten Datenbanken. Wie in der Einleitung dargestellt, wurden in der Praxis solche Lösungen bereits verworfen, da keine zentrale Autorität zur Bereitstellung existiert. State Machines lagern die Daten dezentral, jedoch gibt es weiterhin eine zentrale Instanz, die den Datenstand zusammenführt und nach außen kommuniziert. Mit Bitcoin wurde das System der Blockchain eingeführt. Bitcoin spezialisiert sich auf die Bereitstellung einer Krypto-Währung und bietet keine Kapazität für die Abbildung von Geschäftsprozessen in Smart Contracts. Ethereum bietet die Kapazität der öffentlichen Abbildung von Smart Contracts. Die Konsensfindung findet auf Ledger-Ebene statt, also über die gesamte Kette. Hyperledger Fabric bietet als modular verwendbares System die Möglichkeit einzelne Stakeholder am Informationsaustausch teilhaben zu lassen. Die Konsensfindung findet im Gegensatz zu Bitcoin und Ethereum nicht gleichermaßen über die gesamte Kette von Transaktionen sondern spezifisch pro Transaktion (demnach auf Transaktionsebene) statt, so dass die Charakteristika der jeweiligen Geschäftsbeziehung berücksichtigt werden.

Ansatz	Dezentral	Trust-free	Kapazität	Permissioned	Konsensfindung
ERP-System	✗	✗	✓	✗	✗
State-machine replication (Schneider, 1990)	✓	✗	✓	✗	✗
Bitcoin (Nakamoto, 2008)	✓	✓	✗	✗	Ledger-Ebene
Ethereum (Buterin, 2014)	✓	✓	✓	✗	Ledger-Ebene
Hyperledger Fabric (Androulaki et al., 2018)	✓	✓	✓	✓	Transaktionsebene

Tabelle 4.2: Abgrenzung zentraler und dezentraler Systeme

4.3 Aircraft Surplus Part Network

Nach Hevner et al. (2004) sind Beiträge zur Forschung der Wirtschaftsinformatik dann sinnvoll, wenn sie auf eine Geschäftsanforderung angewandt werden können. Die Autoren greifen auf existierende Methoden zurück, um gestaltungsorientiert einen digitalen Transaktionsfluss auf Basis einer Fabric Blockchain, eine Lösung für die Spezifika der sicherheitsrelevanten Flugzeugersatzteile, vorzuschlagen. Dazu wird der darunterliegende Geschäftsprozess und anschließend ein Transaktionsflussmodell vorgestellt, das den aus dem Geschäftsprozess resultierenden Herausforderungen begegnet.

4.3.1 Einführung in den Verwendungsprozess von Flugzeugersatzteilen

Um teure Ausfallzeiten zu vermeiden, sind Fluggesellschaften auf eine reibungslose und sichere Ersatzteilversorgung angewiesen. Teile werden zeitnah gegen lagernden Ersatz ausgetauscht (sog. Pull and Replace). Überschüssige (sog. Surplus)-Teile werden auf Sekundärmärkten gehandelt. Die Teile bewegen sich in einer „Alternative high tech closed loop supply chain“ (Blumberg, 2004), der neue Teile zugeführt und wertlose Teile zur Verschrottung entnommen werden[7]. Eine besondere Teilegruppe sind sicherheitsrelevante Teile. Diese sind im Durchschnitt vier

[7] Die Zusammenhänge zwischen Teilewert und Verwendung werden ausführlich in Wickboldt und Kliewer (2018b) beschrieben.

Mal so hoch bewertet wie die restlichen Ersatzteile. Für diese Teilegruppe ist eine lückenlose Dokumentation aller Werkstattereignisse im Lebenszyklus, eine sogenannte Back-to-Birth-Dokumentation, notwendig, um eine reibungslose Funktion des Teils sicherzustellen (nachfolgend B2B-Teile). Abb. 4.1 zeigt die Abschnitte des Teile-Verwendungsprozesses in der Luftfahrtindustrie, die durch eine lückenhafte B2B-Dokumentation betroffen sind (blau/grau hervorgehoben). Zum einen werden B2B-Teile, die Lücken in der Dokumentation aufweisen, sicherheitshalber verschrottet und stehen für ein Replace nicht mehr zur Verfügung. Zum anderen verzichten einige MRO-Anbieter auf die Beschaffung von B2B-Teilen aus dem Sekundärmarkt, da die Historie im Zweifel nicht nachvollziehbar ist. Sowohl durch die Verschrottung als auch durch das Fehlen der B2B-Teile auf dem Sekundärmarkt entstehen Ineffizienzen.

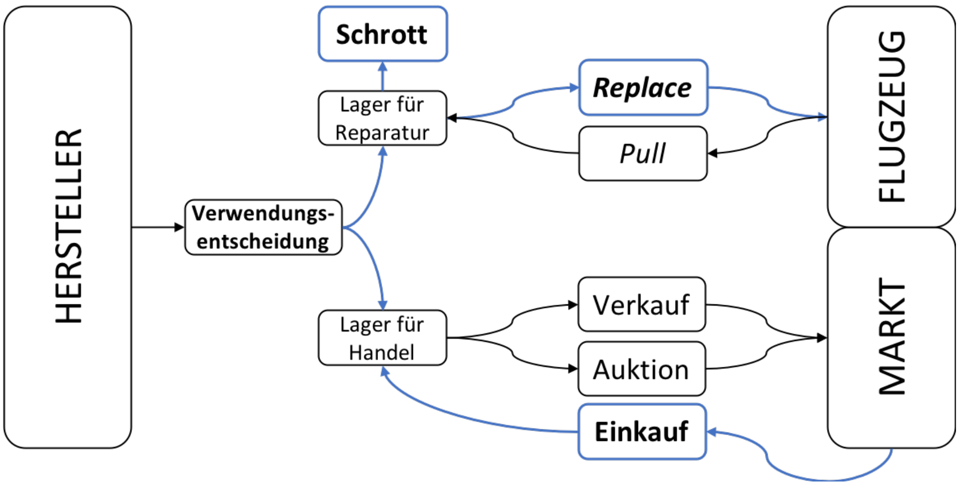

Abbildung 4.1: B2B-Teile im Surplus-Verwendungsprozess

Einen Beitrag zur Unsicherheit bei der Dokumentation der Teilehistorie leistet der derzeit analoge Prozess der Verwaltung von Reparaturzertifikaten. Diese Zertifikate werden dezentral erstellt, auf Papier ausgedruckt und anschließend dem Teil beigelegt. Geht eines dieser Zertifikate verloren, ist das Teil wertlos. Zudem beeinflusst die Reputation der Zertifikat-ausstellenden Werkstatt den Wert des Teiles, was Anreize zur Manipulation geben könnte.

4.3.2 Transaktionsflussmodell auf Basis von Hyperledger Fabric

Diesen Herausforderungen kann durch den Einsatz einer digitalen Plattform begegnet werden. Die Bereitstellung dieser einheitlichen Plattform an zentraler Stelle erfordert ein großes Maß an Vertrauen der Stakeholder an die Organisation, die alle

Daten sammelt, um diese Plattform betreiben zu können. Des Weiteren würde diese zentrale Instanz eine signifikante Macht in der Branche gewinnen, was zu Korruption und Missbrauch beim Umgang mit diesen Daten führen könnte. Nicht zuletzt wäre ein zentral bereitgestelltes System ein sogenannter *Single Point of Failure*, was das Gesamtsystem der Ersatzteilverwendung verwundbar werden lässt. Das nachfolgende dezentrale Blockchain-basierte Transaktionsflussmodell adressiert diese Herausforderungen und die daraus abgeleiteten Anforderungen A1 bis A6.

Nachfolgend wird der Transaktionsfluss im Detail erläutert, darunter die verschiedenen Stakeholder im Netzwerk, die Art und Weise, wie diese Stakeholder Zugriff und Einfluss auf das Netzwerk haben, und wie eingegebene Informationen geordnet, validiert und dem Netzwerk zur Verfügung gestellt werden. Abeyratne und Monfared (2016) haben ein ähnliches Konzept zur Nachverfolgung von Ereignissen in der Herstellung von Pappkartons auf Basis einer Ethereum-Blockchain entwickelt. Allerdings bietet die Ethereum-Blockchain kein Berechtigungssystem. Aus diesem Grund orientieren sich die Autoren in der Beschreibung zwar am Aufbau der in Abeyratne und Monfared (2016) vorgestellten Architektur, entwickeln jedoch das neue Modell auf Methoden von Hyperledger Fabric.

Das Teil (T) hat einen Zustand (Z), wird durch eine Werkstatt (W) repariert und erhält anschließend den Zustand (Z'). Der Reparaturvorgang (R) und Zustand Z' werden durch ein Zertifikat (C) bescheinigt. Das Zertifikat C enthält Informationen über das Teil T selbst (Herstellungsdatum, Flugstunden, Zustand Z), und über die durchgeführte Reparatur R durch Werkstatt W zum Zustand Z'. Nach der Reparatur R wird das Zertifikat C an die Blockchain angehängt, so dass die Historie der Werkstattereignisse und Zustände für alle berechtigten Stakeholder sichtbar ist. Die Stakeholder können Händler, Regulatoren und Werkstätten sein. Folgende Annahmen werden dabei getroffen:

- Der Client (hier W) ist registriert und hat von einer Zertifizierungsstelle kryptografisches Material zur Unterzeichnung der Transaktion und Authentifizierung in der Blockchain erhalten.
- Die sogenannte *Endorsement Policy* (Richtlinie zur Billigung des Vorgangs) besagt, dass drei Stakeholder, davon mindestens zwei fremde Peers die Transaktion freigeben müssen.
- Der Chaincode enthält Regeln, wie z.B.:
  - W hat Berechtigung für R und kann C ausstellen,
  - T ist nur für bestimmte R geeignet. T muss daher in der passenden W repariert werden,
  - die Historie von T ist mit dessen Parametern konform,

- alle 1,5 Jahre erfolgt ein C-Check des Flugzeugs; alle 6 Jahre oder nach 30.000 Flugstunden erfolgt ein D-Check.[8]

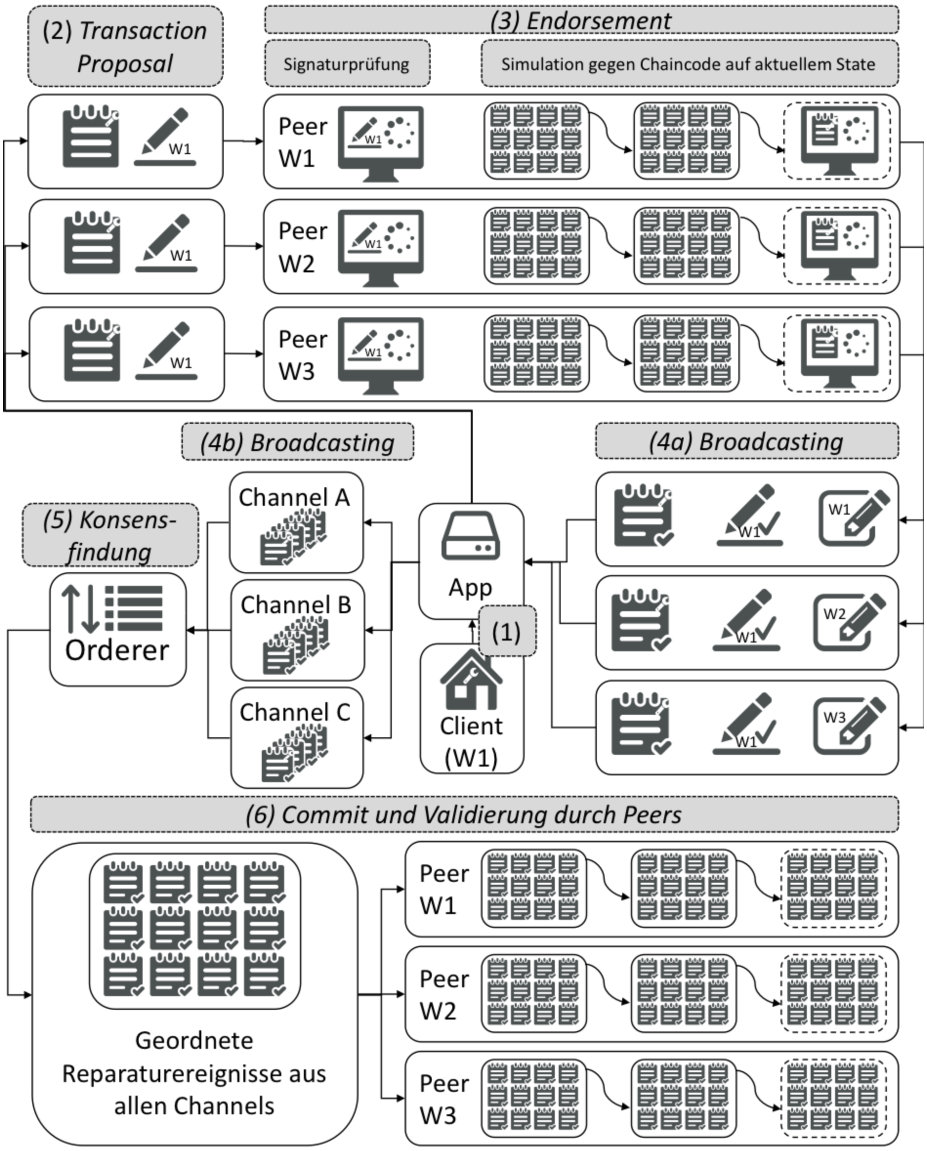

Abbildung 4.2: Transaktionsflussmodell für die Dokumentation von Werkstattereignissen in Hyperledger Fabric

[8]Die Vorgänge C-Check und D-Check werden unter anderem hier näher erläutert: `https://www.lufthansatechnik.com/de/aircraft-maintenance`.

Ein konzeptioneller Transaktionsfluss[9] für die Hinterlegung des Zertifikats für ein Werkstattereignis besteht aus sechs Schritten und ist in Abb. 4.2 dargestellt. Der Client (W1) hinterlegt zunächst das Zertifikat in einer Webanwendung (1). Diese bildet über die Peers die Schnittstelle zur Blockchain. Die hinterlegte Dokumentation wird in der App mit der digitalen Signatur von W1 versehen und als Transaction Proposal an die Endorsing Peers weitergeleitet (2). In dem Modell prüfen gemäß Endorsement Policy insgesamt drei Endorsing Peers, vertretend für die Clients W1, W2 und W3, die Transaktion (3). Dabei wird zuerst die Signatur von W1 überprüft. Anschließend wird die Transaktion gegen den aktuellen State unter Berücksichtigung der im Chaincode hinterlegten Regeln simuliert. Die Simulation ist erfolgreich, wenn das Hinterlegen des neuen Zertifikats mit den Regeln im Chaincode konform ist. Nach erfolgreicher Prüfung gelangen die Ergebnisse des Endorsement, signiert durch den jeweiligen Endorsing Peer, zurück an die App (4a). Die App verteilt die Ergebnisse in die verschiedenen Channel (4b). Die ungeordneten Transaktionen aller Channel werden vom Orderer sortiert und diese Sortierung durch ein Konsensfindungsverfahren legitimiert (5). Die sortierten Transaktionen werden über einen Commit an den Ledger angehängt und an die jeweiligen Peers zurückgegeben. Diese validieren den State und aktualisieren ihn lokal (6).

Zusammengefasst wird der Dokumentationsprozess durch die Verwendung von Hyperledger Fabric nicht nur digitalisiert, sondern die Befolgung der Regularien sichergestellt. Das Zertifikat zum Werkstattereignis wird weiterhin dezentral erstellt. Die automatisierte Berücksichtigung des Chaincode nach dem vorgestellten Transaktionsflussmodell trägt dazu bei, dass eventuelle Lücken bei der Überprüfbarkeit der Einhaltung der Regularien automatisch aufgedeckt werden und es so unmöglich ist nicht-regelkonforme Werkstattereignisse zu dokumentieren.

4.4 Kritische Diskussion und praktische Implikationen

Der in der Einleitung beschriebene derzeitige analoge Prozess zur Dokumentation von Werkstattereignissen birgt die Gefahr von Lücken und Manipulation. Die daraus entstehenden Herausforderungen motivieren zur Entwicklung einer digitalen Lösung zur schnellen (A1), geordneten (A2), persistenten (A3), fälschungssicheren (A4), mit Zugriffsbeschränkungen versehenen (A5) und auf die Regeln des Geschäftsprozesses abgestimmten (A6) Dokumentation. In Abschn. 4.3 dieses Beitrags wurde ein Transaktionsflussmodell zur Speicherung von Zertifikaten für Werkstattereignisse im Laufe des Lebenszyklus eines sicherheitsrelevanten Flugzeugersatzteils entwickelt. Das Modell basiert auf den Verfahren der Blockchain-Technologie, insbesondere Hyperledger

[9] Eine ausführliche Beschreibung des generischen Transaktionsflusses findet sich in der Dokumentation von Hyperledger unter `http://hyperledger-fabric.readthedocs.io/en/latest/txflow.html`.

Fabric, die in Abschn. 4.2 diskutiert wurden. Dieser Beitrag zur Nutzung von Blockchain für den Anwendungsfall der Surplus-Teile-Supply-Chain, kommt dem Aufruf nach „Blockchain applications“ in Zheng et al. (2017) und der Verknüpfung von neuen Technologien mit einem Anwendungsfall im Sinne von Glaser (2017) nach.

Das entwickelte Artefakt ist der konzeptionelle Transaktionsfluss, der aufzeigt, auf welchem Weg Reparaturzertifikate gespeichert werden sollen, um alle sechs Anforderungen zu erfüllen. Wie in Abb. 4.2 zu sehen, konzentriert sich die Konsensfindung auf einen Orderer-Dienst. Nach Rücksprache mit Vertretern der MRO-Branche stellt dies ein realistisches Szenario dar. Bei der Erarbeitung des Modells stellte sich heraus, dass das Problem weniger im mangelnden Vertrauen in eine zentrale Instanz liegt, sondern vielmehr in der Sorge, dass das System einem Single Point of Failure erliegt. Unter diesem Gesichtspunkt ist eine Blockchain-basierte Lösung immer noch sinnvoll, diese wird jedoch nicht die Potenziale des trust-free und der Dezentralität ausnutzen, für die Hyperledger Fabric entwickelt wurde. Gefahren bei einer Einführung könnten in einer unzureichenden Akzeptanz und mangelnden Durchsetzbarkeit des digitalisierten Geschäftsprozesses bei allen Teilnehmern sowie bestehenden Zielkonflikten zwischen den Stakeholdern liegen – während MRO-Anbieter an einer langfristigen Verwendung von Teilen, einer Minimierung der Verschrottungsrate und einem sicheren und effizienten Sekundärmarkt interessiert sind, haben Erstausrüster von dem bisher ineffizienten Prozess durch eine höhere Nachfrage nach Neuteilen profitiert.

Das Modell ließe sich gegebenenfalls auf ähnliche Industrien adaptieren. Denkbar ist die Verwendung unter geringfügiger Anpassung in allen Industrien, die auf eine komplexe und schnelle Versorgung mit sicherheitsrelevanten Ersatzteilen ausgerichtet sind, wie die der Schifffahrt, dem Schienenverkehr oder der Energieerzeugung. Weitere Beiträge zur Wissensbasis würden in der Erforschung der Zusammenhänge zwischen den Herausforderungen einzelner Branchen und deren Adressierung durch Blockchain-basierte Verfahren liegen.

Der nächste Schritt besteht in der Evaluation des Transaktionsflussmodells auf mehreren Ebenen. In Hevner et al. (2004) wird insbesondere festgehalten, dass die Evaluation neuer IT-Artefakte auf zwei Ebenen erfolgt. Die quantitative Evaluation umfasst analytische Simulationen und numerische Vergleiche mit alternativen Designs. Eine zusätzliche qualitative Evaluation erfolgt über die Anwendung im unternehmerischen Kontext. In einem nächsten Schritt sollte ein Simulationsframework entwickelt werden, um das Transaktionsflussmodell zu testen. Gewonnene Erkenntnisse sollten in die Verfeinerung des hier vorgestellten Modells und weitere Anpassungen auf die jeweilige Anwendungsdomäne fließen. Die anstehende Evaluierung in der Praxis zu einem Proof of Concept soll in Kooperation mit einem MRO-Anbieter erfolgen, was der Aufforderung nach „Blockchain testing“ aus Zheng et al. (2017) nachkäme.

Literatur

Abeyratne, S. und Monfared, R. (2016). “Blockchain ready manufacturing supply chain using distributed ledger”. In: *International Journal of Research in Engineering and Technology* 5.09, S. 1–10. DOI: 10.15623/ijret.2016.0509001.

Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., Enyeart, D., Ferris, C., Laventman, G. und Manevich, Y. (2018). “Hyperledger fabric: a distributed operating system for permissioned blockchains”. In: *Proceedings of the Thirteenth EuroSys Conference*. ACM, S. 30. ISBN: 1450355846. DOI: 10.1145/3190508.3190538.

Blumberg, D. F. (2004). *Introduction to Management of Reverse Logistics and Closed Loop Supply Chain Processes*. CRC Press. ISBN: 9780203502938. DOI: 10.1201/b12427.

Buterin, V. (2014). Web Page. DOI: citeulike-article-id:14229692. URL: https://github.com/ethereum/wiki/wiki/White-Paper (besucht am 17. 04. 2018).

Cachin, C. (2016). “Architecture of the Hyperledger blockchain fabric”. In: *Workshop on Distributed Cryptocurrencies and Consensus Ledgers*.

Christidis, K. und Devetsikiotis, M. (2016). “Blockchains and Smart Contracts for the Internet of Things”. In: *IEEE Access* 4, S. 2292–2303. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2016.2566339.

Eaves, A. H. C. und Kingsman, B. G. (2004). “Forecasting for the ordering and stock-holding of spare parts”. In: *Journal of the Operational Research Society* 55.4, S. 431–437. DOI: 10.1057/palgrave.jors.2601697.

Ethereum (2017). *Consortium Chain Development*. Electronic Article. URL: https://github.com/ethereum/wiki/wiki/Consortium-Chain-Development (besucht am 29. 06. 2018).

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 2017*. ISBN: 978-0-9981331-0-2. URL: http://hdl.handle.net/10125/41339.

Group, H. A. W. (2018). *Hyperledger Architecture*. Electronic Article. URL: https://www.hyperledger.org/wp-content/uploads/2018/04/Hyperledger_Arch_WG_Paper_2_SmartContracts.pdf (besucht am 29. 06. 2018).

Hevner, A., March, S., Park, J. und Ram, S. (2004). “Design Science in Information Systems Research”. In: *MIS quarterly* 28.1, S. 75–105. DOI: 10.2307/25148625.

Nakamoto, S. (2008). “Bitcoin: A Peer-to-Peer Electronic Cash System”. In: URL: https://bitcoin.org/bitcoin.pdf (besucht am 14. 09. 2017).

Nofer, M., Gomber, P., Hinz, O. und Schiereck, D. (2017). “Blockchain”. In: *Business & Information Systems Engineering* 59.3, S. 183–187. ISSN: 2363-7005. DOI: 10.1007/s12599-017-0467-3.

Renesse, R. van und Guerraoui, R. (2010). “Replication techniques for availability”. In: *Replication: Theory and Practice*. Hrsg. von B. Charron-Bost, F. Pedone und A. Schiper. Replication: Theory and Practice. Springer-Verlag Berlin Heidelberg. Kap. 2, S. 19–40. ISBN: 978-3-642-11294-2. DOI: 10.1007/978-3-642-11294-2.

Schneider, F. B. (1990). “Implementing fault-tolerant services using the state machine approach: A tutorial”. In: *ACM Computing Surveys (CSUR)* 22.4, S. 299–319. ISSN: 0360-0300. DOI: 10.1145/98163.98167.

Szabo, N. (1997). *Formalizing and Securing Relationships on Public Networks*. Electronic Article. URL: https://ojphi.org/ojs/index.php/fm/article/view/548/469 (besucht am 29. 06. 2018).

Wickboldt, C. und Kliewer, N. (2018b). “Value Based Pricing meets Data Science: A Concept for Automated Spare Part Valuation”. In: *Multikonferenz Wirtschaftsinformatik (MKWI) 2018*. ISBN: 978-3-935786-72-0. URL: http://mkwi2018.leuphana.de/wp-content/uploads/MKWI2018_Band1.pdf.

Zheng, Z., Xie, S., Dai, H., Chen, X. und Wang, H. (2017). “An overview of blockchain technology: Architecture, consensus, and future trends”. In: *Big Data (BigData Congress), 2017 IEEE International Congress on*. IEEE, S. 557–564. ISBN: 1538619962. DOI: 10.1109/BigDataCongress.2017.85.

Chapter 5

Blockchain for Workshop Event Certificates - A Proof of Concept in the Aviation Industry

Abstract

For the maintenance of a fleet, e.g. for air, sea or rail traffic, Maintenance Repair and Overhaul (MRO) companies depend on a safe and fast supply of spare parts. Particular importance lies in safety-relevant spare parts, which must have a complete life cycle documentation of all workshop events in order to be approved for further use. Current documentation practice is characterized by manual recording and analogue storage of workshop certificates, which leads to a high degree of uncertainty due to manipulable, error-prone and incomplete documentation and thus unnecessary scrapping of valuable spare parts. Because of the lack of a central authority that could provide a corresponding IT system, digitization of the documentation process was not yet possible.

This work follows design science principles to create, evaluate and present a blockchain-based IT artefact that enables the digitization of the documentation process based on Hyperledger Fabric. Our research responds to shortcomings in the current analogue business process and to calls for practical use cases in the blockchain research field. Relevance is ensured by close cooperation with domain experts from the aviation industry. Projectable patterns in the documentation process are identified to enable an application of the artefact to a broad solution space.

Keywords: Blockchain, Hyperledger Fabric, Workshop Events, Aviation Industry.

Contents

5.1 Introduction

The Blockchain technology is referred to as a "trust machine" (Economist, 2015). The technologies which Blockchain is based on are rather old. The idea of Smart Contracts has already been introduced more than 20 years ago (Szabo, 1997). Nevertheless, since the introduction of a combination of these technologies by Nakamoto (2008), research of what kind of problems can already be solved with Blockchain has been increased. Still, Glaser (2017) states that "Blockchain is an innovative technology in search of use cases".
This applies to many industries and also to the aviation industry. Within the industry, there are multiple ambitions to experiment with blockchain technology. Air France KLM is partnering with universities, manufacturers and software developers to establish use cases to improve maintenance processes and workflows, e.g. getting a replacement part to Air France (Bellamy, 2017). Lufthansa Industry Solutions is discussing the application in aircraft maintenance in general within the Maintenance Repair and Overhaul (MRO) industry (Evers, 2018). Also consulting companies like Accenture are discussing blockchain applications not only in maintenance but also in ticketing, loyalty programs and security (Accenture, 2018).
The International Air Transport Association (IATA) states that component maintenance costs make up the second largest portion (24%) of airlines' direct maintenance costs (IATA, 2015). IATA as a certification authority discusses now digital platforms which enable industry partners in the value chain to engage in business in a safe and secure manner (IATA, 2018b). One critical part of this value chain is the maintenance repair and overhaul of aircraft spare parts. The whole industry is characterized by isolated systems, for example in trading of surplus spare parts (Wickboldt and Kliewer (n.d.), Wickboldt and Kliewer (2018b)) but also in the maintenance process itself. Safety-related parts require a particular documentation of the parts history and maintenance actions related to it. Life Limited Parts (LLPs) cannot be utilized without a traceable history. Incomplete documentation histories are causing high costs as the value of these parts is usually in the six digits. This documentation is performed manually and saved on paper documents. IATA (2015) states that "Synchronizing paper work (job cards with "dirty finger prints") and computerized record system can introduce delays in the cycle. Paperless operations will improve this process and reduce these potential delays." Wickboldt and Kliewer (2018a) already argued that there is no central institution, which could handle the data if this documentation process is digitized. In order to address the need for such a system we propose a proof of concept of Wickboldt and Kliewer (2018a).
Within this scope, the contribution of our research is threefold: First, it addresses the calls for practical use cases in the blockchain research field by providing insights at the interface between technology, markets and business models (Nofer et al., 2017) and follows the call for "blockchain applications" in Zheng et al. (2017) and

the linking of new technologies with an application case in the sense of Glaser (2017). Second, we provide a proof of concept for a conceptual model which is discussed in Wickboldt and Kliewer (2018a). Third, we identify generalizable patterns in the documentation process so that the design developed here can be used to address similar problems with little effort.
In order to arrive at these contributions the remainder of this paper is structured according to the design science research approach from Gregor and Hevner (2013) as follows: In Section 5.2, we provide a systemic review on the status quo of the problem space and existing solutions. We also give insight into the special features of the business process and by that identify research gaps in the existing knowledge base to handle these features. Section 5.3 introduces the method selection. In Section 5.4 the design search process as well as the resulting artefact is described. The proposed design is constantly questioned, tested against requirements and constraints of the business process (Hevner et al., 2004) to satisfy the laws in the existing environment (Simon, 1996). Results of the evaluation of the artefact are presented in Section 5.5. Section 5.6 discusses the research findings. We conclude and give an outlook in Section 5.7.

5.2 Distributed Ledger Technology

This section provides a brief introduction to blockchain-based systems, outlines the workshop event documentation use case and identifies the research gap. Therefore, Section 5.2.1 reviews the state of the art of Distributed Ledger Technology (DLT) and discusses advantages and disadvantages of the systems. Section 5.2.2 illustrates the use case for documentation of workshop events in the aviation industry and its practical challenges. We close this section by deriving research questions to improve the current documentation process with a decentralized system and the requirements and metrics to achieve that improvement.

5.2.1 State of the Art Technology

DLT is an umbrella term that includes the commonly used term Blockchain. Even before the idea of Blockchain (Nakamoto, 2008), Szabo (1997) introduced the concept of Smart Contracts for the formalized and secure exchange of contracts via the internet which can be triggered by certain events. Due to the fact that these contracts are autonomous and do not rely on each other, Smart Contracts become an alternative for established intermediaries such as banks or marketplace operators (Glaser, 2017). This is particularly relevant for environments where there is no central intermediary. Buterin (2014) goes beyond the financial industry and proposes a blockchain model specializing in the exchange of Smart Contracts. It acts as a trusted distributed application and has the task of providing secure (business)

relationships within a distributed system. This is very similar to the well-known State Machine Approach, a general method for implementing fault-tolerant services in distributed systems (Schneider, 1990). A State Machine on the one hand consists of state variables which determine the systems state and on the other hand commands that change the state of the system. A component (e.g. a Smart Contract) is considered to be defective if its behaviour no longer matches its specification. The distribution over several systems should lead to the fact that in case of a fault only the respective part of the system is affected, and, in this way, the fault tolerance of the overall system is increased. This is vitally important for environments which cannot tolerate a single point of failure. Among other things, the method is based on consensus between distributed systems, for example under the tolerance of Byzantine Faults (Dolev and Strong, 1983).
Blockchains differ from traditional distributed State Machines as follows: Not only one application runs on the distributed system, but many do simultaneously. Applications can be dynamically created and modified by any stakeholder. The code of such an application is not trustworthy and therefore potentially harmful due to the lack of a central intermediary. By achieving consensus by many parties and via an algorithm, the responsibility but also power of finding consensus and by that control of the information is decentralized. This is vitally important for environments which lack of trust between actors. Blockchains have the following characteristics (Zheng et al., 2017):

- Decentralisation: In conventional, centralized systems, transactions are processed via an intermediary who validates the transaction. In the blockchain, these confidence-building intermediaries are replaced by a consensus mechanism so that a blockchain is trust-free. The data is distributed among the nodes of the network.
- Persistence: Once the consensus has been found, it is not possible to change the data status retroactively. Blocks with transactions that contradict the consensus are not accepted.
- Anonymity: In principle, it is possible to carry out transactions exclusively via the disclosure of an address generated at the beginning. Further personal data is not necessary due to the trustfree property.
- Verifiability: Every transaction is digitally signed. Due to data persistency, all transactions remain traceable.

Blockchain systems can be divided into the domains public, consortium and private. All records in the public blockchain are visible to the public and anyone could participate in the consensus process (Zheng et al., 2017). One example of this is

Bitcoin, the first and most prominent application of Blockchain in the financial industry (Nakamoto, 2008). Conversely, only a group of preselected nodes would participate in the consensus procedure of a consortium blockchain. With private blockchain, only the nodes of a particular organization are allowed to join the consensus process. A private blockchain operates in a proprietary network because it is completely controlled by nodes of an organization. The consortium blockchain, in which consensus is agreed upon by nodes from multiple selected organizations, is partially decentralized. Because records are stored on a large number of nodes, it is virtually impossible to manipulate transactions in a public blockchain. Conversely, transactions in a private blockchain or a consortium blockchain could easily be manipulated, as there is only a limited number of participants. In case of a permissioned blockchain, all participants are known, therefore manipulation could be sanctioned. The efficiency correlates negatively to the number of consensus finding nodes (miners) in the system. Many miners provide for a long process of finding a consensus while few miners come faster to a consensus.

Consortium blockchains are designed for use in a restricted circle of stakeholders. Ethereum has provided tools for the development of consortium blockchains (Ethereum, 2017) but lacks the possibility to privately exchange information for heterogeneous business processes in channels between a subset of participants. There is also a fork of Ethereum called Quorum[1] for a permissioned blockchain, developed by J.P. Morgan and focused on the financial industry. Finally, Hyperledger Fabric is a modular and extensible open source system for providing and operating permissioned blockchains for business consortia which provides more flexibility than Ethereum. It is one of the Hyperledger projects operated by the Linux Foundation (Androulaki et al., 2018). One of its key features is the Chaincode. A Chaincode typically implements the business logic on which the stakeholders of the network have agreed. In this sense, it is a Smart Contract. The other one is the Hyperledger Fabric Transaction Flow which is divided into three steps and can be executed on different peers in the network. A Committer starts transactions, manages the ledger and the status. An Endorser receives transaction proposals, simulates them on the basis of the Chaincode or rejects them. An Orderer includes released transactions in the blockchain and orders them. Non-permissioned blockchains like Ethereum, where participants are not necessarily known, consensus is found via algorithms like Proof of Stake or Proof of Work. As a permissioned blockchain, Hyperledger Fabric finds consensus via endorsing transactions. Due to the recognition of the participants, abuses can be punished. Greenspan (2015) proposes a framework called Multichain which is similar to Hyperledger Fabric but differs in predefined transaction types instead of Chaincode.

In the following, a use case is used to investigate the extent to which the prop-

[1] `https://www.jpmorgan.com/global/Quorum`

erties of DLT are suitable for addressing the causes that previously prevented the digitalization of the business process.

5.2.2 Use Case: Documentation of Workshop Events in the Aviation Industry

The aviation industry consists of a restricted circle of stakeholders, specifically business processes around maintaining aircrafts. Stakeholders include the original equipment manufacturers, mechanics, warehouse employees, workshop employees, airline employees, traders and government authorities. Every stakeholder has different tasks and rights within the industry. Each party must follow certain processes and release certain data to document workshop events. Given its importance for aircraft safety, enormous cost effects, large size of stakeholders involved and the need for highly reliable documentation, we inquire a key use case in aircraft maintenance. The documentation of workshop events in the aviation industry is selected as a use case for this proof of concept because it is highly critical for the continuous use of core resources in the aviation industry and is discussed in further detail below.

There are two reasons why a component (combination of parts) or part needs to be sent to a workshop, which are either condition-related or lifetime-related. First, a part has to be replaced condition-related if it fails or certain parameters are exceeded. This is done according to a Reliability-Centered Maintenance (RCM) approach. RCM is comprehensively described by Nowlan and Heap (1978) as an approach which directs maintenance efforts at those parts where reliability is critical. This replaces the old belief that every part of a complex system has a fixed life limit at which an overhaul is needed which was the basis for simply scheduled based maintenance programs (Ben-Daya et al., 2009). Studies in the airline industry showed that scheduled overhauls didn't have large impact on overall reliability, in fact a large portion of items had no “wearout zone”, hence their performance is not correlated to the age (Nowlan and Heap, 1978). Second, time limits however can be discerned in soft life limits which are defined by the operator and hard time limits which are mostly determined by the manufacturer. The operator determines a schedule for maintaining LLPs which is fitted to the aircraft and documented in an Air Operators Certificate (IATA, 2015). The goal of these maintenance programs is to ensure a continued airworthiness.

A component is identified by part number, part description, part serial number, and life limited schedule intervals, in particular hours since new (HSN) and cycles since new (CSN). In order to maximize efficiency, the operator intends to maximize the utilization between these intervals. That includes avoiding a pull and replace of LLPs before their schedule (e.g. unscheduled condition-based repairs). IATA advises operators to introduce soft time limits based on experience that triggers a removal before a failure is expected (IATA, 2015). This is done in practice using a

Weibull statistical analysis. The goal of such analysis is to find the time when it is most efficient to perform the overhaul (Love and Guo, 1996).
A component is unserviceable if it is in a condition that does not allow operative use. A repair leads from an unserviceable to a serviceable condition. An overhaul exceeds these aspirations by restoring the component to "zero time" according to the relevant manual. After a complete overhaul, the component is tested to the same requirements as a new item (IATA, 2009). The focus within the component repair cycle lies on a reasonable short turnaround time in other words the elapsed time between the removal from the aircraft and the return to a serviceable condition.
Figure 5.1 describes the ideal component repair cycle according to IATA. If an aircraft part fails or is due to repair or overhaul, it is removed by a mechanic and is marked as unserviceable. This unserviceable part is stored in a warehouse until a repair shop is ready to repair it. The repair shop performs a repair or an overhaul to set the parts state back to serviceable. This part is then stored in a warehouse until a mechanic installs it in an aircraft. Discussions with domain experts show that usually one part goes through the repair cycle while the aircraft is provided with an equivalent part for a quick pull and replace.

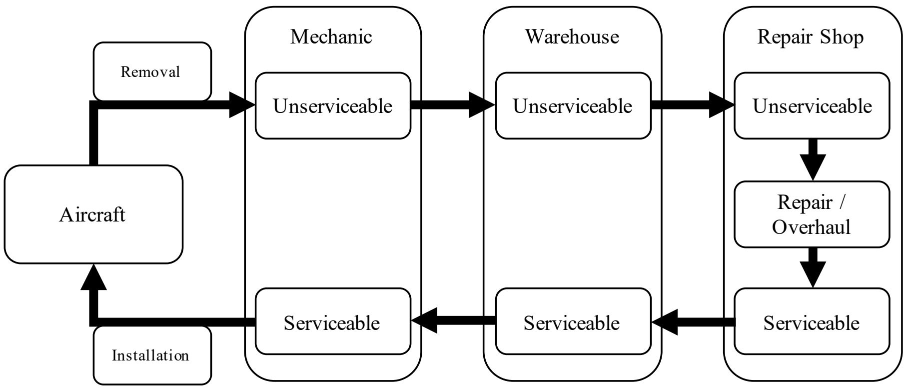

Figure 5.1: Ideal Component Repair Cycle based on IATA (2015)

Especially for LLPs a complete documentation of workshop history, metrics and ownership is needed. Currently, these events are documented via certificates which are printed out on paper and put next to the part. These certificates can be divided into multiple categories, such as *Life Limited Parts Profile* which documents metrics such as total operating time, total operating cycles, the parts cycle limit and time limit of all parts which are combined to a component. There are also *Authorized Release Certificates* which ensure that a part is conform to its design and in a condition for safe operation. This process is still performed analogue due to the fact that there is no central authority which could host a digital system to collect and store the data. The lack of trust in the industry leads to a limited exchange of information. The documentation history is not transferred to the next owner until there is a change of ownership, provided it is still complete.

5.2.3 Existing Artefacts and Research Gap

There are already systems which are targeted to overcome lacking trust by decentralizing the information exchange within a supply chain. Realized case studies for the tracking of goods via a blockchain are Everledger and Provenance (Provenance, 2015). Everledger[2] is a project for tracking diamonds along the value chain based on Hyperledger Fabric. This is intended to ensure that no so-called blood diamonds are traded in a comprehensible manner. Provenance is a company that develops applications for tracking goods, especially in the food industry, based on Ethereum blockchains. Notheisen, Cholewa, et al. (2017) propose a blockchain-based system to track ownership of cars and document the vehicle history for market participants, authorities and third parties. They make use of Ethereum blockchain technology, which is, as a public blockchain, suitable for that business process.
The areas of application for these systems distinguish from our case mainly due to their public nature. However, the documentation of workshop events in the aviation industry requires depicting a central authority that does not exist in this way. Because of the heterogeneity of the documentation certificate types, a system which is able to apply non-deterministic business rules based on the current situation is needed. The identified research gaps lead to the following research questions (RQ):

- (RQ1) How can the documentation process be digitized in a decentralized data model?
- (RQ2) What kind of system realizes a decentralization of the information exchange?

To address these research questions and make a contribution which has practical implications, this system is a proof of concept of already discussed requirements (R) (Wickboldt and Kliewer, 2018a):

- (R1) High speed so as not to interfere with the process of using parts.
- (R2) The possibility to merge events from different sources in an orderly way.
- (R3) A persistent storage of information.
- (R4) A counterfeit-proof solution, which may work without trust if necessary.
- (R5) Limited visibility, i.e. a solution with access authorizations.
- (R6) A solution that is capable of mapping rules of the business process.

[2] `https://www.everledger.io`

An evaluation of the designed artefact requires measurement (March and G. F. Smith, 1995). In order to ensure a contribution which has measurable success, we talked to domain experts from the aircraft MRO industry to compile relevant metrics. Additional to the requirements R1-R6 the following metrics (M) have to be met:

(M1) Sales transactions must be executed immediately, others should be executed within 10 minutes. Confirmation of cycles per hour is only necessary at certain times; the speed of those queries is of secondary importance.

(M2) The system needs to be capable of handling a throughput of 5 Mio. transactions per year[3] or ∼10 transactions per minute or ∼0.16 transactions per second (tps).

5.3 Methodology: Design Science Approach

To structure the creation, evaluation and presentation of this proof of concept, we rely on the Design Science Research (DSR) approach proposed by Hevner et al. (2004). As seen in similar research papers (e.g. Notheisen, Cholewa, et al. (2017)), we summarize the mapping of our research against DSR guidelines of Hevner et al. (2004) in tabular form (see Table 5.1). The resulting IT artefact is a proof of concept prototype which aims to replace a paper-based documentation process with a decentralized, persistent and permissioned blockchain. In order to ensure efficiency and efficacy of the system, we develop metrics and then constantly re-evaluate the artefact against these metrics (March and G. F. Smith, 1995).

[3] This is a projection elaborated with domain experts based on their own number of transactions and their market share.

Guideline	Contribution
Design as an artefact	The outcome of our research is a proof of concept prototype that implements a blockchain-based IT artefact which executes the documentation process for workshop events in the aviation industry.
Problem relevance	Our research questions respond to the need to digitize a process that, due to the absence of a central institution, has not yet been digitizable (RQ 1). Furthermore, we react to the necessity of mapping heterogeneous documentation requirements, which today can be seen in the form of different certificate types (RQ 2). Additionally, we ensure relevance by continuously working with domain experts of the aircraft MRO industry.
Design evaluation	According to Hevner et al. (2004) we perform an analytical evaluation by examining the structure (system architecture) and dynamic (performance measures) of the artefact. Because of the highly innovative characteristics of this artefact, we are performing a descriptive evaluation by constantly exchanging information with domain experts in the MRO industry. The evaluation of the IT artefact is done in an artificial environment following a Technical Risk & Efficacy evaluation strategy based on Venable et al. (2016).
Research contributions	The contribution of our research is threefold: First, our research responds to the calls for practical use cases in the blockchain research field. We are providing insights at the interface between technology, markets and business models (Nofer et al., 2017) and make a practical contribution to the still small knowledge base regarding the blockchain technology and its solutions space. This research follows the call for "blockchain applications" in Zheng et al. (2017) and the linking of new technologies with an application case in the sense of Glaser (2017). Second, we provide a proof of concept for a conceptual model which is discussed in Wickboldt and Kliewer (2018a). Third, we identify projectible patterns (Baskerville and Pries-Heje, 2014) in the documentation process that offers a broad solutions space. Hence, it may be used to address comparable problems in other sectors.
Research rigor	To ensure our research's rigor, we fall back on well-established frameworks in Hevner et al. (2004), Gregor and Hevner (2013) and Venable et al. (2016) which enable us to formulate the creation and evaluation of our artefact in a generalizable way and in the sense of commonly known DSR. Additionally, we embed our research into the already developed guidelines for architectural decisions of Glaser (2017) and Xu et al. (2017)
Design as a search process	In order to arrive at the research contributions, we build on existing literature about blockchain-based transaction systems, such as Nakamoto (2008) and Androulaki et al. (2018) and continuously evaluate and search for satisfactory solutions (Simon, 1996) throughout the development process.
Communication of research	To clarify the insights of our research and maximize its potential impact in technology-oriented as well as management-oriented audiences, we structure our work according to Gregor and Hevner (2013) and utilize the case of the documentation of workshop events in the aviation industry to illustrate the business environment around the developed artefact. In order to address technology-oriented audiences, we present a detailed description of the software architecture in Unified Modeling Language (UML). To address management-oriented audiences, we furthermore describe the underlying business process and embed industry-specific business processes into the process management literature. Finally, we prove the effectiveness of our artefact by discussing potentials and limitations of our solution with domain experts of the aircraft MRO industry.

Table 5.1: Mapping of our IT artefact against the DSR guidelines of Hevner et al. (2004)

5.4 A Blockchain-based Certification Storage System

This section deals with the presentation and justification of the design decisions for the IT artefact. The artefact itself is then described based on its software architecture.

5.4.1 Blockchain design decisions

According to Gregor and Hevner (2013), the process of design search (development) which led to the discovery of the artefact design is described. Based on Xu et al. (2016) and Xu et al. (2017), we present the following design decisions.

Blockchain design decisions

Decision 1 (Transaction processing rate): As described in Section 5.2.3, we expect 0.16 tps. In a single channel environment, Hyperledger Fabric is capable of 140 tps and up to 2,250 tps if performance optimizations are applied (Thakkar et al., 2018) which meets R1 and M2 (transaction speed). M1 (transaction prioritization) is also addressed with this decision as there is no prioritization needed with the current performance expectations.
Decision 2 (Consensus protocol): We choose the Hyperledger Fabric Transaction Flow to meet R6 (business process). In Hyperledger Fabric trust can be created in the endorsement and not with a consensus protocol.

Application design decisions

Decision 1 – Data Structure: Master data is stored off-chain, all data of the certificates is stored on-chain to guarantee persistency through decentralization (R3). This ensures small data size on the blockchain which keeps the systems responsiveness at a high level.
Decision 2 – Blockchain Scope: A consortium blockchain is needed to meet R2 (confluence) and R4 (trust). However, in this early proof of concept, which is evaluated in an artificial environment, we choose only one Orderer node, which technically makes this system a private blockchain.
Decision 3 – Single/multiple chains: For easier chain and permission management we choose a single chain environment. Hyperledger manages the communication via channels, thus creating a blockchain of subordinate areas (n channels) to meet R6 (business process). In this phase we store one transaction per block to ensure no time-outs due to the low expected transaction frequency of 0.16 tps. However, bundling of multiple transactions has been tested and is possible to realize scalability.
Decision 4 – Protocol Configuration: We choose Checkpointing (Hyperledger, 2018) in Hyperledger Fabric, meaning all network participants have to accept the trans-

actions up to the checkpoint as valid and irreversible. As all network participants are known and the Chaincode applies for everyone alike, trust is generated through the traceability of every transaction. In Hyperledger (2018) a consensus is defined as "the full-circle verification of the correctness of a set of transactions comprising a block". As in this state of the proof of concept as simple private blockchain only one Orderer is needed. This meets R6 (business process).
Decision 5 – Authorization: In order to meet R5 (permission), we chose a permissioned blockchain. Hyperledger Fabric has been proven to be a suitable solution for supply chain use cases. Quorum, as a fork of Ethereum is permissioned but as a tokenized system focusses on the financial industry. It also has only simple permission levels while Hyperledger Fabric has fine-grained permissions. We chose Hyperledger Fabric because of the possibility to design the system architecture in a modular way. A detailed justification for the choice of Hyperledger Fabric is provided in Wickboldt and Kliewer (2018a).

5.4.2 Software Architecture

The proposed artefact is a Blockchain-based Certification Storage System (BCS System). Figure 5.2 shows a high-level dataflow diagram during the process of storing a certificate. This dataflow is active every time a certificate is created, changed or queried which is ideally any time a workshop event occurs. The user has to be logged in at the frontend and must have the required authorizations.

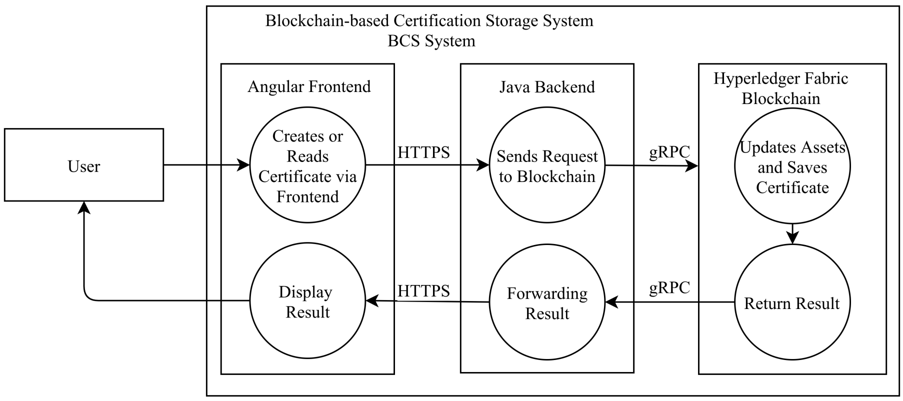

Figure 5.2: Dataflow diagram for certificates in Blockchain-based Certification Storage System

When the certificate is created in the frontend, a form is transferred via encrypted https to the Java Backend. At this point, all certificate information is available on the backend where the form data is then processed by a controller, so that it becomes clear which function is called now, and which data is given with this

call. The requests are then sent to the blockchain via gRPC which is a protocol for calling distributed systems. Within the Hyperledger Fabric Blockchain, the certificate enters the Hyperledger Fabric Transaction Flow. The transaction flow was adapted to the use case in Wickboldt and Kliewer (2018a). This flow is divided in four steps. First, the data from the form is forwarded to the endorsing peers and simulated against the Chaincode. All endorsed transactions are then ordered by the Orderer node and broadcasted to the peers. The ordered transactions are then validated, conflicting transactions are eliminated. The result is returned to the peers, the persistent state is then saved. The information is then retrieved from the Hyperledger Fabric Blockchain via gRPC and is forwarded by the Java Backend to the Angular Frontend via https. The user finally sees the result in the frontend. Figure 5.3 shows the data model, generic data in grey, status reports in green, sale and owner documents in blue and information of an aircraft part in orange.
The Generic Certificate carries generic information for all certificate types and is referenced via the *serial number*[4] of the part. It can be enhanced by detailed information regarding the certificate type. The most common documents and certificates are explained in more detail: A Bill of Sale documents the price for a change of ownership. A Generic Report documents the technical inspection after a sale. The Aircraft Part certificate stores technical and commercial information and is called if an overview of the part is needed. It has the following attributes: A *serialnumber* is the primary key for a specific part. *partOf* describes the relation for a part to a larger component. A component could consist of (*consistsOf*) multiple parts. A part is characterized by the parameters *cyclessincenew* and a *cyclelimit* as well as *timesincenew* (HSN) and a *timelimit* in hours. Equal parts have the same *partnumber*. The *partdescription* helps the user to get a quick information about the parts characteristics. A part always has a current *partOwner* and a *partStatus* (e.g. overhauled, serviceable). Record of Inspection documents the inspection of the part and its measurements and dimensions. A Life Limited Parts Profile is used to update CSN and HSN. A Life Limited Part Usage Report is done on component level. A Certificate for Repairs stores information about workshop events. A Storage Certificate captures the current storage state. Information of a check before a resale is stored in a Material Certificate. A Life Limited Parts List lists information of all parts which are part of a component.

[4] We have deliberately chosen *String* as data type in order to take account of the existing nomenclature in the industry.

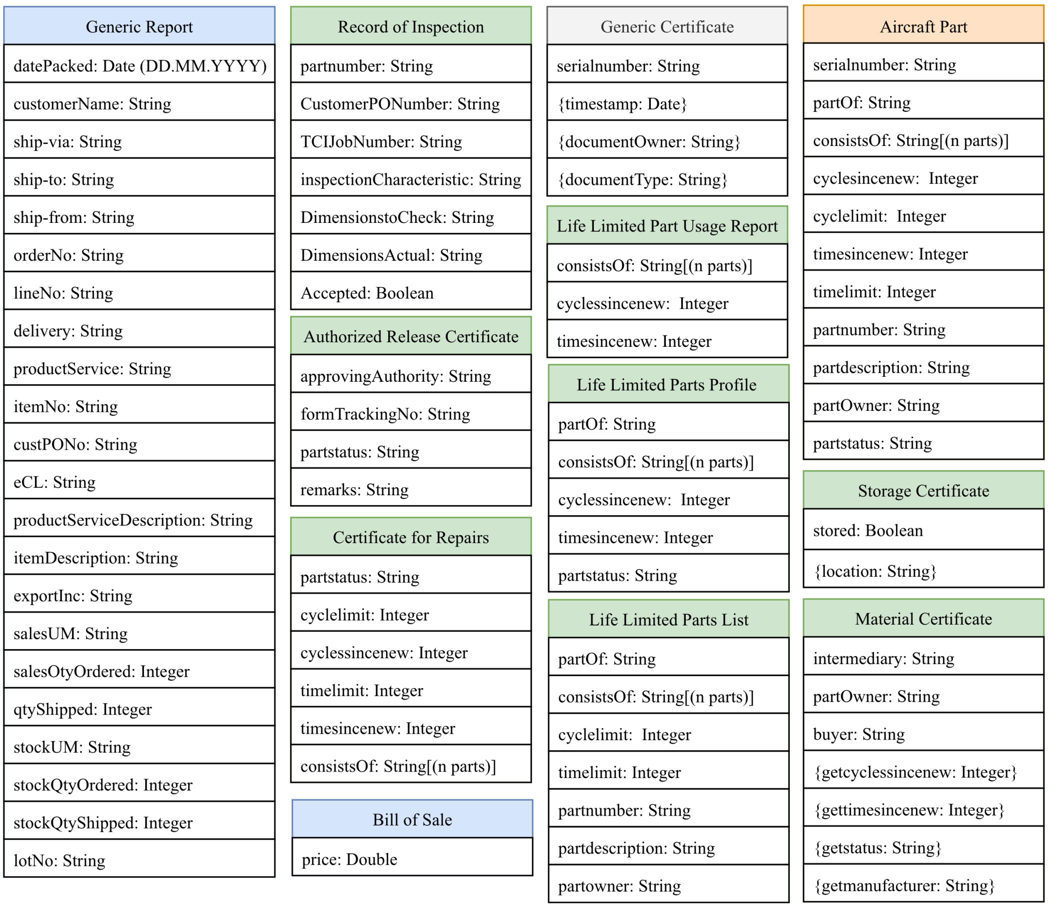

Figure 5.3: Data model of the Blockchain-based Certification Storage System

The following methods (starting with a capital letter) implemented in Chaincode are able to be triggered via the frontend and expect parameters (starting with lower case).
CreateAircraftPart takes all arguments of *Aircraft Part* to create a new part certificate. If the aircraft part with the index serialnumber does not already exist, a new aircraft part certificate is created and stored with its *serialnumber* in the blockchain.
ReadAircraftPart takes a serialnumber. If a correct *serialnumber* has been passed, and the aircraft part exists in the blockchain, it is returned in the form of a JSON string or otherwise throws an error.
CreateCertificate takes *cert_serialnumber*, *documentType*, *documentOwner* and *details* to create a certificate for the corresponding aircraft part, if it exists.
cert_serialnumber is a concatenation of cert ID and *serialnumber* as primary key for a certificate for one specific part. *details* contains all attributes from the data model

for the respective certificate.
GetCertificateHistoryForAircraftPart takes a *serialnumber*. The history of all certificates belonging to the referenced aircraft part is returned.
Figure 5.4 shows a sequence diagram for creating a valid certificate in more detail. The user enters the data in the frontend and submits the form to trigger one of the four mentioned methods (e.g. *CreateCertificate*). The request is sent as a JSON string to the backend via a REST API. The REST API constructs the proposal for transaction. The backend executes the request and transmits the information to a peer in the network. The backend then contacts endorsing peers and attempts to sign the transaction so that the endorsement policy is met. If it is fulfilled, a request is sent to the Orderer which forwards the transaction to all peers in the network (channel) and saves the change persistently in the respective state database of every endorsing peer (in this case CouchDB). After the update a response is written by the blockchain and sent back to the backend. This information is then forwarded to the user.

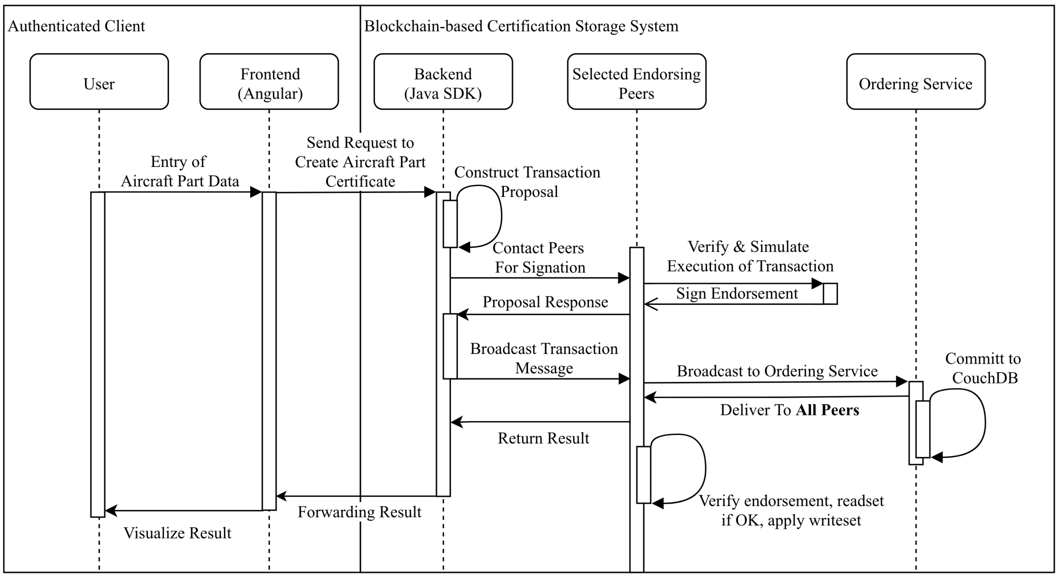

Figure 5.4: Sequence diagram for a valid transaction on the example of creating a certificate

To summarize, we presented an IT artefact which follows the design decisions in Section 5.4.1. The proof of concept prototype introduced in this section enables a persistent and decentralized documentation process. The system is running on Hyperledger Fabric and allows users to initially store workshop certificates or append them to an existing history. In total, we provide a solution for both research questions posed in Section 5.2.3. The use case in the aviation industry highlights the eligibility of the BCS System.

5.5 Evaluation

In order to ensure efficacy and efficiency of the artefact, we provide an evaluation against the requirements and measures presented in Section 5.2.3. The following evaluation assures rigor of our research. As March and G. F. Smith (1995) describe it, each artifact which has been "built" must be evaluated scientifically. Gregor and Hevner (2013) propose "final summative tests in case studies or experiments, expert reviews, simulations, statistics on usage data for implemented systems" and by that proof validity ("works and does what it is meant to do") and utility ("assesses whether the achievement of goals has value outside the development environment"). To achieve that, this evaluation is based on the Framework for Evaluation in Design Science (FEDS) which is a novel framework suited to use in DSR. Venable et al. (2016) state that "evaluation should be relevant, rigorous and scientific". In order to arrive at these targets, this chapter addresses the functional purpose of the evaluation (formative or summative) and the paradigm of the evaluation (artificial or naturalistic) described in Venable et al. (2016).
The major design risk lies in technical difficulties with the blockchain technology which is still in an early stage. As the quality of the artefact itself stands now for evaluation, an application in a real setting would be too expensive and risky. Because of these circumstances, we chose a "Technical Risk & Efficacy" evaluation strategy (Venable et al., 2016) which emphasizes artificial formative evaluations done in an iterative fashion. Because the BCS System is an early stage proof of concept prototype, we provide a formative evaluation in order to provide a foundation for continuous improvement and adaption (Wiliam and Black, 1996) of the artefact. Wickboldt and Kliewer (2018a) provided an initial design which has been the basis for implementation of the BCS System. This article contributes by providing an intermediate evaluation of the current implementation of the BCS System. Before releasing the artefact into field testing (Hevner and Chatterjee, 2010), we provide an empirical artificial evaluation in a controlled laboratory environment. To assure rigor in the assessment of the efficacy of the artefact, we abstract from the natural setting by restricting to selected transaction types which are discussed with domain experts.
The possibility to merge events from different sources in an orderly way (R2) is guaranteed through the use of an Orderer node. Information is stored persistently in a decentralized way by every endorsing peer (R3). The implementation of methods via Chaincode ensures a counterfeit-proof solution (R4). Limited visibility (R5) is ensured by restricting the frontend with an authorization mechanism and encrypt-

ing the data on the blockchain. The combination of Chaincode and implementing existing certificate types in our data model, we ensure to provide a solution that is capable of mapping rules of the business process (R6). The issue of speed and responsivity is of great importance for the acceptance of an IT system. That is why we address the requirement of high speed (R1) in more detail.
The measurements of the BCS System are benchmarked against the practical requirements defined in Section 5.2.3. To measure the performance of the system, we sent 5.000 SOAP requests for *createAircraftPart*. The BCS System is running on an Ubuntu Server 18.04 LTS with an Intel(R) Xeon(R) Gold 6140 CPU @ 2.30GHz (2 Cores) and 8 GB RAM. Storage is realized via a Serial Attached SCSI (SAS). The BCS System currently handles 7.99 tps. To compare this number, Ethereum currently performs at 6.5 tps[5]. One transaction has a size of ∼6 kb which leads to a size of ∼30 megabytes after 5,000 transactions or theoretically around 26 GB after 5 Mio. transactions which are projected annually. Over a course of 5000 transactions, the number of tps increases from 5 to 8 within the first 30 transactions. This might be due to the usage of SAS instead of a SSD which needs a moment to contact and spin up the hard disks. After that the tps rate is stable which indicates scalability of the system. The BCS System currently meets M2 which requires 0.16 tps. By exceeding the requirements for speed by a factor of 50, it is currently not necessary to prioritize the individual transactions. All transactions can be processed in the same way to still meet M1. Nevertheless, there is still room for performance improvement. To classify the numbers we refer to Thakkar et al. (2018) who did performance benchmarking on Hyperledger Fabric systems and provided insights about optimization leverage. They are reaching tps from 140 to 2.250. Right now, the BCS System takes every transaction into a single block. This of course takes time to transfer through the sequence shown in figure 5.4. In comparison to similar systems such as Notheisen, Cholewa, et al. (2017), no miners have to be rewarded therefore no tokens are being used to validate a transaction.

5.6 Discussion

This proof of concept shows that many of the shortcomings in documenting the aircraft component repair cycle can be overcome by taking advantage of blockchain technology. The identified research gap is bridged by digitizing the business process and storing the documents in a blockchain. The consistency and validity of the transactions are ensured by the endorsement process in the BCS System. By using an Orderer, the certificates for the documentation of workshop events as well as transfers of ownership are persistently recorded. The execution of the transaction against the Chaincode matching the certificate ensures that all stakeholders follow

[5] `https://www.etherchain.org`, continuously changing tps, accessed on 11/23/2018

the same rules. Malicious behavior and inadequate usage are ruled out by design of the algorithm (in the sense of Beck et al. (2016)) which is in this case implemented in the Chaincode. In this way, the BCS System ensures credibility without sharing more information than is strictly necessary for the transaction. Due to the decentralization of the system, no stakeholder is given the power to manipulate transactions. In addition, the uniform and tamper-proof filing of information provides regulatory authorities with a transparent basis for checking compliance with regulations. The use of the Hyperledger Fabric platform is constantly growing due to additional use cases. The high degree of dissemination ensures that, for example, performance benchmarks that go far beyond the horizon of our use case are valid. In this way, we can ensure that our design decisions do not limit the scalability of the system.

Limitations on this approach lie in the current state of digitization of aviation data. James Kornberg, director of innovation of the Air France KLM business unit, summarizes these limitations: "In the aviation industry we still have a lot of our data that is not digitized (..) blockchain cannot be applied on nonelectronic data." (Bellamy, 2017). The development of standards that are necessary for the productive use of the system requires the effort of all involved. Although the BCS System itself does not allow for technical tampering, the responsibility lies on the clear formulation of the Chaincode that prevents potential fraud attempts. Glaser (2017) legitimately states that the institution that sets the Chaincode has a major central influence on the execution of the system. Furthermore, the information in the BCS System can only carry as much truth as the data that the respective stakeholder enters into the form. One way to reduce the human risk factor is to shift information input from humans to the device itself through the combination of Blockchain and Internet of Things (Christidis and Devetsikiotis (2016); Zhang and Wen (2017)). For example, it is conceivable that spare parts could be equipped with NFC technology, which enables the automated and input-secure transmission of metrics such as CSN and HSN through the use of Internet of Things. The digitization of the documentation process could be the foundation for sophisticated predictive maintenance beyond the boundaries of the individual companies. Rogue units, which are components that are frequently removed within a short time, could be analyzed regarding the removal causes. Workshops with domain experts show that the digitization of the documentation process alone, which can be facilitated by DTL, is perceived as great utility from the stakeholders.

5.7 Conclusion

The outcome of this work is a blockchain-based IT artefact applied to a use case in the aviation industry and responds to the call for "blockchain applications" in Zheng et al. (2017) and the linking of new technologies with a use case in the sense

of Glaser (2017). The artefact developed is a proof of concept for the decentralized storage of certificates which is discussed in Wickboldt and Kliewer (2018a) that proves the possibility of digitizing the documentation process using the example of the aviation industry. We generalize patterns (Baskerville and Pries-Heje, 2014) in the documentation process which enables researchers to use our artefact and adapt it to comparable problems in other industries.
RQ1 (digitizing documentation process) was addressed in Section 5.4 by presenting a data model that takes the attributes of today's analogue certificates into account. Decentralization of information exchange without providing more information than needed under rules that are applicable to everyone has been realized through the BCS System which is based on the Hyperledger Fabric Blockchain framework and therefore addresses RQ2 (decentralized system). The current analogue process for the documentation of workshop events described in Section 5.2 bears the risk of gaps and manipulation. The resulting challenges were addressed by the development of a digital solution for fast (R1), orderly (R2), persistent (R3), counterfeit-proof (R4), restricted (R5) and business process compliant (R6) documentation.
During talks with domain experts, it turned out that the problem is not so much the lack of trust in a central authority, but rather the concern that the system would succumb to a single point of failure. From this point of view, a blockchain-based solution is still useful, but it will not exploit the potential for a trust-free decentralization for which Hyperledger Fabric was developed. Risks of an introduction could lie in a lack of acceptance and enforceability of the digitized business process among all participants as well as existing conflicting goals between the stakeholders - while MRO providers are interested in a long-term use of parts, a minimization of the scrapping rate and a secure and efficient secondary market, original equipment manufacturers profited from the insecure documentation process, which prevents a minimization of the scrapping rate and a secure secondary market due to more transparent information.
The BCS System could be adapted to any documentation process which is characterized by multiple stakeholders, multiple types of certificates and decentralized authorities. Examples are other transport industries such as shipping, rail transport or industries which also depend on a low repair turnaround time to ensure constant availability, such as the energy sector.
According to the evaluation strategy for "Technical Risk & Efficacy" (Venable et al., 2016), the next step is to evaluate the proof of concept in a more summative way in a more naturalistic environment. This means expanding the circle of domain experts and include a user role model. Further contributions to the knowledge base lie in the exploration of connections between the challenges of individual industries and their addressing by blockchain-based methods.

References

Accenture (2018). *BEYOND THE BUZZ: THE POTENTIAL OF BLOCKCHAIN TECHNOLOGY FOR AIRLINES*. Electronic Article. URL: https://www.accenture.com/us-en/insight-potential-blockchain-technology-for-airlines (visited on 11/27/2018).

Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., Enyeart, D., Ferris, C., Laventman, G., and Manevich, Y. (2018). "Hyperledger fabric: a distributed operating system for permissioned blockchains". In: *Proceedings of the Thirteenth EuroSys Conference*. ACM, p. 30. ISBN: 1450355846. DOI: 10.1145/3190508.3190538.

Baskerville, R. and Pries-Heje, J. (2014). "Design theory projectability". In: *Working Conference on Information Systems and Organizations*. Springer, pp. 219–232.

Beck, R., Czepluch, J. S., Lollike, N., and Malone, S. (2016). "Blockchain-the Gateway to Trust-Free Cryptographic Transactions". In: *ECIS*, ResearchPaper153.

Bellamy, W. (2017). *Air France KLM is Evaluating MRO Potential for Blockchain*. Electronic Article. URL: https://www.aviationtoday.com/2017/10/03/air-france-klm-evaluating-mro-potential-blockchain/ (visited on 11/27/2018).

Ben-Daya, M., Ait-Kadi, D., Duffuaa, S. O., Knezevic, J., and Raouf, A. (2009). *Handbook of maintenance management and engineering*. Vol. 7. Springer.

Buterin, V. (2014). Web Page. DOI: citeulike-article-id:14229692. URL: https://github.com/ethereum/wiki/wiki/White-Paper (visited on 04/17/2018).

Christidis, K. and Devetsikiotis, M. (2016). "Blockchains and Smart Contracts for the Internet of Things". In: *IEEE Access* 4, pp. 2292–2303. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2016.2566339.

Dolev, D. and Strong, H. R. (1983). "Authenticated algorithms for Byzantine agreement". In: *SIAM Journal on Computing* 12.4, pp. 656–666. ISSN: 0097-5397.

Economist, T. (2015). *The promise of the blockchain - The trust machine*. Electronic Article. URL: https://www.economist.com/leaders/2015/10/31/the-trust-machine (visited on 11/27/2018).

Ethereum (2017). *Consortium Chain Development*. Electronic Article. URL: https://github.com/ethereum/wiki/wiki/Consortium-Chain-Development (visited on 06/29/2018).

Evers, M. (2018). *Generating more transparency in aviation with blockchain technology*. Electronic Article. URL: https://www.lufthansa-industry-solutions.com/de-en/solutions-products/aviation/generating-more-transparency-in-aviation-with-blockchain-technology/ (visited on 11/27/2018).

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 2017*. ISBN: 978-0-9981331-0-2. URL: http://hdl.handle.net/10125/41339.

Greenspan, G. (2015). *MultiChain private blockchain—White paper*. Electronic Article. URL: https://www.multichain.com/download/MultiChain-White-Paper.pdf (visited on 11/27/2018).

Gregor, S. and Hevner, A. (2013). "Positioning and presenting design science research for maximum impact." In: *MIS quarterly* 37.2.

Hevner, A. and Chatterjee, S. (2010). "Design science research in information systems". In: *Design research in information systems*. Springer, pp. 9–22.

Hevner, A., March, S., Park, J., and Ram, S. (2004). "Design Science in Information Systems Research". In: *MIS quarterly* 28.1, pp. 75–105. DOI: 10.2307/25148625.

Hyperledger (2018). *Hyperledger Fabric Documentation*. Electronic Article. URL: https://hyperledger-fabric.readthedocs.io/ (visited on 11/27/2018).

IATA (2009). *IATA Technical Reference Manual*. Report. International Air Transport Association. URL: http://www.aerohabitat.eu/uploads/media/29-07-2010_-_IATA_-_Aviation_Technical_Reference_Manual.pdf (visited on 11/07/2018).

— (2015). *Best Practices for Component Maintenance Cost Management*. Report. International Air Transport Association. URL: https://www.iata.org/whatwedo/workgroups/Documents/MCTF/CMCM-2nd-edition.pdf (visited on 11/07/2018).

— (2018b). *Blockchain in Aviation: Exploring the Fundamentals, Use Cases, And Industry Initiatives*. Report. International Air Transport Association. URL: https://www.iata.org/publications/Documents/blockchain-in-aviation-white-paper.pdf (visited on 11/07/2018).

Love, C. and Guo, R. (1996). "Utilizing Weibull failure rates in repair limit analysis for equipment replacement/preventive maintenance decisions". In: *Journal of the operational Research Society* 47.11, pp. 1366–1376. ISSN: 0160-5682.

March, S. T. and Smith, G. F. (1995). "Design and natural science research on information technology". In: *Decision support systems* 15.4, pp. 251–266. ISSN: 0167-9236.

Nakamoto, S. (2008). "Bitcoin: A Peer-to-Peer Electronic Cash System". In: URL: https://bitcoin.org/bitcoin.pdf (visited on 09/14/2017).

Nofer, M., Gomber, P., Hinz, O., and Schiereck, D. (2017). "Blockchain". In: *Business & Information Systems Engineering* 59.3, pp. 183–187. ISSN: 2363-7005. DOI: 10.1007/s12599-017-0467-3.

Notheisen, B., Cholewa, J., and Shanmugam, A. (2017). "Trading Real-World Assets on Blockchain". In: *Business & Information Systems Engineering* 59.6, pp. 425–440. ISSN: 2363-7005.

Nowlan, F. S. and Heap, H. F. (1978). *Reliability-centered maintenance*. Report. United Air Lines Inc San Francisco Ca.

Provenance (2015). *Blockchain: the solution for transparency in product supply chains.* Electronic Article. URL: https://www.provenance.org/whitepaper (visited on 11/07/2018).

Schneider, F. B. (1990). "Implementing fault-tolerant services using the state machine approach: A tutorial". In: *ACM Computing Surveys (CSUR)* 22.4, pp. 299–319. ISSN: 0360-0300. DOI: 10.1145/98163.98167.

Simon, H. A. (1996). *The sciences of the artificial.* MIT press. ISBN: 0262264498.

Szabo, N. (1997). *Formalizing and Securing Relationships on Public Networks.* Electronic Article. URL: https://ojphi.org/ojs/index.php/fm/article/view/548/469 (visited on 06/29/2018).

Thakkar, P., Nathan, S., and Vishwanathan, B. (2018). "Performance Benchmarking and Optimizing Hyperledger Fabric Blockchain Platform". In: *arXiv preprint arXiv:1805.11390.*

Venable, J., Pries-Heje, J., and Baskerville, R. (2016). "FEDS: a framework for evaluation in design science research". In: *European Journal of Information Systems* 25.1, pp. 77–89. ISSN: 0960-085X.

Wickboldt, C. and Kliewer, N. "A framework for value-based pricing in aircraft surplus markets". In: *Multikonferenz Wirtschaftsinformatik (MKWI) 2016.* ISBN: 978-3-86360-132-4. URL: http://www.mkwi2016.de/download/MKWI2016_Research-in-Progress-Poster-Band.pdf.

— (2018a). "Blockchain zur dezentralen Dokumentation von Werkstattereignissen in der Luftfahrtindustrie". In: *HMD Praxis der Wirtschaftsinformatik*, pp. 1–14. ISSN: 1436-3011. URL: https://doi.org/10.1365/s40702-018-00452-y.

— (2018b). "Value Based Pricing meets Data Science: A Concept for Automated Spare Part Valuation". In: *Multikonferenz Wirtschaftsinformatik (MKWI) 2018.* ISBN: 978-3-935786-72-0. URL: http://mkwi2018.leuphana.de/wp-content/uploads/MKWI2018_Band1.pdf.

Wiliam, D. and Black, P. (1996). "Meanings and consequences: a basis for distinguishing formative and summative functions of assessment?" In: *British Educational Research Journal* 22.5, pp. 537–548. ISSN: 0141-1926.

Xu, X., Pautasso, C., Zhu, L., Gramoli, V., Ponomarev, A., Tran, A. B., and Chen, S. (2016). "The blockchain as a software connector". In: *13th Working IEEE/IFIP Conference on Software Architecture (WICSA).* IEEE, pp. 182–191. ISBN: 1509021310.

Xu, X., Weber, I., Staples, M., Zhu, L., Bosch, J., Bass, L., Pautasso, C., and Rimba, P. (2017). "A taxonomy of blockchain-based systems for architecture design". In: *IEEE International Conference on Software Architecture (ICSA).* IEEE, pp. 243–252. ISBN: 1509057293.

Zhang, Y. and Wen, J. (2017). "The IoT electric business model: Using blockchain technology for the internet of things". In: *Peer-to-Peer Networking and Applications* 10.4, pp. 983–994. ISSN: 1936-6442.

Zheng, Z., Xie, S., Dai, H., Chen, X., and Wang, H. (2017). "An overview of blockchain technology: Architecture, consensus, and future trends". In: *Big Data (BigData Congress), 2017 IEEE International Congress on*. IEEE, pp. 557–564. ISBN: 1538619962. DOI: `10.1109/BigDataCongress.2017.85`.

Chapter 6

Benchmarking a Blockchain-based Certification Storage System

Abstract

A comprehensive empirical study is performed to measure the performance of a Blockchain-based Certification Storage System in Hyperledger Fabric. This work is based on a proof of concept in the aviation industry and follows a Technical Risk & Efficacy evaluation strategy to determine the utility derived from the use of the artefact. Relevant tuning parameters for performance and scalability as well as bottlenecks are identified. The impact of configuration parameters such as blocksize, transaction arrival rate and number of concurrent users on the systems performance is investigated. Observations show that demands at throughput above system limits lead to transaction failures. Contributed are a repeatable process to performance sensitivity analysis and recommendations for configuring a Blockchain-based Certification Storage System for stable but high performance. The results can be used as a basis for optimizing the performance of similar systems.

Keywords Blockchain · Hyperledger Fabric · Benchmarking · Maintenance Events

Contents

6.1 Introduction and Related Work

Software Performance Engineering or Benchmarking includes efforts to describe performance changes in the system. These activities are divided into two approaches, an early predictive model (C. Smith, 2002) and a late measurement approach (Woodside et al., 2007). Barber (2004) argues that predictive model based approaches depend on a significant amount of emprical data which is rarely available when designing a new and innovative system. That is the reason why measurement approaches are popular. Examples for the measurement approach are Arlitt et al. (2001) for a large web-based shopping system or Avritzer et al. (2002) for a large industrial system.

In the aviation industry, safety-related spare parts require sophisticated and complete documentation of workshop events. Due to trust issues and no central institution which could handle digital workshop event certificates, this process hasn't been digitized yet. A blockchain-based concept to overcome these issues has been proposed in Wickboldt and Kliewer (2018a). This concept has been implemented and presented as an IT artefact called Blockchain-based Certification Storage System (BCSS) (Wickboldt and Kliewer, 2019). Workshops with domain experts show that the solution needs to be capable of handling an average of ∼0.16 transactions per second (tps) which equals around five million transactions a year.

A measurement approach is performed to ensure that the BCSS meets the requirements of the underlying business process. March and G. F. Smith (1995) state that the evaluation of a designed artefact requires measurement. Rudimentary benchmarks have been performed as part of the proof of concept in Wickboldt and Kliewer (2019) which showed that at least 7.99 tps are reachable. To ensure scalability, this report contributes comprehensive benchmarking results of the BCSS. An analytical evaluation in the sense of Hevner et al. (2004) is performed by examining the performance characteristics. This work extends the evaluation of Wickboldt and Kliewer (2018a) and Wickboldt and Kliewer (2019) on the path of the Technical Risk & Efficacy evaluation strategy (Venable et al., 2016). The step of performance benchmarking in an artificial environment is taken to prepare the transfer to a naturalistic environment. As the proof of concept in Wickboldt and Kliewer (2019) delivered first insights into performance, summative evaluations in an artificial environment are performed to gain deep understanding about blockchain configuration parameters and their impact on performance.

In order to arrive at these contributions the remainder of this report is organized as follows. Section 6.2 gives an overview about the benchmarking framework and experimental setup as well as configuration parameters of the benchmark. Benchmarking results are presented and discussed in section 6.3. A conclusion of this report and an outlook to future research is given in section 6.4.

6.2 Performance Benchmarking with Hyperledger Caliper

This contribution is orientated at the work of Thakkar et al. (2018) who provide insights about performance measurement in Hyperledger Fabric Blockchain systems using Hyperledger Caliper[1] which is a benchmark tool for Hyperledger Fabric. Reports produced by Caliper include:

R1 Throughput measured in transactions per second (tps),

R2 failed transactions due to timeouts,

R3 transaction latency,

R4 resource utilization.

Four methods for performing transactions to write and read workshop event certificates are implemented in BCSS. These methods receive parameters about the status of the spare part. *writeAsset* receives information via the frontend like serial number, Cycles Since New, Hours Since New, Cycle Limit, Hour Limit, Part Number, Part Description, Part Owner and Part Status. *writeCert* receives a serial number of the certificate, certificate type, certificate owner and details corresponding to the certificate. *queryAsset* receives the serial number. Using this, a JSON string with the information of the part is returned. *queryCert* gets the serial number of a part and returns the history of all certificates for that part.

For the purpose of this benchmark, a part filled with random data is generated (*writeAsset*). Next, a certificate for a workshop event for this part is generated (*writeCert*). The third and fourth operation are reading the asset's information (*queryAsset*) and reading the certificate's information (*queryCert*). Figure 6.1 visualizes a sequence of a Caliper benchmark.

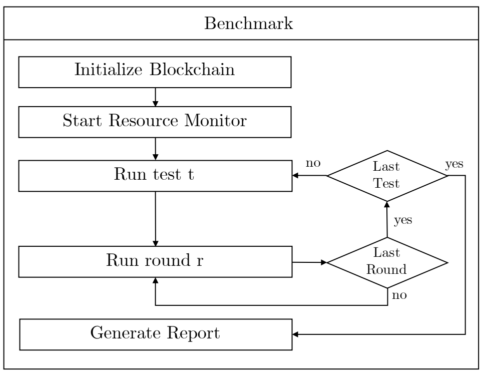

Figure 6.1: Sequence Diagram for Caliper Benchmark based on Caliper (2019)

[1] https://www.hyperledger.org/projects/caliper

The blockchain is initialized, the resource monitore is started. The configuration file[2] is consulted for the first test t. A test varies by block size and number of users whereas a round consists of multiple benchmarks for different transaction arrival rates (TAR). All rounds r are performed. This is done for all tests. After completion of the last test, a html benchmark report is generated.

6.2.1 Configuration Parameters and Experimental Setup

As this is a benchmark of a proof of concept in an artificial environment, a measurement approach is done on a development system. Tested are four operations on the BCSS. First, an asset is written onto the blockchain. Second, a certificate is issued on that asset. Next, the asset is queried. The fourth operation is querying the written certificate. The system comprises the following components: Intel(R) Xeon(R) Gold 6140 CPU @ 2.30GHz, 8 GB DDR4 RAM (ECC), 320 GB Serial Attached SCSI running Ubuntu 18.04 LTS. Hyperledger Fabric is running in version 1.4.

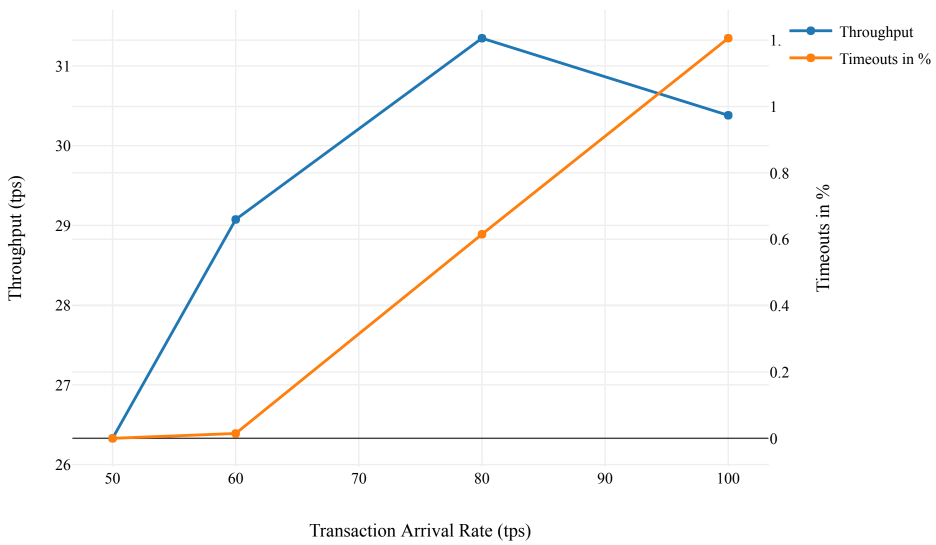

Figure 6.2: Transaction Arrival Rate on Throughput and Timeouts

[2]Sourcecode 6.1 in the appendix shows an example configuration file

In a preliminary analysis, the following has been observed (see figure 6.2): Average throughput doesn't necessarily increase with a higher TAR measured in tps. Transactions are increasingly terminated by timeout larger arrival rates. Only at an arrival rate of more than 50 tps, parameters like block size and number of users play a role to prevent timeouts. Throughput declines with arrival rates greater than 80 tps. The following parameters are therefore set for a sensitivity analysis: The length of the queue for benchmarked transactions is set to 1000 to ensure the benchmark runs some time to produce reliable and stable results. Benchmark runs are run with TAR of 50, 60, 80 and 100 tps.

A benchmark is conducted for block sizes of $2^n, n \in \{0, 1, 2, \dots, 10\}$ over a number of users $2^n, n \in \{0, 1, 2, \dots, 6\}$. A network with two organisations with each two peers is used. The benchmark runs are started via algorithm 1.

Algorithm 1 Start Benchmark

```
for All Block Sizes do
    for All User Sizes do
        procedure Set Blocksize(blockSize, FabricConfigurationFile)
            Open FabricConfigurationFile
            MaxMessageCount ← blockSize
            Close FabricConfigurationFile
            Rewrite Genesis Block(FabricConfigurationFile)
        end procedure
        procedure Set Number of Clients(numberClients, CaliperConfigurationFile)
            Open caliper-config.json
            number ← numberClients
            Close caliper-config.json
        end procedure
        procedure Shutdown Blockchain
            Shut down Docker Containers
            Remove Local State
            Remove Chaincode
            Sleep 60 Seconds
        end procedure
        procedure Startup Blockchain
            Startup Docker Containers
            Create Channel
            Join Peers to Channel
            Install Chaincode
            Instantiate Chaincode
        end procedure
        Start Benchmark(CaliperConfigurationFile)
    end for
end for
```

For every permutation the block size is set in the Hyperledger Fabric blockchain configuration. Timeouts are set to standard values of Hyperledger Fabric, for validating a transaction it is set to 60 seconds, for writing a block it is set to 10 seconds.

The genesis block is written respecting the block size of the Fabric configuration file. The number of clients is written into the Caliper benchmark configuration. After setting all parameters, eventually running blockchain services are shutdown. A sleep after executing the command ensures that all docker containers are indeed stopped. The blockchain is started, building onto the new genesis block. The channel for communication is created. All 4 peers join the channel. The chaincode for writing and reading certificates is installed and instantiated. Finally, the benchmarking is started with the current Caliper configuration.

This algorithm runs until every permutation of block size and current users has been benchmarked. All reports R1 to R4 are then retrieved by a Python Jupyter Notebook[3] for analysis. The results of the benchmark are presented in the following section.

6.3 Benchmarking Results

As Hyperledger Caliper delivers data about throughput (R1), failure rate (R2) latency (R3) and resource utilization (R4), this section presents and discusses benchmarking results for these four performance indicators. Throughput and latency results are presented along the parameters block size and number of concurrent clients. Performance measures are presented regarding memory and utilized disk space along tested block sizes.

Figure 6.3 shows Pearson correlation coefficients between the benchmarking results. A correlation coefficient of 1 suggests that the respective results correlate absolutely with eachother. A correlation coefficient of -1 means that they correlate absolutely against eachother. Block size and throughput are strongly positively (0.53) correlated to eachother, suggesting a larger block size improves the throughput. Average latency and throughput have a very strong negative correlation (-0.84), suggesting that higher throughput decreases latency. Higher user count leads to a higher probability for timeouts (correlation of 0.31). On the one hand, a larger arrival rate is slightly positive correlated with higher throughput (0.13). On the other hand, it is also positively correlated (0.14) with timeouts, suggesting that a load above the performance limits of the system lead to timeouts.

Next, numerical results of the benchmarks are presented and discussed using line plots. The data is grouped by the desired transaction arrival rate. The raw numbers are shown in a table next to the figure.

[3] https://jupyter.org

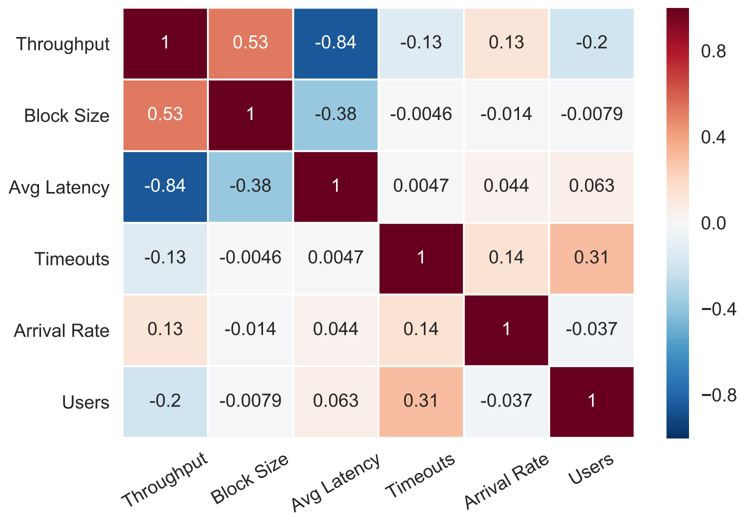

Figure 6.3: Correlation Heatmap for Benchmarking Results

6.3.1 Transaction Throughput and Failed Transactions

Transaction output is measured in transactions per second. Measured is the mean throughput along all blocksizes and calculated for any tested user count. Timeouts are calculated in % of all transactions. These two measured variables are considered together, since for a BCSS in production a high throughput only makes sense without transaction failures.

Concurrent Users

As shown in figure 6.4, the systems performance is decreasing with an increasing number of clients which are active at the same time. Read operations are generally quicker (∅ 32.17tps) than write operations (∅ 26.34 tps).

Concurrent Users	Throughput 100 tps write	Throughput 100 tps read	Throughput 80 tps write	Throughput 80 tps read	Throughput 60 tps write	Throughput 60 tps read	Throughput 50 tps write	Throughput 50 tps read
1	31.22	37.57	30.79	36.44	28.59	34.88	25.74	32.92
2	30.51	36.49	29.99	34.55	28.64	33.62	25.06	30.28
4	30.39	35.78	29.82	34.61	28.06	33.32	24.94	27.78
8	28.82	34.66	27.84	34.15	26.16	32.55	24.4	27.36
16	25.59	34.02	24.8	32.8	23.84	31.62	23.56	26.64
32	19.79	29.81	20.01	28.94	21.92	25.66	21.72	25.58

Table 6.1: Mean Throughput per Number of Users

Maximum throughput of 37.57 tps for reading at a TAR of 100 and 31.22 tps for

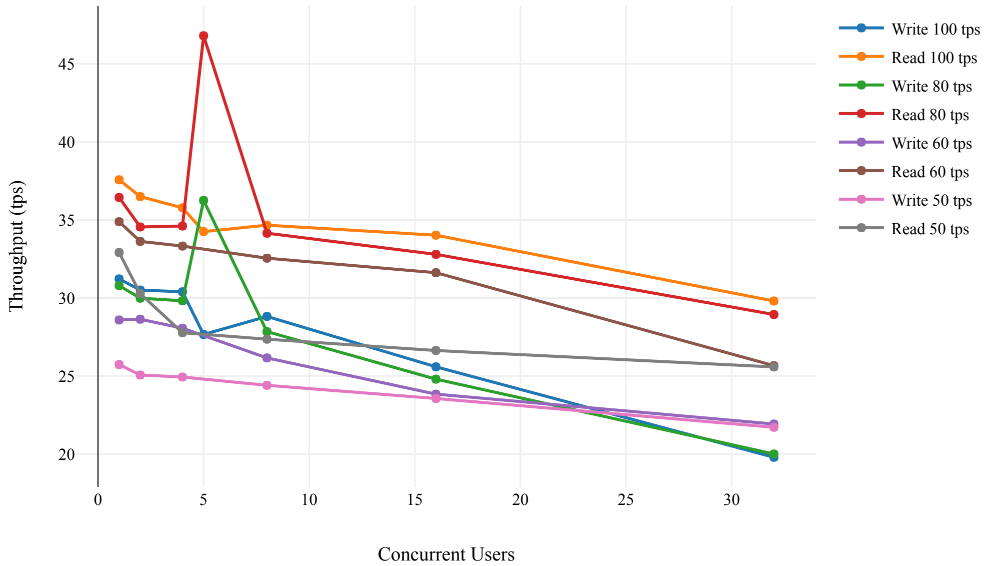

Figure 6.4: Mean Throughput per Number of Users

writing also at a TAR of 100 is reached with only one concurrently active client. If 32 users are active the minimum throughput is reached for reading at 25.58 tps (TAR 50) and 19.79 for writing (TAR 100).

Figure 6.5 shows the number of transactions which are not handled because of a timeout. Until 4 concurrent users, there are no timeouts at all for reading and writing operations. Writing operations tend to timeout more significantly, especially at a higher transaction arrival rate. The highest timeout rate for reading is at 100 tps desired throughput and 32 users (0.43%). The highest timeout rate for writing is also at 100 tps desired throughput and also 32 users (15.21%). The presented performance does not necessarily depend on the number of users but on the number of transactions in the transaction queue. This relationship between performance and fulfillment rate is related to the overall performance of the host system. If too many transactions are sent to the orderer with a high TAR or high number of users, transactions which are not handled after 60 seconds get discarded to ensure a steady performance overall.

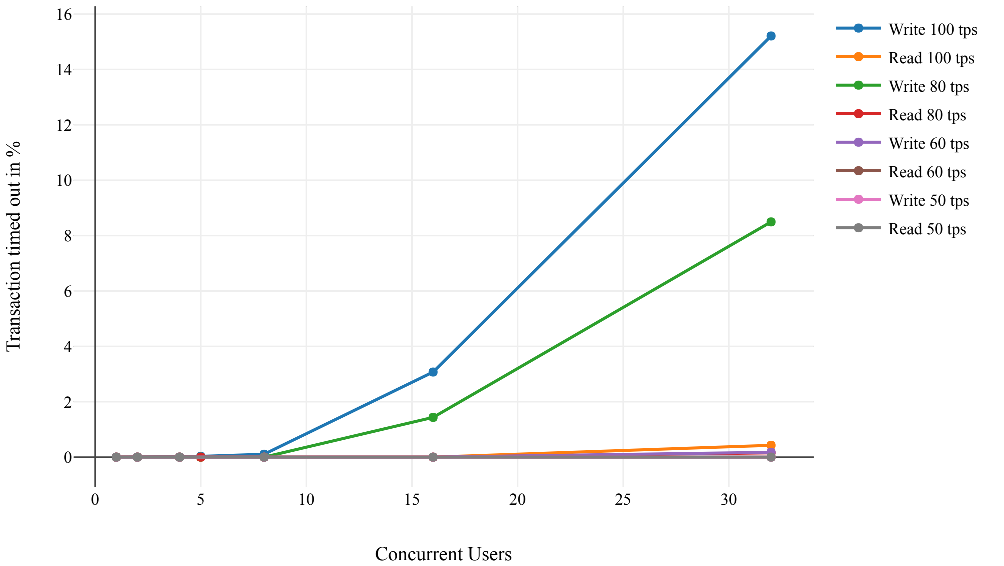

Figure 6.5: Mean Timeouts per Number of Users

Concurrent Users	Timeout % 100 tps write	Timeout % 100 tps read	Timeout % 80 tps write	Timeout % 80 tps read	Timeout % 60 tps write	Timeout % 60 tps read	Timeout % 50 tps write	Timeout % 50 tps read
1	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0
4	0	0	0	0	0	0	0	0
8	0.11	0	0	0	0	0	0	0
16	3.07	0	1.44	0	0	0	0	0
32	15.21	0.43	8.49	0.15	0.18	0	0	0

Table 6.2: Mean Timeouts per Number of Users

To summarize, high TAR and low number of clients lead to a high throughput. Timeouts can be prevented until 16 users and a TAR of 60 tps. This leads to a performance of 23.84 tps for writing and 31.62 tps for reading.

Block Size

Figure 6.6 shows the relation between transaction throughput and block size. Up to a block size of 128 throughout increases.

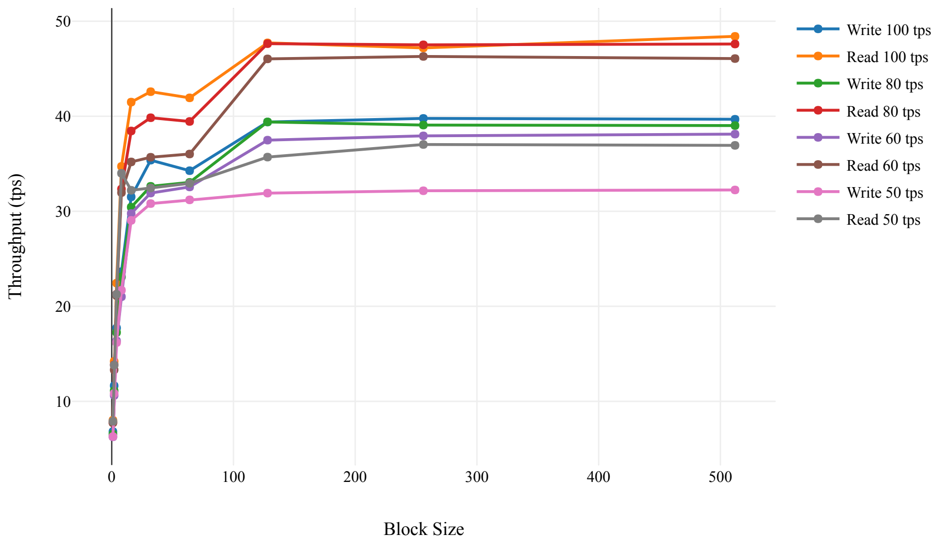

Figure 6.6: Mean Throughput and Block Size

Block Size	Throughput 100 tps write	Throughput 100 tps read	Throughput 80 tps write	Throughput 80 tps read	Throughput 60 tps write	Throughput 60 tps read	Throughput 50 tps write	Throughput 50 tps read
1	6.82	8.01	6.56	7.75	6.27	7.76	6.29	7.9
2	11.64	14.21	11.12	13.84	10.62	13.32	10.83	13.82
4	17.72	22.44	17.25	21.28	16.36	21.12	16.18	21.26
8	23.65	34.72	23.1	32.33	21.01	31.97	21.69	34
16	31.52	41.48	30.43	38.45	29.77	35.19	29.04	32.21
32	35.38	42.58	32.62	39.84	31.92	35.68	30.81	32.45
64	34.27	41.93	33.04	39.45	32.56	36.02	31.2	32.94
128	39.39	47.72	39.4	47.63	37.48	46.03	31.91	35.7
256	39.77	47.19	39.07	47.51	37.93	46.29	32.16	37.03
512	39.68	48.4	39.02	47.6	38.12	46.07	32.24	36.93

Table 6.3: Mean Throughput and Block Size

Maximum reading throughput of 48.4 tps is reached at TAR of 100 tps and a block size of 512. Maximum writing throughput of 39.77 tps is reached at TAR of 100 tps and a block size of 256. Minimum reading throughput of 7.75 tps is reached at TAR of 80 tps and a block size of 1. Minimum writing throughput of 6.27 tps is reached at TAR of 60 tps and a block size of 1. Also here, a higher transaction arrival rate results in higher throughput. Reading is faster than writing.

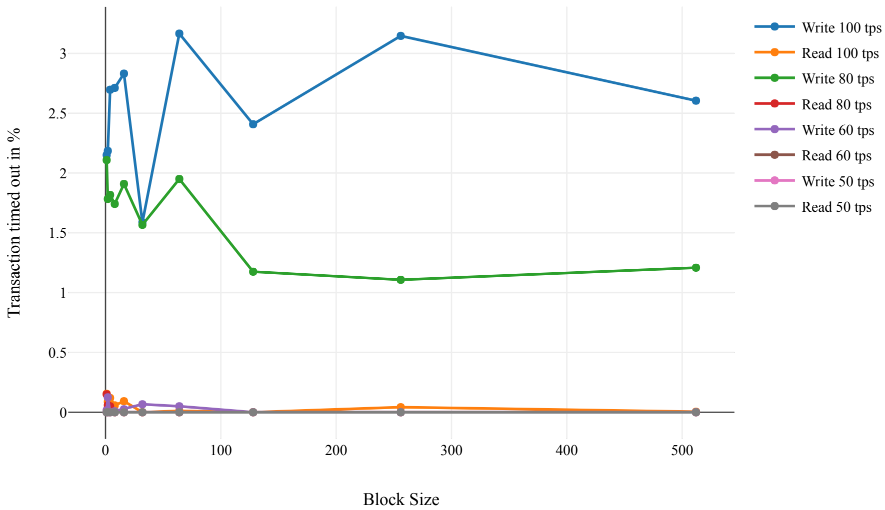

Figure 6.7: Mean Timeouts and Block Size

Block Size	Timeout % 100 tps write	Timeout % 100 tps read	Timeout % 80 tps write	Timeout % 80 tps read	Timeout % 60 tps write	Timeout % 60 tps read	Timeout % 50 tps write	Timeout % 50 tps read
1	2.15	0.15	2.11	0.15	0.02	0	0	0
2	2.18	0.09	1.78	0.06	0.12	0	0	0
4	2.7	0.12	1.82	0.05	0	0	0	0
8	2.71	0.06	1.74	0	0.01	0	0	0
16	2.83	0.09	1.91	0	0.02	0	0	0
32	1.59	0	1.57	0	0.07	0	0	0
64	3.17	0.01	1.95	0	0.05	0	0	0
128	2.41	0	1.18	0	0	0	0	0
256	3.15	0.04	1.11	0	0	0	0	0
512	2.6	0	1.21	0	0	0	0	0

Table 6.4: Mean Timeouts and Block Size

Figure 6.7 shows the relation between block size and failed transactions. Maximum timeouts for reading of 0.15% are reached at a TAR of 100 and block size of 1. The most timeouts for writing (3.17%) are experienced with a TAR of 100 and a block size of 64. Block sizes equal to or greater than 128 deliver stable performance at a TAR of 60 for reading and writing. Below that block sizes, timeouts occur. It can also be seen that runs with a transaction arrival rate greater 80 have significant timeouts, independent from block size.

For a high throughput a high TAR of 100 and large blocks of at least 256 transactions are recommended. Additionally, to prevent timeouts, block sizes smaller than 128 and TAR larger than 60 should be avoided. A block size of 128 and an expected arrival rate of transactions of 60 still leads to 37.48 tps in writing and 46.03 of reading which is close to the maximum.

6.3.2 Transaction Latency

Transaction latency is an indicator for the time an issued transaction is completed and a response is available to the application that issued the transaction. It is measured in miliseconds (ms). Latency depends on the amount of transactions in a block but also on concurrent users.

Block Size

As figure 6.8 shows, average transaction latency decreases rapidly until 32 transactions per block. This is true for all transaction arrival rates. Maximum average latency of 64.15 ms is reached at a TAR of 80 for reading operations with one transaction per block. Writing with a TAR of 60 and a block size of 1 leads to the maximum average latency of 81.46 ms. Minimum average latency of 2.9 ms is reached with a TAR of 50 and block size of 32 for reading and with a TAR of 50 and block size of 512 for writing (11.41 ms).

Block Size	Avg Latency 100 tps write	Avg Latency 100 tps read	Avg Latency 80 tps write	Avg Latency 80 tps read	Avg Latency 60 tps write	Avg Latency 60 tps read	Avg Latency 50 tps write	Avg Latency 50 tps read
1	75.27	63.55	76.55	64.15	81.46	62.4	78.09	59.36
2	45.64	35.01	46.8	35.48	47.17	34.99	44.87	31.64
4	29.74	21.25	30.08	21.9	30.78	20.02	29.02	17.7
8	21.77	13.7	21.47	13.36	22.35	12.84	19.91	9.61
16	18.01	9.53	17.78	10.13	17.09	8.02	15.21	4.93
32	15.41	7.98	15.84	7.47	14.89	5.75	12.59	2.9
64	15	7.87	14.91	7.33	13.74	5.93	11.8	3.57
128	14.31	7.86	13.78	7.2	13.45	5.11	12.01	3.69
256	14.17	7.87	14.54	7.56	13.26	4.94	11.84	3.49
512	14.12	7.83	14	7.4	13.36	5.17	11.41	3.81

Table 6.5: Average Latency and Block Size

If the recommended parameters from the throughput and failure tests of a TAR of

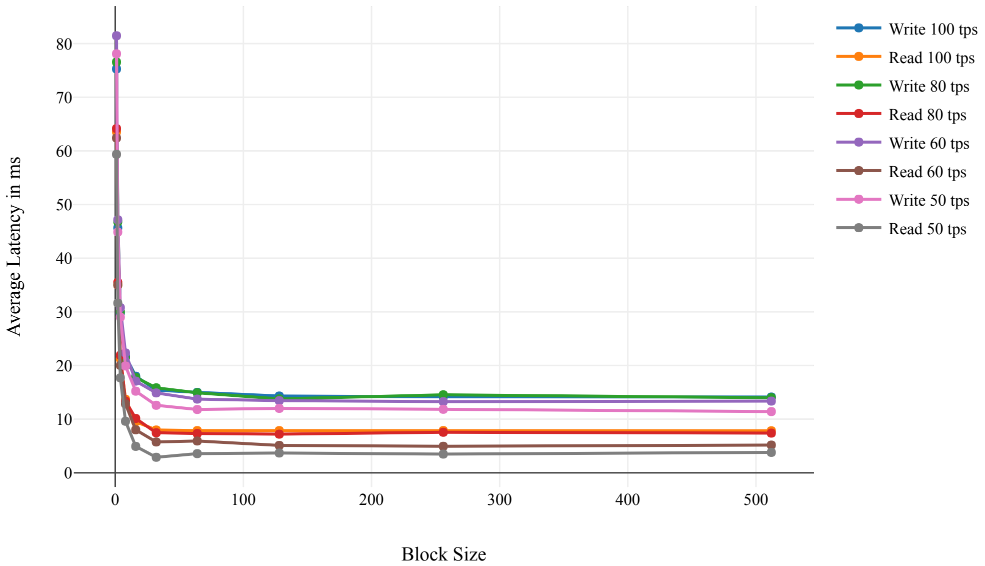

Figure 6.8: Average Latency and Blocksize

60 and a block size of 128 are considered, an average latency of 13.45 ms for writing and 5.11 ms for reading is observerd which is close to the minimum observations.

Concurrent Users

Figure 6.9 shows the relation between the number of concurrent users and the application's latency. Transaction latency increases with an increase of concurrent users in the network. This is especially true for writing operations across all transaction arrival rates.

Maximum average latency of 20.21 ms is reached at a TAR of 80 and 32 concurrent clients for reading operations. Maximum average latency for writing is reached at a TAR of 60 and 32 clients. Minimum average latency for reading of 11.6 ms is reached at 50 tps and one user. With the same parameters the minimum average latency for writing of 21.72 ms is reached.

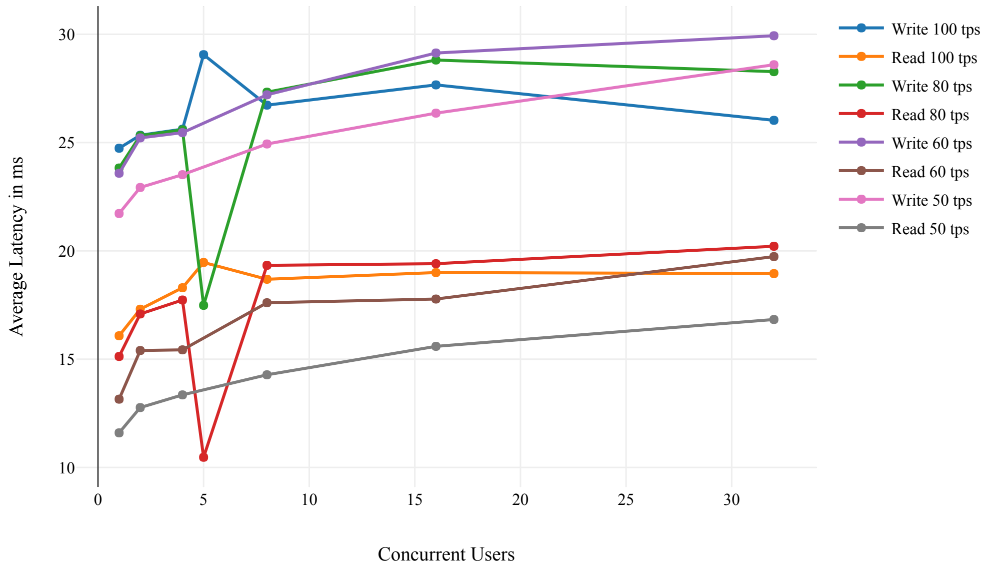

Figure 6.9: Average Latency and Number of Users

Concurrent Users	Avg Latency 100 tps write	Avg Latency 100 tps read	Avg Latency 80 tps write	Avg Latency 80 tps read	Avg Latency 60 tps write	Avg Latency 60 tps read	Avg Latency 50 tps write	Avg Latency 50 tps read
1	24.74	16.08	23.83	15.12	23.58	13.15	21.72	11.6
2	25.34	17.31	25.32	17.09	25.22	15.4	22.93	12.76
4	25.62	18.3	25.59	17.73	25.46	15.43	23.52	13.35
8	26.73	18.69	27.33	19.33	27.21	17.6	24.93	14.28
16	27.66	19	28.81	19.41	29.14	17.78	26.36	15.59
32	26.03	18.95	28.27	20.21	29.93	19.74	28.59	16.83

Table 6.6: Average Latency and Number of Users

The recommended user limit for preventing timeouts is 16 at 60 tps. If these parameters are taken into account a latency of 29.14 ms for writing and 17.78 ms for reading can be achieved.

6.3.3 Resource Utilization

Figure 6.10 shows the relation between block size and the highest number of RAM used within a benchmark run. It is observerd that a higher block size leads to a larger allocation of RAM. But this is only true until a block size of 128. Larger block sizes don't necessarily require more RAM.

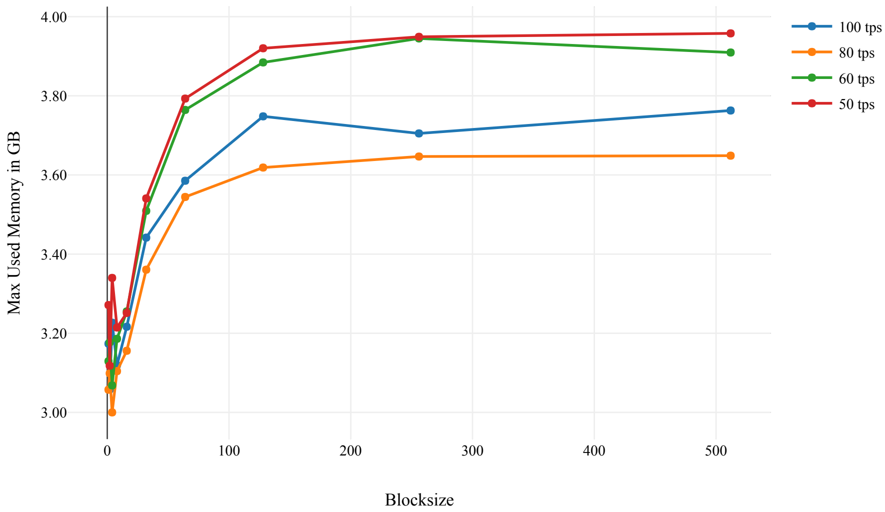

Figure 6.10: Maximum Used Memory per Block Size

Block Size	Memory in GB 100 tps	Memory in GB 80 tps	Memory in GB 60 tps	Memory in GB 50 tps
1	3.17	3.06	3.13	3.27
2	3.06	3.1	3.18	3.12
4	3.23	3	3.07	3.34
8	3.12	3.1	3.19	3.21
16	3.22	3.16	3.26	3.25
32	3.44	3.36	3.51	3.54
64	3.59	3.54	3.76	3.79
128	3.75	3.62	3.88	3.92
256	3.7	3.65	3.94	3.95
512	3.76	3.65	3.91	3.96

Table 6.7: Maximum Used Memory per Block Size

Used memory is at maximum of 3.96 GB in benchmarking rounds using a blocksize of 512 at 50 tps and lowest with 3.0 GB using a block size of 4 at 80 tps. This is probably due to the high throughput and low time outs which leads to the highest number of successfull transactions in the memory. The highest memory usage is seen with target arrival rates of 50 and 60 tps. This are rates at which there are no timeouts. Higher arrival rates of 80 and 100 tps use less memory but don't finish all transactions within time.

Figure 6.11 shows the relation between used disk space and block size. Larger block sizes lead to a lower need for disk space per transaction.

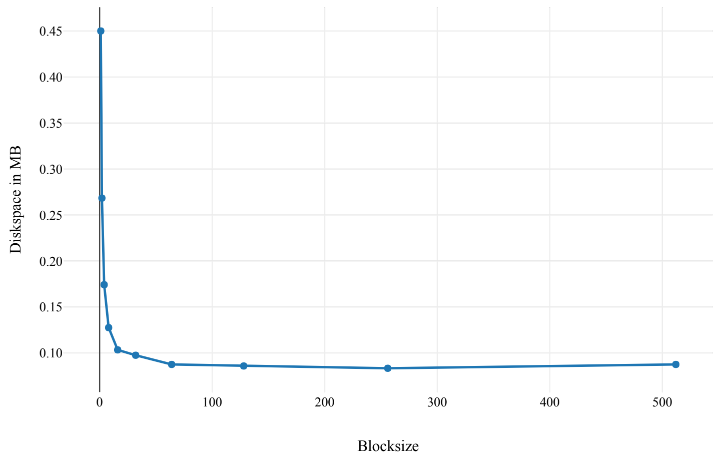

Figure 6.11: Used Disk Space per Blocksize

Block Size	Used Disc Space per Transaction in MB
1	0.45
2	0.27
4	0.17
8	0.13
16	0.1
32	0.1
64	0.09
128	0.09
256	0.08
512	0.09

Table 6.8: Used Disk Space per Blocksize

The least amount of disk space of 0.08 MB is needed when configuring the system with 256 transactions per block. Block sizes equal to or larger than 16 don't consume more than 0.1 MB of disk space.

Dependent on the performance of the hard drive but true in any case: Transactions with less need for disk space are obviously faster written and also faster read than transactions which require more disk space. This information is viable to understand why operations (especially writing) take longer when the block sizes are smaller. Regarding throughput and failure rate, a block size of 128 has been recommended. This would take 0.09 MB per transaction.

6.4 Conclusion and Outlook

A measurement approach of a Blockchain-based Certification Storage System has been performed to test the system for performance in an artificial environment. The artefact based on Hyperledger Fabric has been benchmarked with Hyperledger Caliper. Tested was a single orderer environment with four peers on one host. Results were given regarding throughput, failure rate, latency and resoure utilization.

It has been shown that the systems performance is decisively dependent on configuration parameters. Reading operations are generally faster than writing operations. A higher block size leads to a higher throughput and to less latency until the performance limits of the host system is reached. The most important factor for systems stability is transaction arrival rate, meaning the rate at which transactions reach the BCSS but not necessarily the rate the system can process them. If this rate is higher than what the system is capable of, timeouts increase and throughput decreases. It is therefore of great importance to scale the host system to the business network's requirement. This report has shown that the tested host is able to process 16 concurrent users with an arrival rate up to 60 tps. The recommended block size is 128 as there is no significant increase in throughput for larger block sizes without sacrificing system stability. With this configuration 37.48 tps when writing and 46.03 when reading are achievable on this system.

The benchmarking results are coherent with similar benchmarks. Timeouts could be decreased using 4 cores instead of 2 (Thakkar et al., 2018). Also throughput is greatly affected by the number of CPU cores. Thakkar et al. (2018) oberserve, considering their own configuration, 32 tps at 2 CPU cores and 848 tps at 16 cores. This is especially true when dividing communication into separate channels. However the main trend, that the number of concurrent users correlates negatively with the performance is also true in their work.

Although the throughput meets the requirements, the stable performance is only given to 16 simultaneous users. This is not enough for a realistic environment. As the research follows a Technical Risk & Efficacy evaluation strategy, future research should move the evaluation to a more naturalistic environment. Technically, this means using multiple servers to expand performance limits and increasing the number of peers. In this case, ordering transactions can't be done on one host but can be harmonized by a service called *kafka*[4]. Considering that the current system can deliver a competitive throughput, it can be assumed that the use of a BCSS to manage certificates of workshop events is scalable through further decentralization and thus provides a solid basis for real-world use cases.

[4] `https://hyperledger-fabric.readthedocs.io/en/release-1.4/kafka.html`

References

Arlitt, M., Krishnamurthy, D., and Rolia, J. (2001). "Characterizing the scalability of a large web-based shopping system". In: *ACM Transactions on Internet Technology* 1.1, pp. 44–69. ISSN: 1533-5399.

Avritzer, A., Kondek, J., Liu, D., and Weyuker, E. J. (2002). "Software performance testing based on workload characterization". In: *Proceedings of the 3rd international workshop on Software and performance*. ACM, pp. 17–24.

Barber, S. (2004). "Creating effective load models for performance testing with incomplete empirical data". In: *Web Site Evolution, Sixth IEEE International Workshop on (WSE'04)*. IEEE, pp. 51–59.

Caliper, H. (2019). *Hyperledger Caliper Architecture*. Electronic Article. URL: `https://hyperledger.github.io/caliper/docs/2_Architecture.html` (visited on 03/10/2019).

Hevner, A., March, S., Park, J., and Ram, S. (2004). "Design Science in Information Systems Research". In: *MIS quarterly* 28.1, pp. 75–105. DOI: `10.2307/25148625`.

March, S. T. and Smith, G. F. (1995). "Design and natural science research on information technology". In: *Decision support systems* 15.4, pp. 251–266. ISSN: 0167-9236.

Smith, C. (2002). "Software performance engineering". In: *Encyclopedia of Software Engineering*.

Thakkar, P., Nathan, S., and Vishwanathan, B. (2018). "Performance Benchmarking and Optimizing Hyperledger Fabric Blockchain Platform". In: *arXiv preprint arXiv:1805.11390*.

Venable, J., Pries-Heje, J., and Baskerville, R. (2016). "FEDS: a framework for evaluation in design science research". In: *European Journal of Information Systems* 25.1, pp. 77–89. ISSN: 0960-085X.

Wickboldt, C. and Kliewer, N. (2018a). "Blockchain zur dezentralen Dokumentation von Werkstattereignissen in der Luftfahrtindustrie". In: *HMD Praxis der Wirtschaftsinformatik*, pp. 1–14. ISSN: 1436-3011. URL: `https://doi.org/10.1365/s40702-018-00452-y`.

— (2019). "Blockchain for Workshop Event Certificates - A Proof of Concept in the Aviation Industry". In: *Proceedings of the 27th European Conference on Information Systems (ECIS)*. ISBN: 978-1-7336325-0-8. URL: `https://aisel.aisnet.org/ecis2019_rp/35/`.

Woodside, M., Franks, G., and Petriu, D. C. (2007). "The future of software performance engineering". In: *2007 Future of Software Engineering*. IEEE Computer Society, pp. 171–187.

Appendix

```
{"blockchain": {
		"type": "fabric",
		"config": "./fabric.json"
	},
	"test": {
		"name": "simple",
		"description" : "Caliper Benchmark that writes Assets, certs, retrieves Assets and Certs.",
		"clients": {
			"type": "local",
			"number": 1
		},
		"rounds": [
		{
			"label" : "open",
			"txNumber" : [1000, 1000],
			"rateControl" : [{"type": "fixed-rate", "opts": {"tps" : 100}}, {"type": "fixed-rate", "opts": {"tps" : 100}}],
			"callback" : "benchmark/simple/writeAsset.js"
		},
		{
			"label" : "open",
			"txNumber" : [1000, 1000],
			"rateControl" : [{"type": "fixed-rate", "opts": {"tps" : 100}}, {"type": "fixed-rate", "opts": {"tps" : 100}}],
			"callback" : "benchmark/simple/writeCert.js"
		},
		{
			"label" : "query",
			"txNumber" : [1000, 1000],
			"rateControl" : [{"type": "fixed-rate", "opts": {"tps" : 100}}, {"type": "fixed-rate", "opts": {"tps" : 100}}],
			"arguments": {  "partSerialNumber": "1"},
			"callback" : "benchmark/simple/queryAsset.js"
		},
		{
			"label" : "query",
			"txNumber" : [1000, 1000],
			"rateControl" : [{"type": "fixed-rate", "opts": {"tps" : 100}}, {"type": "fixed-rate", "opts": {"tps" : 100}}],
			"arguments": {  "partSerialNumber": "1"},
			"callback" : "benchmark/simple/queryCert.js"
		}]
	},
	"monitor": {
		"type": ["docker", "process"],
		"docker":{
			"name": ["peer0.org1.example.com", "peer1.org1.example.com", "peer0.org2.example.com", "peer1.org2.example.com", "orderer.example.com"]
		},
		"process": [
		{
			"command" : "node",
			"arguments" : "local-client.js",
			"multiOutput" : "avg"
		}],"interval": 1}}
```

Source Code 6.1: Example Caliper Benchmark Configuration File

Chapter 7

Decentralized Workshop Event Documentation with Hyperledger Fabric

Abstract

Due to the transparency and redundancy of blockchain, previously required intermediaries become obsolete, which enables the technology to align and digitize business processes in sectors with strong competition and low trust. On the example of the aviation industry, this contribution considers the application of blockchain in supply chains. Complete life cycle documentation is mandatory for safety-related aircraft parts. This work presents and evaluates an IT artefact storing the information of workshop event certificates in a Hyperledger Fabric blockchain. Our research responds to the calls for practical applications in the blockchain research field. An existing proof of concept is advanced towards a more naturalistic environment by decentralizing the system. The results suggest a growth in performance in regard to transaction throughput, latency and memory usage by distributing the system on different physical machines. Projectable patterns are identified that can be applied to a broad solution space in different industry sectors.

Keywords Blockchain, Hyperledger Fabric, Workshop Events, Aviation Industry, Benchmarking, Evaluation

Contents

7.1 Introduction

Blockchain technology as a combination of existing technologies including encryption, decentralized data management and consensus building was introduced 2008 (Nakamoto, 2008). It is the underlying core mechanism for Bitcoin and was first implemented in 2009. Blockchain is characterized by its decentralization, anonymity, auditability, persistency and therefore its ability to save costs and improve efficiency (Zheng et al., 2017). Today, blockchain technology is broadly considered a promising technology that can be applied to manifold applications beyond cryptocurrencies (Androulaki et al. (2018); Zheng et al. (2017)).

Due to the transparency and redundancy of blockchain, previously required intermediaries become obsolete. This enables the technology to be used in several business areas. There is already a great demand for blockchain solutions in a variety of industries, such as healthcare, fintech, real estate and government (Christidis and Devetsikiotis, 2016). There are also applications regarding the supply chain of the aviation industry (Madhwal and Panfilov, n.d.).

The *International Air Transport Association* (IATA) states that component maintenance costs account for the second largest share (24%) of airlines' direct maintenance costs (IATA, 2015). As a result of its special properties, the aviation industry is particularly suitable for digitizing processes with the help of blockchain. The safety of aircrafts and their components plays a central role in the aviation industry: To ensure this safety, a complete life cycle documentation for safety-relevant parts is required which is essential for further usage. The current analog process for documentation of workshop events is error-prone and leads to the devaluation of high-priced aircraft parts. Due to the intense competitive situation and the large number of stakeholders involved, the implementation of a central digital solution has not been successful yet. Therefore, Wickboldt and Kliewer (2018a) propose a decentralized solution based on the blockchain technology. The implementation of a Proof of Concept with the blockchain framework Hyperledger Fabric is described in Wickboldt and Kliewer (2019) called Blockchain-based Certification Storage System (BCSS). The artefact is benchmarked and evaluated in a laboratory environment in Wickboldt (2019).

Within this scope, the contribution of our research is threefold: First, by advancing a blockchain-based IT artefact towards a naturalistic environment, this research responds to the calls for applications in the blockchain research field to a great extent. Second, we find projectible patterns (Baskerville and Pries-Heje, 2014) in the evaluation of the instantiation that may be used to solve problems in other sectors. Third, the efficacy and efficiency of the proof of concept is proven in a more naturalistic environment.

This contribution builds on existing literature on a proof of concept and extends the evaluation according to the Framework for Evaluation in Design Science Research (FEDS) by Venable et al. (2016). To support evaluation research design decisions, the Technical Risk & Efficacy evaluation strategy has been chosen in Wickboldt and Kliewer (2019) and Wickboldt (2019) and is continued in this contribution. With this, the artefact is first developed and evaluated in an artificial laboratory environment and afterwards iteratively extended by more naturalistic requirements from the business perspective. In order to advance this artefact towards a more naturalistic environment, this paper aims to improve the artefact to increase its potential to be used in a productive system. Building upon the work of Wickboldt and Kliewer (2018a) and Wickboldt and Kliewer (2019), the present paper forms the fourth evaluation episode in the FEDS evaluation process. In this context, the goal is to increase the technical depth by addressing additional business needs by further decentralizing the blockchain network.

We rely on the design science research approach from Gregor and Hevner (2013) to structure the remainder of this paper as follows: In section 7.2 we give a systematic overview of the problem area and existing solutions. Furthermore, we provide insights into the business process for workshop event documentation in the aviation industry, as the underlying business environment for the development and evaluation of the BCSS. The third section presents the research methods used. Section 7.4 describes the design research process as an evolution of BCSS towards the resulting artefact. Design decisions made are permanently questioned and subjected to the requirements and restrictions of the business process (Hevner et al., 2004) in order to ultimately respect the laws of the existing environment (Simon, 1996). The results of the evaluation of the artefact are presented in detail in Section 7.5. Section 7.6 discusses and classifies these results. The section concludes with practical implications. This article concludes with a conclusion and outlook in Section 7.7.

7.2 Problem Description and State of the Art

7.2.1 Blockchain

Blockchain is a term for a distributed ledger, with the aim of correctly documenting a common state for all participants in a decentral network. The blockchain is a chain of *transactions* that are persisted in an immutable *ledger* in the form of data blocks (Swan, 2015). This network is maintained by mutually untrusting *peers*, while each peer owns a copy of the ledger. Transactions are validated by peers that execute a *consensus protocol*. In order to achieve consistency, the transactions are grouped into blocks and they are connected by a hash function.

The research field around Blockchain technology calls for "blockchain applications" (Zheng et al., 2017). Due to their globalized nature, business processes in supply chain and logistics are particularly in need for digitization. Dyer and Singh (1998) state that network participants agree to share company-owned insights if they can be certain that this is not shared with the competitors. Hence, there is a trust issue which must be overcome. Overcoming this trust issue was impossible when relying on a centralized solution. There are first applications of blockchain technology which currently transforms the industry.

The application of blockchain technology in maritime logistics has already been investigated and mapping of logistics processes via blockchain has been considered (Stahlbock et al., 2018). Various possible applications of blockchain in supply chains have already been researched, e.g. systems have been proposed which support the supply chains of agricultural products in the Asian region with blockchain with regard to traceability (Kumar and Iyengar (2017); Tian (2016)). In the energy sector, supply chain solutions on the basis of the blockchain technology are also considered more closely (Albrecht et al. (2018); Sikorski et al. (2017)). Moreover, there are concepts regarding the supply chain of the aviation industry (Madhwal and Panfilov, n.d.).

According to Zheng et al. (2017) blockchains fall into one of three categories: public, permissoned and consortium. This paper focuses on the consortium and private blockchain. They are characteized by the fact that network participants, in contrast to public blockchains, are not pseudonymized. In addition, they are permissioned, efficient and the consensus algorithm is either executed partially decentralized through the nodes of selected organizations, or centralized by a single instance. A recognized and common framework for private and consortium blockchain is Hyperledger Fabric.

7.2.2 Hyperledger Fabric

Hyperledger Fabric is an open source blockchain framework implementation and is run by the Linux Foundation. It is characterized primarily by its modular architecture, as well as its extensibility. Unlike public blockchains, which allow unknown entities to participate in the network, Hyperledger Fabric networks require participants to register with a trusted Membership Service Provider (MSP). Furthermore, Fabric is the first system to enable the development of distributed applications using General Purpose Languages without an existing technical dependency on a cryptocurrency (Androulaki et al., 2018). The framework is characterized by an

Execute-Order-Validate Blockchain architecture that differs from the regular Order-Execute design. The implementation of business processes is realized with the help of Smart Contracts or the chaincode. A Fabric blockchain consists of a network of nodes, which get their identity from an MSP. According to Androulaki et al. (2018), a node in a fabric network can have the following roles:

- *Clients* create transaction proposals for execution, orchestrate execution, and send transactions over the network to the Orderers.
- *Peers* execute and validate transaction proposals. All peers manage the blockchain ledger, that is, the entire transaction history and the current status of the ledger. Only a subset of the peers execute transaction proposals, which are called endorsers.
- *Orderer Service Nodes* (OSNs) form the Ordering Service and determine the sequence of the transactions that are attached to the ledger.

In Fabric, a peer consists of various components. For data persistence, each peer has a key-value database such as LevelDB or CouchDB. The ledger component manages the ledger and enables simulation, validation and ledger update phases. To achieve greater scalability, the peers use a gossip protocol to efficiently distribute the state of the ledger and data in the channel. Furthermore, peers have an endorser and committer component, which for example achieve the execution and validation of the chain code (Androulaki et al., 2018). Fabric also offers a built-in *Certificate Authority* (CA). A CA distributes digitally signed certificates, together with a public key, to the network participants. If a participant trusts the CA and knows the public key, it can verify the signature of another network participant using the key and ensure network membership.

The corresponding transaction flow diagram in Fabric is shown in Figure 7.1. The starting point is a client that initiates a transaction by sending a request. An application uses one of the supported SDK's (Node, Java or Python) and API's to create the transaction request. The request contains a function in the chaincode, which is to be called. The SDK ensures that the request has the correct format and is sent via gRPC, along with the user's cryptographic credentials, to the peers on the network. The endorsers then verify the transaction and, if correct, execute it. The chaincode is executed together with the transmitted arguments and the result is calculated from the current status of the ledger and returned as a response. The following step verifies this response and checks whether the endorsement policy has been met. The application also sends the transaction request and transaction response to the Orderer. The Orderer sorts the transactions chronologically according to channels and creates blocks with the transactions for each channel. A channel is a private subnetwork for the communication of at least two network participants.

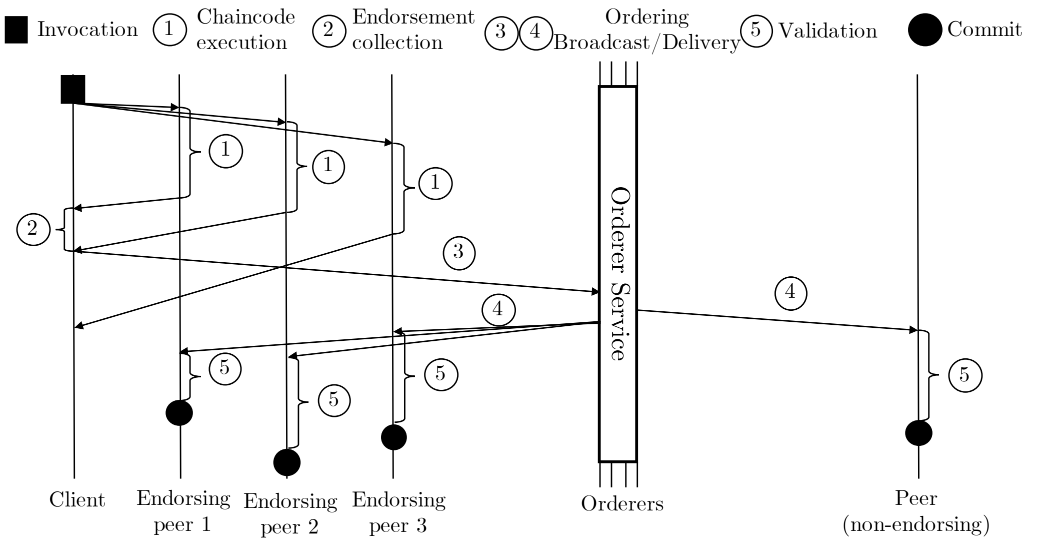

Figure 7.1: Transaction Flow Diagram

The transactions in the blocks are then validated and submitted. Finally, the ledger is invariably updated with the new transactions. A transaction is always triggered when a business process takes place. In the example of the aviation industry, this sequence would always run through when a stakeholder issues a certificate to an aircraft part and enters it into the BCSS.

7.2.3 Use Case: Documentation of Workshop Events in the Aviation Industry

The Maintenance Repair and Overhaul (MRO) industry in the aviation industry consists of a limited number of interest groups, in various business processes related to aircraft maintenance. Stakeholders include mechanics, original equipment manufacturers, workshops, airlines, traders and government agencies. Each stakeholder has different roles and has access to proprietary information within the industry. Each party must follow individual processes and release different data to document workshop events. Wickboldt and Kliewer (2019) propose a proof of concept for this use case, as a result of its high cost effects, the large number of stakeholders involved, the importance for aircraft safety and the need for highly reliable documentation.

Aircraft spare parts are distinguishable into safety-related and non-safety-related. Safety-related parts include turbines, landing gear and control components. Non-safety related parts are for example seats and other parts of the aircraft interior.

The focus of this work is on safety-relevant parts whose technical condition significantly influences flight safety. These parts require complete lifecycle documentation, also known as back-to-birth (BtB) documentation in order allow continued use and trade. Airlines often outsource maintenance of these parts and enter into contracts with *Maintenance Repair and Overhaul* (MRO) providers to ensure safe, fast and efficient maintenance (Wickboldt and Kliewer, 2018a). The documentation is currently being carried out with the help of certificates. These certificates are issued in paper form and physically kept in a container. Each workshop or trade event which requires documentation, another document is created and kept together with the rest of the documents. This method is highly susceptible to failure, as loss or damage to at least one of these documents results in immediate devaluation of the aircraft part. Without complete documentation, it is no longer allowed to trade the parts, or to build in an airplane. These requirements are mandated by the respective aviation authorities including the European Union Aviation Safety Agency (EASA, 2015) in Europe and the Federal Aviation Administration (FAA, 2005) in the USA. Failure of documentation results into scrapping of the parts which in turn leads to huge losses for the airlines, as the value of these parts, according to industry experts, take on average a six-figure amount. Furthermore, industry experts state that about 50% of all traded life limited parts have an erroneous certificate documentation and must be scrapped.

In order to support and represent the interest of airlines, the organization IATA was founded in 1945. It acts as a counterweight to the *International Civil Aviation Organization* (ICAO), which represents mainly state interests in aviation. Since its founding, IATA has coordinated and standardized many aspects of airline activities (Doganis, 2013). The aim of IATA is to represent, guide and serve the air transportation industry[1]. To this end, the organization develops global and commercial standards on which the aviation industry is built. One of these proposed processes is the Ideal Component Repair Cycle, which describes the cost-efficient course of repair of an aircraft part (IATA, 2015). The pattern process is shown in Figure 7.2. The goal of the process is to minimize turnaround time (TAT). This is defined as the elapsed time between removal of a component from the aircraft and return to the operator with the status functional capability. The amount of the TAT determines the cost of a repair, since the corresponding part indirectly incurs costs during this time due to the required replacement.

In each step of the process described in Figure 7.2, a certificate is issued. The following types of certificates have been identified by reviewing example certificates provided by domain experts: Parts List Report, Parts Profile Report, Parts Usage

[1] `https://www.iata.org/about/Pages/mission.aspx`

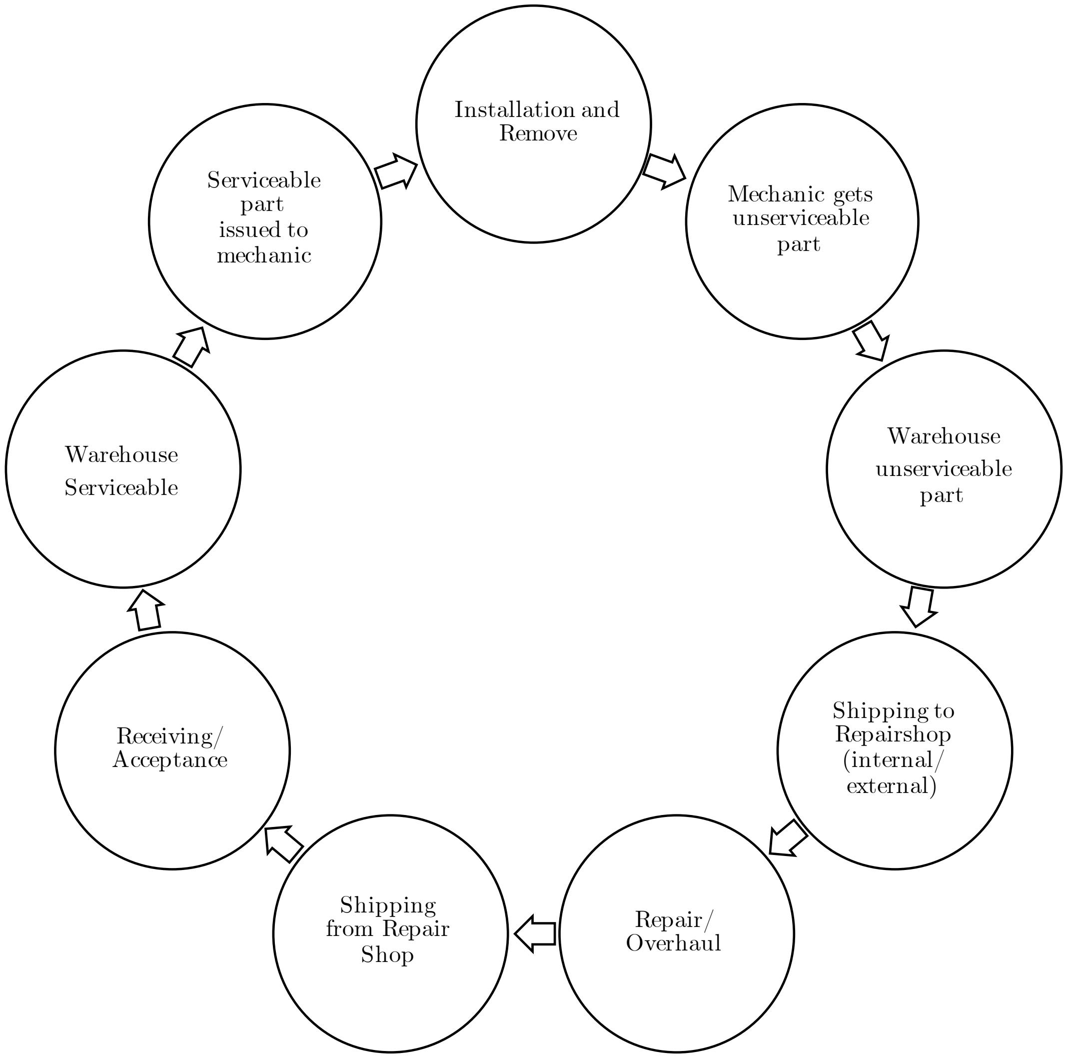

Figure 7.2: IATA Component Repair Cycle, based on IATA (2015)

Report, Record of Inspection, Material Certificate, Storage Report, Bill of Sale, Generic Report and Repair Report. These certificates are only a subset of the actual existing certificate types, but they cover the component life cycle adequately. The approach presented here was validated by means of a truthful documentation of an aircraft part from practice. The Parts List Report is a certificate that a manufacturer produces to prove the production of an aircraft part. Based on the model process of the IATA, the involved actors and roles can be determined for the identified certificates. In the case of the Parts List Report, this is the manufacturer.

The Parts Profile Report is used to document the installation and removal of aircraft parts. The person in charge is a mechanic. In Parts Usage Report, airlines document the life of aircraft parts. This is divided into cycles since new, i.e. the number of flight cycles since production and the time since new, i.e. the time of production. The Record of Inspection is used to document measurements of the dimensions of parts. The actual dimensions are compared with the given dimensions. The Storage Report is evidence of warehousing in a warehouse by a warehouse employee. The Material Certificate and the Bill of Sale document the sale. The Material Certificate is a check before the sale and the Bill of Sale describes the price, the buyer and the seller. The Repair Report confirms the actual repair of aircraft parts. In this case, a broken part is repaired and is then considered serviceable. The Generic Report is a document for the transport. Among other things, the responsible delivery service, the delivery destination and the delivery date are recorded.

7.2.4 Requirements and limitations of the business process

An artefact for the digital documentation of workshop events must meet certain requirements so that it can be used in practice. After consulting domain experts, Wickboldt and Kliewer (2019) suggest the following metrics that are necessary for a successful implementation:

(1) "Sales transactions must be executed immediately, others should be executed within 10 minutes. Confirmation of cycles per hour is only necessary at certain times; the speed of those queries is of secondary importance."

(2) "The system needs to be capable of handling a throughput of 5 Mio. transactions per year or 10 transactions per minute or 0.16 transactions per second (tps)."

Additionally, the following requirements are set:

(3) The system needs to be able to map the stakeholders in the aviation industry in the form of organizations on independent physical machines.

(4) The performance must remain stable within the expected number of users.

(5) The system must support a technical administrator role that is superordinate to other organizations.

7.3 Research Method

For structuring the creation, evaluation and presentation of the artefact, we rely on the *Design Science Research* (DSR) framework (Hevner et al., 2004). The viable IT artefact is a proof a concept prototype based on the blockchain technology that enables the documentation of workshop events in the aviation industry. The problem relevance is given, since a central digital solution has not been successful yet. Moreover, we ensure relevance by working with domain experts in the MRO industry. We contribute to the design evaluation by evaluating the IT artefact in a naturalistic environment. The contribution of our research is threefold: First, by advancing a blockchain-based IT artefact towards a naturalistic environment, this research responds to the calls for applications in the blockchain research field to a great extent. Moreover, according to Hevner et al. (2004) we contribute to the environment by evaluating an artefact that is designed to meet business needs, following the call for "blockchain applications" in Zheng et al. (2017). Second, we find projectible patterns (Baskerville and Pries-Heje, 2014) in the evaluation of the instantiation that may be used to solve problems in other sectors. Third, the efficacy and efficiency of the proof of concept is proven in a more naturalistic environment. We ensure research rigor by using well-established frameworks with Hevner et al. (2004), Gregor and Hevner (2013) and Venable et al. (2016). To get to the research contributions, our work is based on existing literature about blockchain applications (Androulaki et al. (2018); Nakamoto (2008); Thakkar et al. (2018)). Moreover, it is important that every artefact instantiation is scientifically evaluated and researchers should rigorously demonstrate quality, utility and efficacy of evaluation methods (Hevner et al. (2004); March and G. F. Smith (1995); Hevner and Chatterjee (2010)). In order to achieve this, this evaluation is based on the Framework for Evaluation in Design Science Research. We address technology-oriented audiences by providing insight into the technology choices, such as the blockchain framework, Docker Swarm and Apache Kafka. To address management-oriented audiences, we describe the underlying business process in the MRO industry.

The followed trajectory in a DSR project depends on respective needs and available resources and therefore there are different evaluation strategies that can be followed. Venable et al. (2016) state, that the functional purpose of evaluation can be either formative or summative. Moreover, paradigm of evaluation can be either naturalistic or artificial. Each of the possible strategies operate as a progression within these two dimensions from the origin towards a final evaluation. For this artefact, the greatest design risk is the underlying blockchain technology and is technically oriented. Furthermore, evaluations in a real system with the actual companies and stakeholders in the aviation industry are very expensive (Wickboldt and Kliewer, 2019). The quality of the artefact itself is to be evaluated and is the basis for a possible cooperation along the supply chain. Due to these circumstances, the *Technical Risk & Efficacy* evaluation strategy is chosen. The main aspect of this strategy is the early use of formative evaluations, in order to be able to influence design decisions and detect difficulties as early as possible and to therefore reduce costs and risks (Venable et al., 2016). This paper aims to provide a formative evaluation in order to improve the quality of the BCSS. Building upon the work of Wickboldt and

Kliewer (2018a), Wickboldt and Kliewer (2019) and Wickboldt (2019) the present contribution forms the fourth evaluation episode in the FEDS evaluation process. In this context, the goal is to increase the technical depth by addressing additional business needs.

Technically, evaluation is performed by performing benchmarks. This paper is based on the work of Thakkar et al. (2018), who provide insights into performance measurement and optimization leverage in Hyperledger Fabric Blockchain systems using Hyperledger Caliper, a benchmark tool for Hyperledger Fabric which was already used in Wickboldt (2019) for a simple network configuration. This network is now enhanced to meet the additional requirements.

7.4 Artefact Description

7.4.1 A Blockchain-based Certification Storage System

The BCSS consists of three modules that interact with each other. The communication between the modules takes place using the JSON data format. The system implements four methods for performing transactions to write and read workshop event certificates. These methods receive parameters about the status of the spare part. Function F1 *writeAsset* creates a new aircraft part with information such as serial number and part owner. An aircraft part is a digital asset that is stored in the Blockchain. Subsequently, F2 *writeCert* is used to append a certificate to the history of an asset. As a result of a variety of different certificates in the real world, the information is abstracted with a generic *detail* field. The stored data of these methods can be read with functions F3 *queryAsset* and F4 *queryCert*, in order to receive the current state of the digital asset or the complete certificate history.

The following sequence diagram shows the information flow between the individual modules when creating an aircraft part in the blockchain (see Figure 7.3). It shows the process from the input of the user in the frontend to the feedback from the blockchain. After the registered user has filled out and confirmed a form for creating a new aircraft part using the user interface, the entered information is transferred to the *createCertificate* method implemented in the backend in form of a JSON string. A transaction request is then created from the transferred data. This contains additional information, such as the chaincode id with the desired version and the identification of the client. The transaction request is then sent to one or, if necessary, several peers. The method for creating a certificate is now executed in the chaincode. If this has been successfully completed, the peers execute this transaction. The current state of the database is considered to generate transaction results. This result includes a write and read set, as well as a response value. The transaction is then verified by at least one endorsing peer. For example, the endorsing policy checks whether the sender has sufficient rights to send the transaction. If an existing authorization can be verified, a successful response to the backend, and distributed it to all peers using the Orderer. The Ordering Service takes care of the chronological order of the transactions within the channels. It creates the blocks with the transactions for each channel. The transactions in the blocks are also marked for validity. Each peer now adds the new block to its chain and updates the changes to the valid transactions in its state database (here CouchDB). An event about the completion and writing of the transaction is then created and the result sent to the backend. This is then forwarded to the frontend and the user is presented with feedback on the success of the transaction via the user interface.

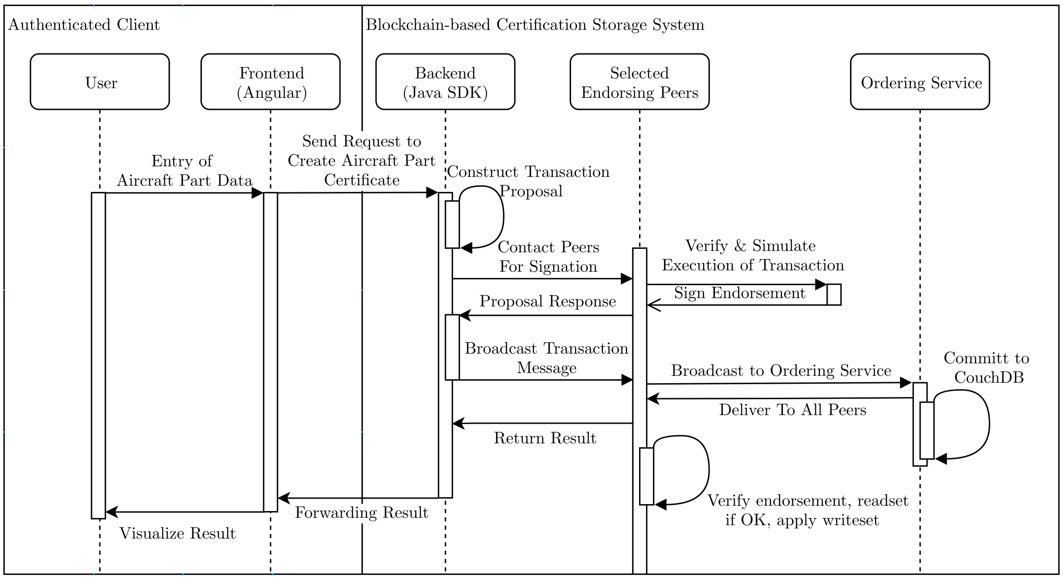

Figure 7.3: Sequence Diagram

7.4.2 Decentralized Architecture

A central aspect of blockchain is decentralization. Based on the Technical Risk & Efficacy evaluation strategy, the artefact is first developed and evaluated in an artificial laboratory environment and afterwards iteratively extended by more naturalistic requirements from the business perspective. For the purpose of simplicity, the artefact has been evaluated formative and artificial on a single machine (Wickboldt, 2019). With this next evaluation step, the BCSS is now distributed over several machines. Hyperledger Fabric offers the possibility to operate the Ordering Service in a Kafka mode. A setup based on Apache Kafka will be used to order the transactions from the peers. This instance of Apache Kafka offers consistency despite node crashes and scalable publish-subscribe messaging (Androulaki et al., 2018). In Fabric, each channel is assigned to a separate single-partition topic in Kafka. The OSNs function as proxies between Kafka and peers. Moreover, they are independent of Kafka and could run on other physical machines. The Kafka setup consists of multiple nodes in a Kafka Cluster and a *ZooKeeper* ensemble. The minimum number of nodes in the Kafka Cluster is four, in order to be fault tolerant. If one of the brokers crashes, all existing channels are readable and writable, and it is possible to create new channels. The number of nodes in the ZooKeeper ensemble is either three, five, or seven. In order to avoid split-brain scenarios, it must be an odd number, and it has to be larger than one to avoid having a single point of failure.[2]

In order to test the effects of decentralization on performance, the complexity of the architecture is iteratively increased. The first step is to run the Kafka cluster on a single machine and compare it with the previous solo mode. The network is then divided between two separate physical machines. The individual components are each operated in docker containers. Docker Swarm is used so that they can still communicate with each other. A virtual network is set up by the first machine M1, the Swarm Manager. Additional machines can join this network with a unique token and the IP address of M1. The containers can then be operated within this network and can be distributed via the hostnames of the machines. The communication between the containers is then within this network as if it were on a single physical machine. As an alternative to this procedure, *Kubernetes* can be used. The containers are orchestrated in a master-slave architecture. Since this requires a further external dependency, which can influence the performance, the Docker Swarm structure was preferred for the sake of simplicity.

Two Fabric Orderers, a Kafka cluster consisting of three ZooKeeper nodes and four Kafka brokers, a CA and two peers are operated on machine M1. Only two peers and one CA are operated on the second machine M2. This setup has a single point on failure with M1, as it contains the entire Kafka cluster. The next step is therefore to examine the performance of a three-part network. The Kafka components are distributed to three machines M1, M2 and M3. Each machine also has two peers and a CA. The machine M1 additionally runs two Orderers. Due to the redundancy and distribution of the Kafka cluster and the CAs, this setup already largely corresponds to a realistic architecture. However, the three machines in this setup are provided by the same provider, which may have a positive effect on latencies. In addition, the performance and the operating system of the machines are almost identical. In the last step the network will be extended to 4 organizations. A fourth physical

[2] `https://hyperledger-fabric.readthedocs.io/en/release-1.4/kafka.html`

machine M4 is added, which is outside the network. Furthermore, the two Orderers are distributed to the machines M1 and M4. From a technical point of view, this structure corresponds to the requirements of reality and could be used in this form in a business application. Figure 7.4 shows the technical architecture with four physical hosts, their respective components and their intercommunication.

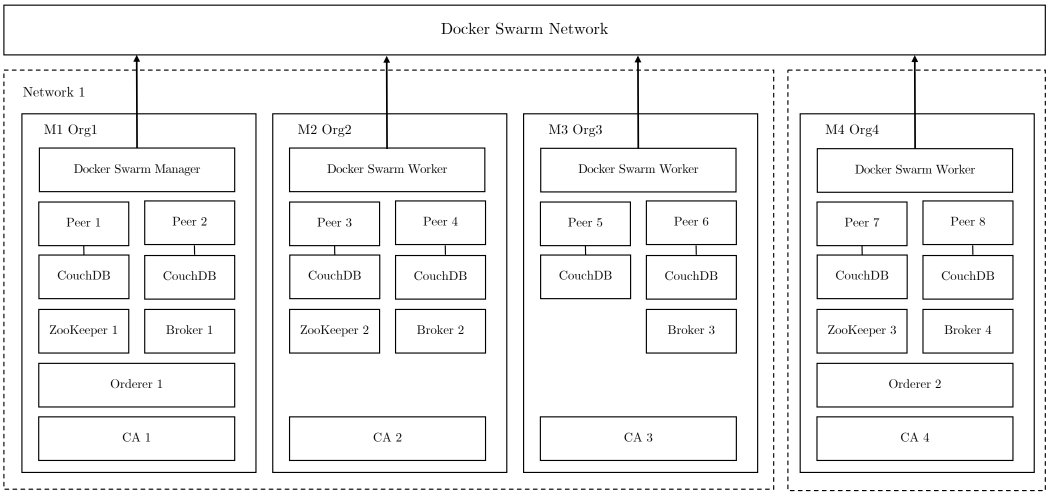

Figure 7.4: System architecture with four physical hosts

7.5 Evaluation

In order to prove that the artefact meets the business requirements, utility, quality, and efficacy of the artefact (Hevner et al., 2004) is demonstrated via performance benchmarks. Benchmarking is performed by using the Hyperledger Caliper framework. Reports created by Caliper include the transaction throughput in second (tps), failed transaction due to timeouts, transaction latency and resource utilization.

As a part of this benchmark, an aircraft part filled with random data is generate with F1. Afterwards, method F2 is used to create a certificate for a workshop event that is attached to the history of the previously created asset. The third and fourth operations are reading the system information (F3) and reading the certification information (F4).

7.5.1 Configuration Parameters and Experimental Setup

The benchmarks are conducted on four machines, where M1, M2 and M3 are in the same network and M4 is running on a separate network. With this, the architecture is more complex and is closer to conditions in a possible real-world application where the machines are completely independent of each other. Therefore, a benchmark with the external machine M4 may problems with higher latencies as a result of the network separation. The machines M1, M2 and M3 have the following components: Intel(R) Xeon(R) Gold 6140 CPU @2.30GHz, 8 GB DDR4 RAM (ECC), 320 GB Serial Attached SCSI. M1 is running on Ubuntu 18.04 LTS while M2 and M3 are running on Debian GNU/Linux 8 (jessie). Again, this aspect serves to ensure that the system is as realistic and heterogeneous as possible. Moreover, Machine M4 has a different hardware configuration and consists of the following: two cores of an Intel(R) Core (TM) i7-3930K CPU @ 3.20GHz, 4GB DDR4 RAM and a 300GB HDD. Hyperledger Fabric is running in version 1.4 on all machines.

This benchmark provides insights about the performance of Fabric under various conditions and helps to understand how certain parameters affect the performance regarding the underlying business process of storing workshop event certificates. Since the performance from a peer's perspective has already been studied thoroughly (Androulaki et al. (2018),Thakkar et al. (2018)), this paper focuses on the Orderer and network perspective. In all tests the endorsement policy is constantly *"OR('Org2MSP.admin')"* and requires a signature. This means that there is no additional computing power required to fulfill the endorsement policy and instead all requests that come from the administrator are fulfilled by default. There is only one channel in all setups and the peer database is CouchDB.

7.5.2 Setup configurations

The tests for each setup are repeated with different configuration parameters. The number of transactions per test round is constantly 1000. The transaction arrival rate, also called send rate, stands for the number of chaincode calls that Hyperledger Caliper sends to the system per second. In Hyperledger Fabric the Orderers control the number of messages batched into a block. This parameter is called block size. Moreover, the batch timeout is configured to be 2 seconds, which is the amount of time to wait before creating a batch. A benchmark is conducted for the transaction arrival rates $t \in 20, 40, 50, 60, 80, 100$, the number of concurrent users $2^c, c \in 0, 1, 2, 4, 8, 16$ and the block size $2^b, b \in 1, 2, 4, 8, 16, 32, 64, 128, 256, 512$. The configuration of the tested systems is shown in Table 7.1.

Number of Machines	Orderer Mode	Number of Orgs	Number of Peers	Number of Orderers	Number of ZooKeepers	Number of Kafka Brokers
1	Solo	2	4	1 (M1)	-	-
2	Kafka	2	4	2 (M1)	3 (M1)	4 (M1)
3	Kafka	3	6	2 (M1, M2)	3 (M1, M2, M3)	4 (2 on M1, M2, M3)
4	Kafka	4	8	2 (M1, M4)	3 (M1, M2, M4)	4 (M1, M2, M3, M4)

Table 7.1: Setup configurations

7.5.3 Benchmarking Results

Hyperledger Caliper provides data about throughput, latency, failure rate and resource utilization. The results of the benchmarks with these performance indicators are presented in the following. Throughput and latency results are shown along the parameter block size and number of concurrent users. The results of disc utilization are presented along block sizes.

Timeouts per Throughput and Transaction Arrival Rate

Figure 7.5 shows the correlation between transaction throughput and the relative number of timeouts for different transaction arrival rates.

Figure 7.5: Timeouts per Throughput and Transaction Arrival Rate

Setup performances

The results suggest that the timeouts multiply when the maximum system load is reached. For the machines used in this benchmark the average maximum for all tested configurations is at approximately 80 transactions per second.
Figure 7.6 shows the performance of the system environments described in Table 7.1. The transaction throughput is measured for different numbers of concurrent users. The results show that write operations have a lower throughput than read operations and require more computing power. The setup on a single physical machine with the Solo Orderer has the lowest overall throughput with an average of 24.47 tps for read and 17.28 tps for write operations. The single host setup with the Kafka Orderer showed a significant performance growth with an average of 31.30 tps for read and 21.46 tps for write operations, performing 26.39% better on average. By distributing the system on two physical machines, the performance increases to an average of 41.65 tps for read operations and 34.12 tps for write operations. As a result, the mean throughput increased by 43.60% if a second physical machine is used. The results suggest that a distribution of the system on three machines is nearly identical, with a difference of only 0.59 tps on average. The performance increases slightly (15.4%) when the fourth machine is added to the system, which is a result of the faster CPU in M4.

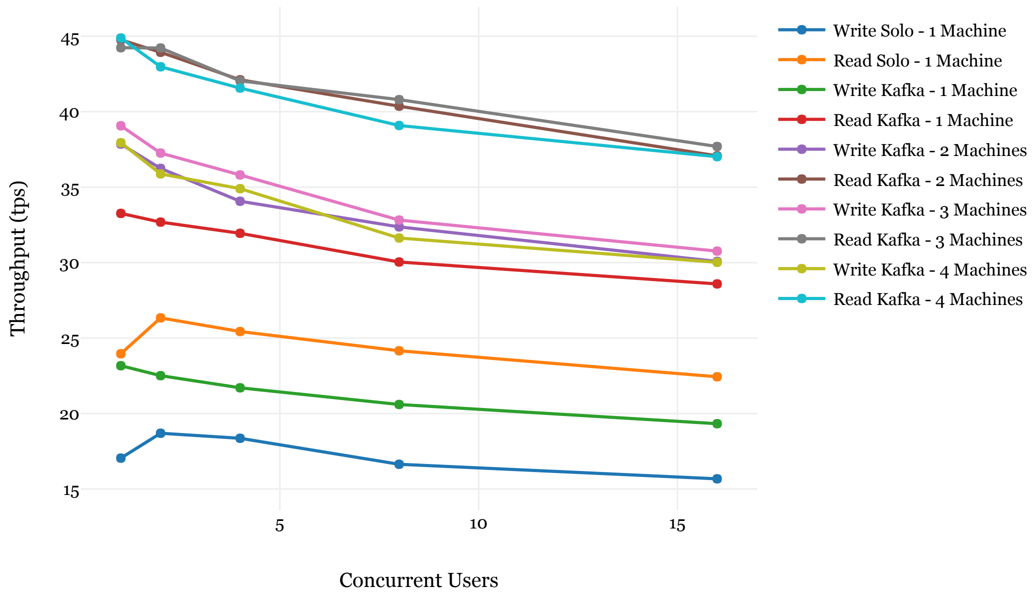

Figure 7.6: Mean Throughput per Number of Users by Setup

As shown in Figure 7.7, the average memory usage is dependent on the block size and the system setup. The maximum was reached with the Solo configuration and a maximum memory usage of 170.48 MB for a block size of 512. The Kafka configuration on 2 physical hosts showed the lowest overall maximum memory usage with 95.28 MB with 2 transactions per block. For the distributed system with 3 and 4 physical machines, the minimum average memory usage is higher. All configurations reach the maximum with a block size of 128.

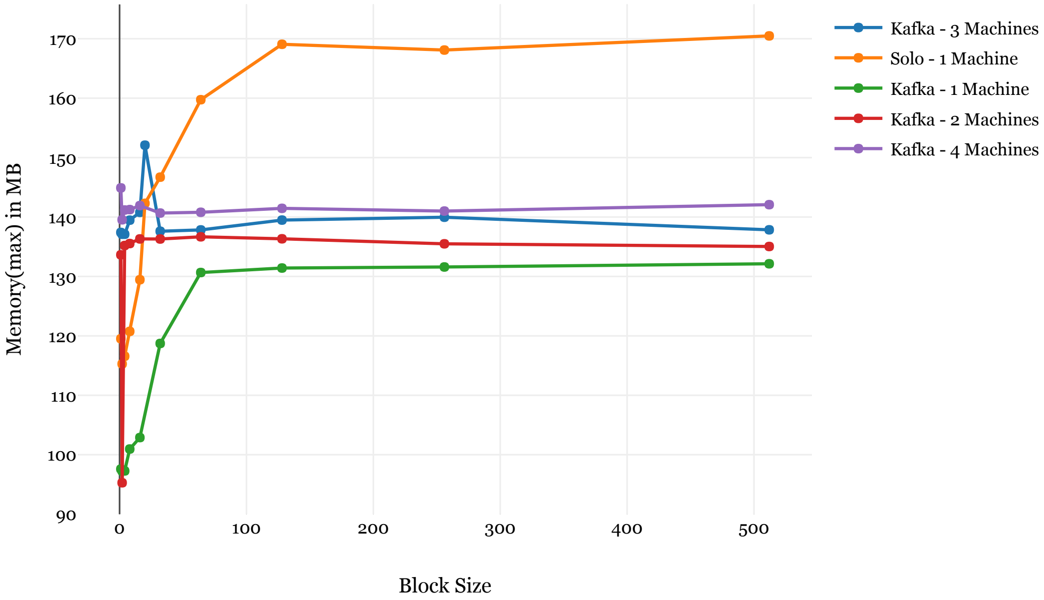

Figure 7.7: Mean Used Memory per Block Size

Parameter configuration

Figure 7.8 shows Pearson correlation coefficients between the benchmarking results. A correlation coefficient of 1 indicates that the respective results correlate totally with each other, while a correlation coefficient of -1 suggests that they correlate totally against each other. The block size and throughput are strongly positively (0.3) correlated to each other, indicating that a larger block size leads to a lower throughput. Average latency and throughput have a strong negative (-0.46) correlation, suggesting that a higher throughput increases the average latency. Arrival rate and throughput are also strongly positively (0.57) correlated to each other. Users and throughput have a slightly positive (0.22) correlation, suggesting that a higher number of users leads to a higher number of timeouts. Moreover, the number of concurrent users and throughput correlated slightly negatively (-0.14) to each other.

A higher number of users therefore indicates a lower throughput.

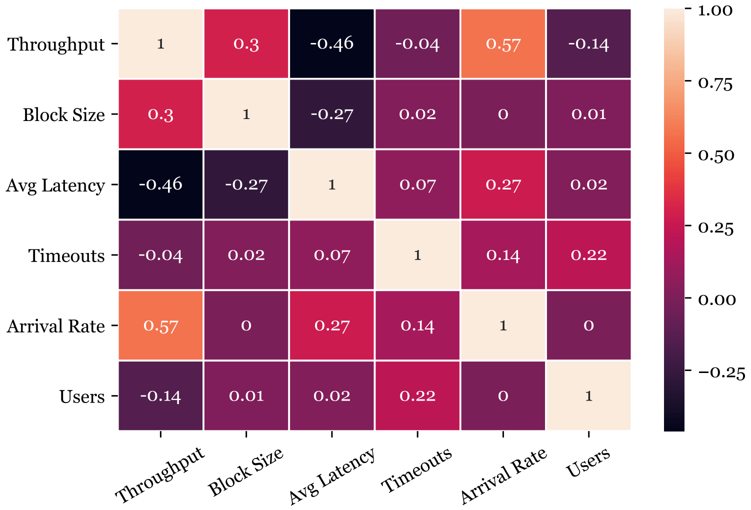

Figure 7.8: Correlation Heatmap for Benchmarking Results

Transaction Throughput and Failed Transactions

Figure 7.9 shows the relation between the average transaction throughput and the number of concurrent users. Write and read operations were tested with transaction send rates of 20, 40, 50, 60, 80 and 100 transactions per second. The results indicate that the systems performance is decreasing slightly with an increasing number of concurrently active clients. The largest decrease in performance was measured for a send rate of 100 read operations per second. The maximum throughput was 64.07 tps on average for 1 user and the minimum was at 49.23 tps for 16 concurrent users, a performance reduction of approximately 23.16%. In contrast, for a send rate of 20 read operations per second the throughput decreases by about 19.40%, which is equal to approximately 3.82 tps.
As shown in Figure 7.10, the number of concurrent users and the number of transaction timeouts have a positive correlation. For write operations with a high send rate of 80 and 100 tps, the number of timeouts increases significantly when testing with more than 16 and more than 8 concurrent users, respectively. Another benchmark with a block size of 128 with the configuration on 4 machines suggested that the maximum number of users for the test systems is 45, with an average fail rate over all send rates of more than 9.46%. Tests with 32 users showed an average fail rate of about 3.43%.

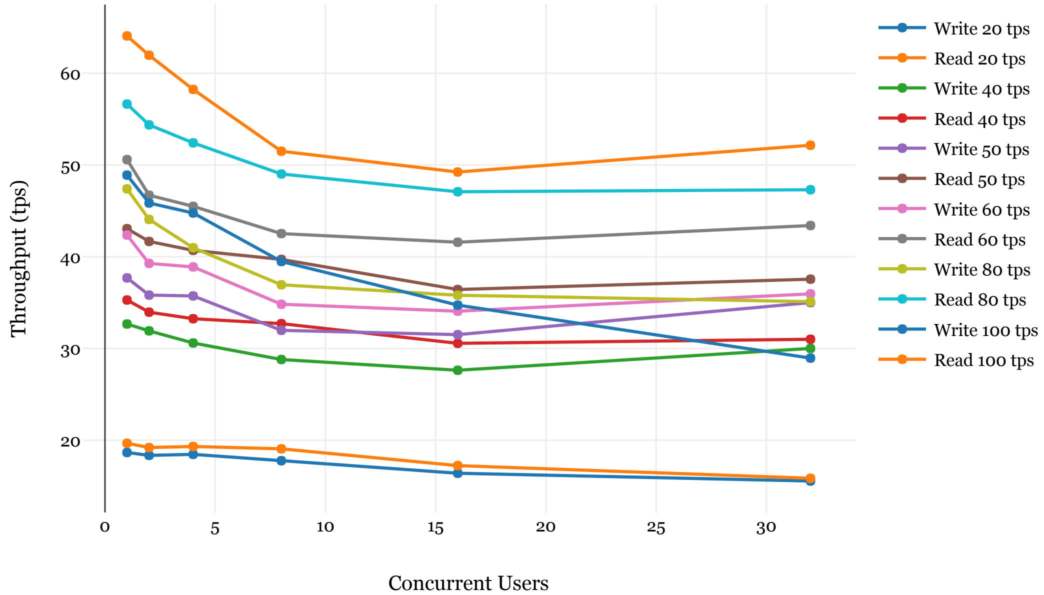

Figure 7.9: Mean Throughput per Number of Users

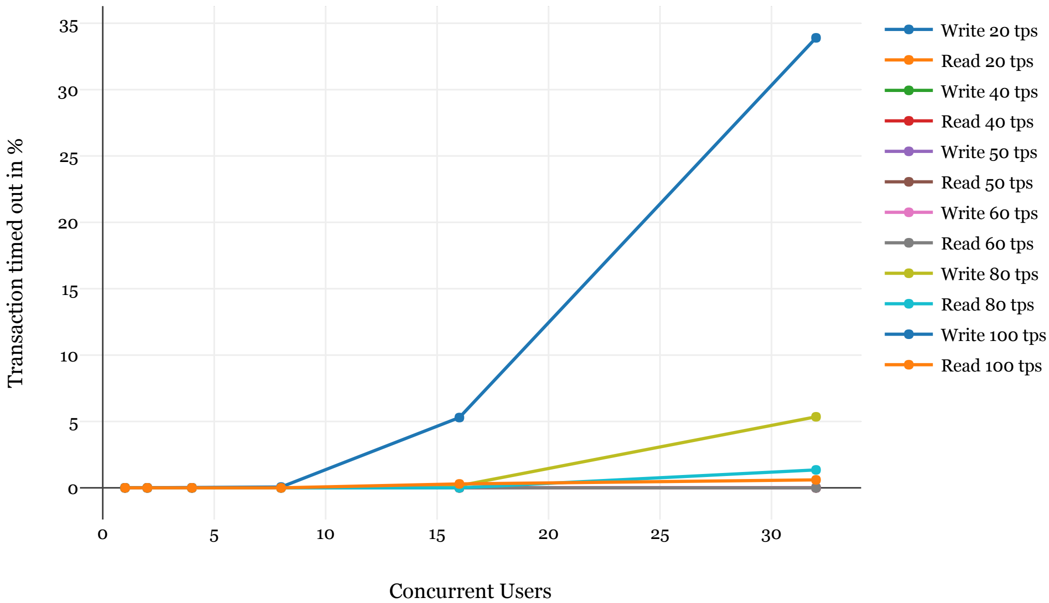

Figure 7.10: Mean Timeouts per Number of Users

Figure 7.11 shows the relation between block size and latency. The results suggest that the average transaction latency decreases significantly until a block size of 32 for all transaction arrival rates. The maximum average latency of 76.95 ms is reached at a send rate of 80 and 1 transaction per block for writing. The minimum average latency of 1.11 ms is reached at a transaction arrival rate of 20 and a block size of 4 for reading.

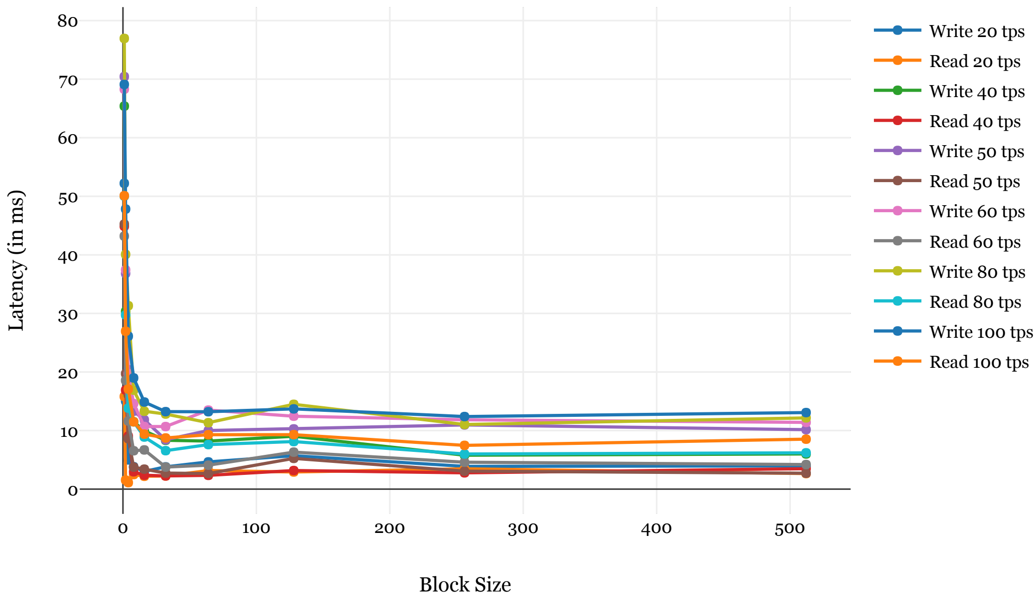

Figure 7.11: Mean Latency per Block Size

7.6 Discussion

Evidence of the evaluation via performance benchmarks proves that the artefact meets the specified requirements. In order to ensure research rigor we compare our results with existing literature. Thakkar et al. (2018) observe that the throughput is highly dependent on the number of CPU cores, which explains the better performance in throughput for distributed setups. This is coherent with our results. Thakkar et al. (2018) suggest that the number of concurrent users correlates negatively with the performance, which could not be confirmed by the conducted benchmarks. In Wickboldt (2019) a setup on a single machine with the Orderer mode Solo was tested. Our results are consistent with this previous work on several points. In Wickboldt (2019) a strong positive correlation (0.53) between the block size and throughput has been observed. Moreover, the benchmarks indicated a strongly negative (-0.84) correlation between the average latency and transaction throughput. These results could be confirmed by the conducted benchmarks within the scope of this paper. In Wickboldt (2019) a stable performance is given to 16 concurrent users in a laboratory environment. These results have been confirmed in a more naturalistic environment. The conducted benchmarks suggest that a higher number of physical machines, organizations and peers have no negative impact on performance. On the contrary, the throughput of distributed systems was significantly higher than single machine environments.

7.7 Conclusions

The outcome of this work includes the advancement of a blockchain-based IT artefact towards a more naturalistic environment and responds to the call for "blockchain applications" in Zheng et al. (2017). We found projectible patterns (Baskerville and Pries-Heje, 2014) that researches can adapt to comparable issues in other sectors. The evaluation proves the efficacy and efficiency of the proof of concept in a naturalistic environment and extends the work of Wickboldt and Kliewer (2018a), Wickboldt and Kliewer (2019) and Wickboldt (2019).

The performance of Hyperledger Fabric was examined using the example of the Blockchain-based Certification Storage System. Hyperledger Caliper was used to run the benchmarks. Various system architecture configurations as well as some crucial parameters were tested. In accordance with the Technical Risk & Efficacy evaluation strategy (Venable et al., 2016), the artefact is first designed in a simple laboratory environment, and then iteratively adds and examines reality requirements. Within this paper, five stages of evaluation were examined. First, a minimal and simplified setup on a single machine was tested. Afterwards the fault tolerant and high-performance Kafka service was introduced. The results indicate a great performance gain when using Kafka compared to the development configuration Solo. Since the underlying business process involves multiple independent stakeholders, a second machine was added to the system in the third evaluation stage. As a result, the transaction throughput increased by 43.60%. In this configuration most components were running on the first machine, which is a *single point of failure* (SPOF). Therefore, a third physical host was introduced in the fourth step. In this setup, there is no SPOF with all components distributed between the machines. There was no measurable performance difference to the prior configuration. In the evaluation stage, the impact of latency was investigated by adding a fourth machine in a different network. The fourth machine had faster hardware, which is why the performance increased slightly.

The requirements (1) and (2) have already been confirmed in prior research and the conducted benchmarks verify the results (Wickboldt, 2019). The requirement (3) was addressed by adding multiple physical machines to the system, which represent the different stakeholder in the aviation industry. Since the performance only increased or stagnated when adding more machines, a higher number of stakeholders is expected to not influence the systems performance negatively. Requirement (4) was addressed by testing the impact of the number of concurrent users on the system. The results suggest a negative impact on the transaction throughput and the relative number of timeouts for high send rates and high number of users. This requirement is fulfilled in consideration of the very low expected throughput and the moderate number of expected users. By showing that the Hyperledger Fabric components are interchangeable between the organizations, the requirement (5) was addressed. A machine with a central CA can be used to represent a technical administrator in the system.

The results indicate a significant increase in transaction throughput, with lesser timeouts using Apache Kafka. The same could be observed by distributing the system over several machines. Due to the overall higher performance of all machines in sum, the latency has also decreased. The network latency did not significantly reduce the performance of the system. The various configurations also found that a block size of 128 provides very good results in terms of latency, throughput and memory usage. Furthermore, the results regarding the influence of the number of users on existing research coincide (Thakkar et al., 2018). Thus, an increasing number of users has a negative effect on the transaction rate. Since the expected transaction arrival rate in a real application is comparably low (2), this should not decrease the performance of the system in a naturalistic environment.

Following the Technical Risk & Efficacy evaluation strategy (Venable et al., 2016), future research should move the system to an even more naturalistic environment and evaluations should be done in a more summative way. This means that further requirements must be set with a larger circle of domain experts. This includes requirements for different user roles and the consensus algorithm between the network participants. The current results indicate that the performance of the proposed system exceeds the requirements and that it can be used as a basis for real-world applications.

References

Albrecht, S., Reichert, S., Schmid, J., Strüker, J., Neumann, D., and Fridgen, G. (2018). "Dynamics of blockchain implementation-a case study from the energy sector". In: *Proceedings of the 51st Hawaii International Conference on System Sciences*. ISBN: 0998133116.

Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., Enyeart, D., Ferris, C., Laventman, G., and Manevich, Y. (2018). "Hyperledger fabric: a distributed operating system for permissioned blockchains". In: *Proceedings of the Thirteenth EuroSys Conference*. ACM, p. 30. ISBN: 1450355846. DOI: 10.1145/3190508.3190538.

Baskerville, R. and Pries-Heje, J. (2014). "Design theory projectability". In: *Working Conference on Information Systems and Organizations*. Springer, pp. 219–232.

Christidis, K. and Devetsikiotis, M. (2016). "Blockchains and Smart Contracts for the Internet of Things". In: *IEEE Access* 4, pp. 2292–2303. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2016.2566339.

Doganis, R. (2013). *Flying off course: The economics of international airlines*. Routledge. ISBN: 1134887787.

Dyer, J. H. and Singh, H. (1998). "The relational view: Cooperative strategy and sources of interorganizational competitive advantage". In: *Academy of management review* 23.4, pp. 660–679. ISSN: 0363-7425.

EASA (2015). *Easa: Acceptable Means of Compliance (Amc) and Guidance Material (Gm)*. Electronic Article. URL: https://www.easa.europa.eu/sites/default/files/dfu/Annex%20I%20to%20Decision%202015029-R%20-%20(AMC-GM%20Part-M).pdf (visited on 04/20/2019).

FAA (2005). *Advisory Circular*. Electronic Article. URL: https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_20-154.pdf (visited on 04/20/2019).

Gregor, S. and Hevner, A. (2013). "Positioning and presenting design science research for maximum impact." In: *MIS quarterly* 37.2.

Hevner, A. and Chatterjee, S. (2010). "Design science research in information systems". In: *Design research in information systems*. Springer, pp. 9–22.

Hevner, A., March, S., Park, J., and Ram, S. (2004). "Design Science in Information Systems Research". In: *MIS quarterly* 28.1, pp. 75–105. DOI: 10.2307/25148625.

IATA (2015). *Best Practices for Component Maintenance Cost Management*. Report. International Air Transport Association. URL: https://www.iata.org/whatwedo/workgroups/Documents/MCTF/CMCM-2nd-edition.pdf (visited on 11/07/2018).

Kumar, M. V. and Iyengar, N. (2017). "A framework for Blockchain technology in rice supply chain management". In: *Adv. Sci. Technol. Lett* 146, pp. 125–130.

Madhwal, Y. and Panfilov, P. B. "BLOCKCHAIN AND SUPPLY CHAIN MANAGEMENT: AIRCRAFTS'PARTS'BUSINESS CASE". In: *28TH DAAAM INTERNATIONAL SYMPOSIUM ON INTELLIGENT MANUFACTURING AND AUTOMATION*. Vol. 28. ISBN: 1726-9679.

March, S. T. and Smith, G. F. (1995). "Design and natural science research on information technology". In: *Decision support systems* 15.4, pp. 251–266. ISSN: 0167-9236.

Nakamoto, S. (2008). "Bitcoin: A Peer-to-Peer Electronic Cash System". In: URL: `https://bitcoin.org/bitcoin.pdf` (visited on 09/14/2017).

Sikorski, J. J., Haughton, J., and Kraft, M. (2017). "Blockchain technology in the chemical industry: Machine-to-machine electricity market". In: *Applied Energy* 195, pp. 234–246. ISSN: 0306-2619.

Simon, H. A. (1996). *The sciences of the artificial.* MIT press. ISBN: 0262264498.

Stahlbock, R., Heilig, L., and Voß, S. (2018). "Blockchain in der maritimen Logistik". In: *HMD Praxis der Wirtschaftsinformatik* 55.6, pp. 1185–1203. ISSN: 1436-3011.

Swan, M. (2015). *Blockchain: blueprint for a new economy.* 1. Beijing u.a.: O'Reilly, XVII, 130 S. ISBN: 978-1-4919-2049-7 1-4919-2049-1. URL: `http://deposit.d-nb.de/cgi-bin/dokserv?id=5125084&prov=M&dok_var=1&dok_ext=htm`.

Thakkar, P., Nathan, S., and Vishwanathan, B. (2018). "Performance Benchmarking and Optimizing Hyperledger Fabric Blockchain Platform". In: *arXiv preprint arXiv:1805.11390*.

Tian, F. (2016). "An agri-food supply chain traceability system for China based on RFID & blockchain technology". In: *13th international conference on service systems and service management (ICSSSM)*. IEEE, pp. 1–6. ISBN: 1509028420.

Venable, J., Pries-Heje, J., and Baskerville, R. (2016). "FEDS: a framework for evaluation in design science research". In: *European Journal of Information Systems* 25.1, pp. 77–89. ISSN: 0960-085X.

Wickboldt, C. (2019). "Benchmarking a Blockchain-based Certification Storage System". In: *Diskussionsbeiträge Fachbereich Wirtschaftswissenschaft.* Vol. 2019/5 : Information Systems. Freie Universität Berlin. URL: `https://refubium.fu-berlin.de/handle/fub188/24484?locale-attribute=en`.

Wickboldt, C. and Kliewer, N. (2018a). "Blockchain zur dezentralen Dokumentation von Werkstattereignissen in der Luftfahrtindustrie". In: *HMD Praxis der Wirtschaftsinformatik*, pp. 1–14. ISSN: 1436-3011. URL: `https://doi.org/10.1365/s40702-018-00452-y`.

— (2019). "Blockchain for Workshop Event Certificates - A Proof of Concept in the Aviation Industry". In: *Proceedings of the 27th European Conference on Information Systems (ECIS)*. ISBN: 978-1-7336325-0-8. URL: `https://aisel.aisnet.org/ecis2019_rp/35/`.

Zheng, Z., Xie, S., Dai, H., Chen, X., and Wang, H. (2017). "An overview of blockchain technology: Architecture, consensus, and future trends". In: *Big Data*

(BigData Congress), 2017 IEEE International Congress on. IEEE, pp. 557–564. ISBN: 1538619962. DOI: `10.1109/BigDataCongress.2017.85`.

Chapter 8

Conclusion

The scope of this thesis lied on the evaluation of methods from data analytics and decentralized ledger technologies by designing IT artefacts and applying them to address problems in existing business processes in the aviation industry. Six publications document the research process and give insights in evaluating methods to find a FMV for spare parts and also preserve this value by ensuring complete lifecycle documentation. Chapter 2 and 3 address problem P1 (A valuation of spare parts is only possible with great effort due to incomplete information about these parts.). Findings of chapter 2 and 3 had been communicated (Peffers et al., 2007) on scientific conferences. Chapter 4, 5, 6 and 7 address problem P2 (The history of workshop and trading events for safety-related spare parts cannot be documented transparently and persistently.). Communcation of the findings had been done to business and scientific audiences.

8.1 Summary of Findings

Chapter 2 embeds an automated valuation in the pricing literature and suggests a system architecture for a Fair Market Evaluator which could be used by end users in the MRO industry. Findings of chapter 2 include that cost-based and competition-based methods are not but value-based methods are feasible to enable an automated valuation of spare parts and that the Fair Market Evaluator adresses shortcomings in manual valuation. Automated valuation sets a basis for part related decision making and by that speeds up global supply chain which could lead to a competitive advantage (Bharadwaj et al., 2013).

Chapter 3 goes deeper into the individual methods within the ETL process and gives recommendations for the classification of spare parts for subsequent valuation. A spare part's lifecycle is aligned to new parts. When the distribution rate of new parts drops, the aircraft types reached a certain age that there is greater demand for used parts. Fewer new parts come onto the market, more parts have to be overhauled and reused. This increases the demand for used spare parts, which in turn influences their value. An IT artefact in the sense of Hevner et al. (2004) and Peffers et al. (2007) is introduced which combines methods from ETL (Kimball and

Ross, 2011), handling missing data (Allison, 2002) and suitable methods for value estimation (Voß and Lessmann, 2013). The ASPV framework gives insight about how to process data through an ETL process as a basis for automated valuation. Key findings include that deduplication of multiple price points per part and per date should be performed by selecting datasets according to business rules. It also turns out that missing data is one of the key obstacles when applying forecasting methods to estimate a FMV. Demonstration of the artefact has been performed on scientific conferences, especially to ensure that the artefact solves the problem of handling missing data. Unfortunately and due to confidentiality of business critical data, publication of numerical results has not been possible until now. A core insight after working with the data on client site of understanding the greatest impact on a spare parts value is that value vastly depends on condition and part history of the particular part.

By analyzing the data and talking to domain experts, it became clear that the value of a spare part depended significantly on its condition and workshop history. This lead to the urge to define objectives for the second artefact which solves the problem of incomplete workshop event documentation and by that ensures the value retention of the spare part.

Chapter 4 introduces a conceptual model for a blockchain-based artefact which digitizes the analogue documentation process and by that solves the problem of missing workshop event certificates. This contributes to the call for "blockchain applications" in Zheng et al. (2017) and the linking of new technologies with a use case in the sense of Glaser (2017). Key findings of chapter 4 include that a permissioned blockchain based on Hyperledger Fabric (Group, 2018) fits the requirements to the solution which are to ensure fast, orderly, persistent, forgery-proof, access-restricted documentation solution that is aligned with the rules of the business process. This also overcomes trust issues when sharing proprietary information (Dyer and Singh, 1998) between supply chain partners. It allows information flows (Klein and Rai, 2009) in order to realize reduced costs of operations, enhanced productivity and improved management of assets.

A proof of concept is introduced as an IT artefact in the sense of Hevner et al. (2004) and Peffers et al. (2007) in chapter 5. The underlying business process, an Ideal Component Repair Cycle, had been analyzed deeply. At every stage in the process a documentation certificate is generated. Knowing this, a datamodel has been created which adresses the business process. The BCSS was introduced, consisting of an Angular Frontend, a Java Backend and a Hyperledger Fabric Blockchain. A simple evaluation showed that the required performance of 0.16 tps can be met and surpassed. This demonstrates that the artefact is able to solve instances of the problem in the sense of Peffers et al. (2007).

Further evaluation (Peffers et al., 2007) by investigating the system's performance to ensure efficacy of the artefact had been done in chapter 6. A sophisticated performance benchmark regarding block size, user count and transaction arrival rate had been done in order to identify the capabilities of the host system and find the right parameters for the Hyperledger Fabric blockchain. It has been shown that already at a dual core host system and with one orderer node, 16 concurrent users with a transaction arrival rate of 60 tps can be handled. A block size of 128 transactions per block has been recommended. With that configuration a throughput is reached which is in line with similar blockchain systems (Thakkar et al., 2018). It also shows that the performance of the BCSS system exceeds the requirements by large, ensuring scalability and future proofing of the solution.

In chapter 7, the system has been further decentralized to perform close to a naturalistic environment. Additional requirements (independent physical machines, stable performance while increasing the number of users, a technical role which is able to administrate the decentralized system) were met. Hyperledger Fabric instances (Orgs) had been distributed over four different machines in two networks. Communication between the Orgs was organized in a Docker Swarm Network. The system was benchmarked again. In comparison the the setup in chapter 6, transaction throughput had been increased while transaction timeouts were decreased. This proofs scalability of the system and pushed the proof of concept further toward a naturalistic environment.

8.2 Future Research

As projectible patterns of the business process of an Alternative Closed Loop Supply Chain had been communicated in chapter 3, the ASPV should be extended by an application on real world data. Prediction models on pricing in secondary markets had been successfully testet by Voß and Lessmann (2013). Spare part classification can be applied to any process which is able to differentiate between vital few and the rest. This classification is used to limit the scope for preparing the data for estimating a FMV. Also, the understanding of the product lifecycle of spare parts helps to enrich the prediction model with further information. The suitability of the application of the presented methods can, as in this case, be done together with domain experts. Even better, and above all more beneficial for the general knowledge base, would be to work with data whose publication does not put the owner at a competitive disadvantage.

The continuation of research at BCSS is much more straightforward due to the lower dependency on data. Following the Technical Risk & Efficacy evaluation strategy, Venable et al. (2016) suggest incrementally taking the IT artefact closer to a naturalistic environment. This includes further scaling of the system and decentralizing it through multiple networks. As a proof of concept has been introduced in chapter 5, it was evaluated deeply in chapter 6 and extended in chapter 7. Next steps include taking the proof of concept to selected users to test it outside the laboratory. User roles within the system should be further investigated with domain experts to increase user acceptance.

In the end, data correctness and data quality depends on user inputs. Therefore, even when the data is digitized, there is uncertainty about the correctness of the output data such as estimated FMVs or maintenance event certificates. Internet of Things will play a larger role in order to avoid such data entry risks. Ideally, the data is not manually entered by the user but automatically sent from the device itself.

Also, as of today, (confidential) data is provided to data silos of analytic companies to be used as training data for prediction algorithms. Projects like Ocean Protocol[1] aim for a new way of data ownership by introducing a decentralized data market place while data ownership remains at the original place. This allows algorithms and prediction models to access shared data without exposing the data into silos. Such approaches might overcome limitations of evaluating IT artefacts using confidential data.

[1] `https://docs.oceanprotocol.com`

References

Allison, P. D. (2002). "Missing data. Series: Quantitative applications in the social sciences". In: *Thousand Oaks. CA*.

Bharadwaj, A., El Sawy, O., Pavlou, P., and Venkatraman, N. (2013). "Digital business strategy: toward a next generation of insights". In: *MIS quarterly*, pp. 471–482. ISSN: 0276-7783.

Dyer, J. H. and Singh, H. (1998). "The relational view: Cooperative strategy and sources of interorganizational competitive advantage". In: *Academy of management review* 23.4, pp. 660–679. ISSN: 0363-7425.

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 2017*. ISBN: 978-0-9981331-0-2. URL: `http://hdl.handle.net/10125/41339`.

Group, H. A. W. (2018). *Hyperledger Architecture*. Electronic Article. URL: `https://www.hyperledger.org/wp-content/uploads/2018/04/Hyperledger_Arch_WG_Paper_2_SmartContracts.pdf` (visited on 06/29/2018).

Hevner, A., March, S., Park, J., and Ram, S. (2004). "Design Science in Information Systems Research". In: *MIS quarterly* 28.1, pp. 75–105. DOI: `10.2307/25148625`.

Kimball, R. and Ross, M. (2011). *The data warehouse toolkit: the complete guide to dimensional modeling*. John Wiley & Sons. ISBN: 1118082141.

Klein, R. and Rai, A. (2009). "Interfirm strategic information flows in logistics supply chain relationships". In: *Mis quarterly*, pp. 735–762. ISSN: 0276-7783.

Peffers, K., Tuunanen, T., Rothenberger, M. A., and Chatterjee, S. (2007). "A design science research methodology for information systems research". In: *Journal of management information systems* 24.3, pp. 45–77.

Thakkar, P., Nathan, S., and Vishwanathan, B. (2018). "Performance Benchmarking and Optimizing Hyperledger Fabric Blockchain Platform". In: *arXiv preprint arXiv:1805.11390*.

Venable, J., Pries-Heje, J., and Baskerville, R. (2016). "FEDS: a framework for evaluation in design science research". In: *European Journal of Information Systems* 25.1, pp. 77–89. ISSN: 0960-085X.

Voß, S. and Lessmann, S. (2013). "Resale Price Prediction in the Used Car Market". In: *Tristan Symposium VIII*.

Zheng, Z., Xie, S., Dai, H., Chen, X., and Wang, H. (2017). "An overview of blockchain technology: Architecture, consensus, and future trends". In: *Big Data (BigData Congress), 2017 IEEE International Congress on*. IEEE, pp. 557–564. ISBN: 1538619962. DOI: `10.1109/BigDataCongress.2017.85`.

Global References

Abe, N. and Kamba, T. (2000). "A Web marketing system with automatic pricing". In: *Computer Networks-the International Journal of Computer and Telecommunications Networking* 33.1-6, pp. 775–788. ISSN: 1389-1286.

Abeyratne, S. and Monfared, R. (2016). "Blockchain ready manufacturing supply chain using distributed ledger". In: *International Journal of Research in Engineering and Technology* 5.09, pp. 1–10. DOI: 10.15623/ijret.2016.0509001.

Accenture (2018). *BEYOND THE BUZZ: THE POTENTIAL OF BLOCKCHAIN TECHNOLOGY FOR AIRLINES*. Electronic Article. URL: https://www.accenture.com/us-en/insight-potential-blockchain-technology-for-airlines (visited on 11/27/2018).

Ackert, S. (2012). "Basics of Aircraft Market Analysis". In: *Aircraft Monitor*.

Albrecht, S., Reichert, S., Schmid, J., Strüker, J., Neumann, D., and Fridgen, G. (2018). "Dynamics of blockchain implementation-a case study from the energy sector". In: *Proceedings of the 51st Hawaii International Conference on System Sciences*. ISBN: 0998133116.

Allison, P. D. (2002). "Missing data. Series: Quantitative applications in the social sciences". In: *Thousand Oaks. CA*.

Anderson, J., Jain, D., and Chintagunta, P. (1992). "Customer value assessment in business markets: A state-of-practice study". In: *Journal of Business-to-Business Marketing* 1, pp. 3–29.

Anderson, J. and Narus, J. (1998). "Business marketing: understand what customers value". In: *Harvard business review* 76, pp. 53–67.

Anderson, J., Thomson, J. B., and Wynstra, F. (2000). "Combining value and price to make purchase decisions in business markets". In: *International Journal of Research in Marketing* 17.4, pp. 307–329. ISSN: 0167-8116.

Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., Enyeart, D., Ferris, C., Laventman, G., and Manevich, Y. (2018). "Hyperledger fabric: a distributed operating system for permissioned blockchains". In: *Proceedings of the Thirteenth EuroSys Conference*. ACM, p. 30. ISBN: 1450355846. DOI: 10.1145/3190508.3190538.

Arlitt, M., Krishnamurthy, D., and Rolia, J. (2001). "Characterizing the scalability of a large web-based shopping system". In: *ACM Transactions on Internet Technology* 1.1, pp. 44–69. ISSN: 1533-5399.

Armstrong, S. (2016). *Eight things you need to know about the future of retail.* Electronic Article. URL: https://www.wired.co.uk/article/wired-retail-2015 (visited on 03/25/2019).

Avritzer, A., Kondek, J., Liu, D., and Weyuker, E. J. (2002). "Software performance testing based on workload characterization". In: *Proceedings of the 3rd international workshop on Software and performance.* ACM, pp. 17–24.

Barber, S. (2004). "Creating effective load models for performance testing with incomplete empirical data". In: *Web Site Evolution, Sixth IEEE International Workshop on (WSE'04).* IEEE, pp. 51–59.

Baskerville, R. and Pries-Heje, J. (2014). "Design theory projectability". In: *Working Conference on Information Systems and Organizations.* Springer, pp. 219–232.

Beck, R., Czepluch, J. S., Lollike, N., and Malone, S. (2016). "Blockchain-the Gateway to Trust-Free Cryptographic Transactions". In: *ECIS*, ResearchPaper153.

Bellamy, W. (2017). *Air France KLM is Evaluating MRO Potential for Blockchain.* Electronic Article. URL: https://www.aviationtoday.com/2017/10/03/air-france-klm-evaluating-mro-potential-blockchain/ (visited on 11/27/2018).

Ben-Daya, M., Ait-Kadi, D., Duffuaa, S. O., Knezevic, J., and Raouf, A. (2009). *Handbook of maintenance management and engineering.* Vol. 7. Springer.

Bharadwaj, A., El Sawy, O., Pavlou, P., and Venkatraman, N. (2013). "Digital business strategy: toward a next generation of insights". In: *MIS quarterly*, pp. 471–482. ISSN: 0276-7783.

Blumberg, D. F. (2004). *Introduction to Management of Reverse Logistics and Closed Loop Supply Chain Processes.* CRC Press. ISBN: 9780203502938. DOI: 10.1201/b12427.

Blythe, J. (2008). *Essentials of marketing.* Pearson Education.

Buterin, V. (2014). Web Page. DOI: citeulike-article-id:14229692. URL: https://github.com/ethereum/wiki/wiki/White-Paper (visited on 04/17/2018).

Cachin, C. (2016). "Architecture of the Hyperledger blockchain fabric". In: *Workshop on Distributed Cryptocurrencies and Consensus Ledgers.*

Caliper, H. (2019). *Hyperledger Caliper Architecture.* Electronic Article. URL: https://hyperledger.github.io/caliper/docs/2_Architecture.html (visited on 03/10/2019).

Chaudhuri, S., Dayal, U., and Narasayya, V. (2011). "An overview of business intelligence technology". In: *Communications of the ACM* 54.8, pp. 88–98. ISSN: 0001-0782.

Chen, H., Chiang, R. H., and Storey, V. C. (2012). "Business intelligence and analytics: From big data to big impact". In: *MIS quarterly* 36.4. ISSN: 0276-7783.

Christidis, K. and Devetsikiotis, M. (2016). "Blockchains and Smart Contracts for the Internet of Things". In: *IEEE Access* 4, pp. 2292–2303. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2016.2566339.

Davenport, T. H. (2006). “Competing on analytics”. In: *harvard business review* 84.1, p. 98. ISSN: 0017-8012.

Deloitte (2019). *Tech Trends 2019: Beyond the digital frontier*. Electronic Article. URL: https://www2.deloitte.com/content/dam/Deloitte/br/Documents/technology/DI_TechTrends2019.pdf (visited on 03/15/2019).

Doganis, R. (2013). *Flying off course: The economics of international airlines*. Routledge. ISBN: 1134887787.

Dolev, D. and Strong, H. R. (1983). “Authenticated algorithms for Byzantine agreement”. In: *SIAM Journal on Computing* 12.4, pp. 656–666. ISSN: 0097-5397.

Dyer, J. H. and Singh, H. (1998). “The relational view: Cooperative strategy and sources of interorganizational competitive advantage”. In: *Academy of management review* 23.4, pp. 660–679. ISSN: 0363-7425.

EASA (2015). *Easa: Acceptable Means of Compliance (Amc) and Guidance Material (Gm)*. Electronic Article. URL: https://www.easa.europa.eu/sites/default/files/dfu/Annex%20I%20to%20Decision%202015029-R%20-%20(AMC-GM%20Part-M).pdf (visited on 04/20/2019).

Eaves, A. H. C. and Kingsman, B. G. (2004). “Forecasting for the ordering and stock-holding of spare parts”. In: *Journal of the Operational Research Society* 55.4, pp. 431–437. DOI: 10.1057/palgrave.jors.2601697.

Economist, T. (2015). *The promise of the blockchain - The trust machine*. Electronic Article. URL: https://www.economist.com/leaders/2015/10/31/the-trust-machine (visited on 11/27/2018).

Ethereum (2017). *Consortium Chain Development*. Electronic Article. URL: https://github.com/ethereum/wiki/wiki/Consortium-Chain-Development (visited on 06/29/2018).

Evers, M. (2018). *Generating more transparency in aviation with blockchain technology*. Electronic Article. URL: https://www.lufthansa-industry-solutions.com/de-en/solutions-products/aviation/generating-more-transparency-in-aviation-with-blockchain-technology/ (visited on 11/27/2018).

FAA (2005). *Advisory Circular*. Electronic Article. URL: https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_20-154.pdf (visited on 04/20/2019).

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 2017*. ISBN: 978-0-9981331-0-2. URL: http://hdl.handle.net/10125/41339.

Greenspan, G. (2015). *MultiChain private blockchain—White paper*. Electronic Article. URL: https://www.multichain.com/download/MultiChain-White-Paper.pdf (visited on 11/27/2018).

Gregor, S. and Hevner, A. (2013). “Positioning and presenting design science research for maximum impact.” In: *MIS quarterly* 37.2.

Grosfeld-Nir, A., Ronen, B., and Kozlovsky, N. (2007). "The Pareto managerial principle: when does it apply?" In: *International Journal of Production Research* 45.10, pp. 2317–2325.

Group, H. A. W. (2018). *Hyperledger Architecture*. Electronic Article. URL: https://www.hyperledger.org/wp-content/uploads/2018/04/Hyperledger_Arch_WG_Paper_2_SmartContracts.pdf (visited on 06/29/2018).

Hastie, T., Tibshirani, R., and Friedman, J. (2009). "The Elements of Statistical Learning: Data Mining, Inference, and Prediction". In: *Springer Series in Statistics*.

Hevner, A. and Chatterjee, S. (2010). "Design science research in information systems". In: *Design research in information systems*. Springer, pp. 9–22.

Hevner, A., March, S., Park, J., and Ram, S. (2004). "Design Science in Information Systems Research". In: *MIS quarterly* 28.1, pp. 75–105. DOI: 10.2307/25148625.

Hinterhuber, A. (2004). "Towards value-based pricing—An integrative framework for decision making". In: *Industrial Marketing Management* 33.8, pp. 765–778. ISSN: 00198501. DOI: 10.1016/j.indmarman.2003.10.006.

— (2008). "Customer value-based pricing strategies: why companies resist". In: *Journal of Business Strategy* 29.4, pp. 41–50. ISSN: 0275-6668. DOI: 10.1108/02756660810887079.

Hyperledger (2018). *Hyperledger Fabric Documentation*. Electronic Article. URL: https://hyperledger-fabric.readthedocs.io/ (visited on 11/27/2018).

IATA (2009). *IATA Technical Reference Manual*. Report. International Air Transport Association. URL: http://www.aerohabitat.eu/uploads/media/29-07-2010_-_IATA_-_Aviation_Technical_Reference_Manual.pdf (visited on 11/07/2018).

— (2015). *Best Practices for Component Maintenance Cost Management*. Report. International Air Transport Association. URL: https://www.iata.org/whatwedo/workgroups/Documents/MCTF/CMCM-2nd-edition.pdf (visited on 11/07/2018).

— (2017). *IATA Annual Review 2017*. Electronic Article. URL: https://www.iata.org/publications/Documents/iata-annual-review-2017.pdf (visited on 03/20/2019).

— (2018a). *AIRLINE MAINTENANCE COST EXECUTIVE COMMENTARY: An Exclusive Benchmark Analysis (FY2017 data) by IATA's Maintenance Cost Task Force*. Electronic Article. URL: https://www.iata.org/whatwedo/workgroups/Documents/MCTF/MCTF-FY2017-Report-Public.pdf (visited on 03/15/2019).

— (2018b). *Blockchain in Aviation: Exploring the Fundamentals, Use Cases, And Industry Initiatives*. Report. International Air Transport Association. URL: https://www.iata.org/publications/Documents/blockchain-in-aviation-white-paper.pdf (visited on 11/07/2018).

Johansson, M., Keränen, J., Hinterhuber, A., Liozu, S., and Andersson, L. (2015). "Value assessment and pricing capabilities? How to profit from value". In: *Journal of Revenue & Pricing Management*.

Kelly, D. B. (2008). "Forecasting aircraft values: An appraiser's perspective". In: *Air-finance annual*.

Kim, H. M. and Laskowski, M. (2018). "Toward an ontology-driven blockchain design for supply-chain provenance". In: *Intelligent Systems in Accounting, Finance and Management* 25.1, pp. 18–27. ISSN: 1055-615X.

Kimball, R. and Ross, M. (2011). *The data warehouse toolkit: the complete guide to dimensional modeling*. John Wiley & Sons. ISBN: 1118082141.

Klein, R. and Rai, A. (2009). "Interfirm strategic information flows in logistics supply chain relationships". In: *Mis quarterly*, pp. 735–762. ISSN: 0276-7783.

Kumar, M. V. and Iyengar, N. (2017). "A framework for Blockchain technology in rice supply chain management". In: *Adv. Sci. Technol. Lett* 146, pp. 125–130.

Love, C. and Guo, R. (1996). "Utilizing Weibull failure rates in repair limit analysis for equipment replacement/preventive maintenance decisions". In: *Journal of the operational Research Society* 47.11, pp. 1366–1376. ISSN: 0160-5682.

Madhwal, Y. and Panfilov, P. B. "BLOCKCHAIN AND SUPPLY CHAIN MANAGEMENT: AIRCRAFTS'PARTS'BUSINESS CASE". In: *28TH DAAAM INTERNATIONAL SYMPOSIUM ON INTELLIGENT MANUFACTURING AND AUTOMATION*. Vol. 28. ISBN: 1726-9679.

March, S. T. and Smith, G. F. (1995). "Design and natural science research on information technology". In: *Decision support systems* 15.4, pp. 251–266. ISSN: 0167-9236.

Markus, M. L. and Rowe, F. (2018). "Is IT Changing the World? Conceptions of Causality for Information Systems Theorizing". In: *Management Information Systems Quarterly* 42.4, pp. 1255–1280.

Meise, C., Wickboldt, C., and Kliewer, N. (2019). "Decentralized Workshop Event Documentation with Hyperledger Fabric (submitted)". In: *Fortieth International Conference on Information Systems*.

Mercer, Z. C. and Brown, T. S. (1999). "Fair Market Value vs. The Real World". In: *Business Valuation Review* 18.1, pp. 16–25.

Merriam-Webster (2019). *Provenance*. Electronic Article. URL: `https://www.merriam-webster.com/dictionary/provenance` (visited on 03/25/2019).

MRO-Network (2015). *Teardown Numbers Attract New Entrants*. Electronic Article. URL: `https://www.mro-network.com/maintenance-repair-overhaul/teardown-numbers-attract-new-entrants` (visited on 03/20/2019).

Naisbitt, J. and Cracknell, J. (1984). *Megatrends: Ten new directions transforming our lives*. Report. Warner Books New York.

Nakamoto, S. (2008). "Bitcoin: A Peer-to-Peer Electronic Cash System". In: URL: `https://bitcoin.org/bitcoin.pdf` (visited on 09/14/2017).

Neely, A. and Jarrar, Y. (2004). “Extracting value from data–the performance planning value chain”. In: *Business Process Management Journal* 10.5, pp. 506–509.

Nofer, M., Gomber, P., Hinz, O., and Schiereck, D. (2017). “Blockchain”. In: *Business & Information Systems Engineering* 59.3, pp. 183–187. ISSN: 2363-7005. DOI: 10.1007/s12599-017-0467-3.

Notheisen, B., Cholewa, J., and Shanmugam, A. (2017). “Trading Real-World Assets on Blockchain”. In: *Business & Information Systems Engineering* 59.6, pp. 425–440. ISSN: 2363-7005.

Notheisen, B., Hawlitschek, F., and Weinhardt, C. (2017). “Breaking down the blockchain hype–towards a blockchain market engineering approach”. In: *European Conference on Information Systems.*

Nowlan, F. S. and Heap, H. F. (1978). *Reliability-centered maintenance.* Report. United Air Lines Inc San Francisco Ca.

Östlin, J., Sundin, E., and Björkman, M. (2009). “Product life-cycle implications for remanufacturing strategies”. In: *Journal of cleaner production* 17.11, pp. 999–1009.

Peffers, K., Tuunanen, T., Rothenberger, M. A., and Chatterjee, S. (2007). “A design science research methodology for information systems research”. In: *Journal of management information systems* 24.3, pp. 45–77.

Provenance (2015). *Blockchain: the solution for transparency in product supply chains.* Electronic Article. URL: https://www.provenance.org/whitepaper (visited on 11/07/2018).

Quantas (2016). *The A, C and D of aircraft maintenance.* Electronic Article. URL: https://www.qantasnewsroom.com.au/roo-tales/the-a-c-and-d-of-aircraft-maintenance/ (visited on 03/10/2019).

Renesse, R. van and Guerraoui, R. (2010). “Replication techniques for availability”. In: *Replication: Theory and Practice.* Ed. by B. Charron-Bost, F. Pedone, and A. Schiper. Replication: Theory and Practice. Springer-Verlag Berlin Heidelberg. Chap. 2, pp. 19–40. ISBN: 978-3-642-11294-2. DOI: 10.1007/978-3-642-11294-2.

Ronen, B., Spiegler, I., and Kozlovsky, N. (2005). “The use, misuse and abuse of the Pareto principle”. In: *Working paper.*

Schneider, F. B. (1990). “Implementing fault-tolerant services using the state machine approach: A tutorial”. In: *ACM Computing Surveys (CSUR)* 22.4, pp. 299–319. ISSN: 0360-0300. DOI: 10.1145/98163.98167.

Sikorski, J. J., Haughton, J., and Kraft, M. (2017). “Blockchain technology in the chemical industry: Machine-to-machine electricity market”. In: *Applied Energy* 195, pp. 234–246. ISSN: 0306-2619.

Simon, H. A. (1996). *The sciences of the artificial.* MIT press. ISBN: 0262264498.

Smith, C. (2002). “Software performance engineering”. In: *Encyclopedia of Software Engineering.*

Stahlbock, R., Heilig, L., and Voß, S. (2018). "Blockchain in der maritimen Logistik". In: *HMD Praxis der Wirtschaftsinformatik* 55.6, pp. 1185–1203. ISSN: 1436-3011.

Swan, M. (2015). *Blockchain: blueprint for a new economy.* 1. Beijing u.a.: O'Reilly, XVII, 130 S. ISBN: 978-1-4919-2049-7 1-4919-2049-1. URL: http://deposit.d-nb.de/cgi-bin/dokserv?id=5125084&prov=M&dok_var=1&dok_ext=htm.

Szabo, N. (1997). *Formalizing and Securing Relationships on Public Networks.* Electronic Article. URL: https://ojphi.org/ojs/index.php/fm/article/view/548/469 (visited on 06/29/2018).

Thakkar, P., Nathan, S., and Vishwanathan, B. (2018). "Performance Benchmarking and Optimizing Hyperledger Fabric Blockchain Platform". In: *arXiv preprint arXiv:1805.11390.*

Tian, F. (2016). "An agri-food supply chain traceability system for China based on RFID & blockchain technology". In: *13th international conference on service systems and service management (ICSSSM).* IEEE, pp. 1–6. ISBN: 1509028420.

Tilson, D., Lyytinen, K., and Sørensen, C. (2010). "Research commentary—Digital infrastructures: The missing IS research agenda". In: *Information systems research* 21.4, pp. 748–759. ISSN: 1047-7047.

Venable, J., Pries-Heje, J., and Baskerville, R. (2016). "FEDS: a framework for evaluation in design science research". In: *European Journal of Information Systems* 25.1, pp. 77–89. ISSN: 0960-085X.

Venkatraman, N. (1994). "IT-enabled business transformation: from automation to business scope redefinition". In: *Sloan management review* 35.2, pp. 73–88. ISSN: 0019-848X.

Voß, S. and Lessmann, S. (2013). "Resale Price Prediction in the Used Car Market". In: *Tristan Symposium VIII.*

Wickboldt, C. (2019). "Benchmarking a Blockchain-based Certification Storage System". In: *Diskussionsbeiträge Fachbereich Wirtschaftswissenschaft.* Vol. 2019/5 : Information Systems. Freie Universität Berlin. URL: https://refubium.fu-berlin.de/handle/fub188/24484?locale-attribute=en.

Wickboldt, C. and Kliewer, N. "A framework for value-based pricing in aircraft surplus markets". In: *Multikonferenz Wirtschaftsinformatik (MKWI) 2016.* ISBN: 978-3-86360-132-4. URL: http://www.mkwi2016.de/download/MKWI2016_Research-in-Progress-Poster-Band.pdf.

— (2018a). "Blockchain zur dezentralen Dokumentation von Werkstattereignissen in der Luftfahrtindustrie". In: *HMD Praxis der Wirtschaftsinformatik*, pp. 1–14. ISSN: 1436-3011. URL: https://doi.org/10.1365/s40702-018-00452-y.

— (2018b). "Value Based Pricing meets Data Science: A Concept for Automated Spare Part Valuation". In: *Multikonferenz Wirtschaftsinformatik (MKWI) 2018.* ISBN: 978-3-935786-72-0. URL: http://mkwi2018.leuphana.de/wp-content/uploads/MKWI2018_Band1.pdf.

Wickboldt, C. and Kliewer, N. (2019). "Blockchain for Workshop Event Certificates - A Proof of Concept in the Aviation Industry". In: *Proceedings of the 27th European Conference on Information Systems (ECIS)*. ISBN: 978-1-7336325-0-8. URL: https://aisel.aisnet.org/ecis2019_rp/35/.

Wiliam, D. and Black, P. (1996). "Meanings and consequences: a basis for distinguishing formative and summative functions of assessment?" In: *British Educational Research Journal* 22.5, pp. 537–548. ISSN: 0141-1926.

Woodside, M., Franks, G., and Petriu, D. C. (2007). "The future of software performance engineering". In: *2007 Future of Software Engineering*. IEEE Computer Society, pp. 171–187.

Wyman, O. (2018a). *GLOBAL FLEET & MRO MARKET FORECAST COMMENTARY*. Electronic Article. URL: https://www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2018/January/2018-2028_Global_Fleet_MRO_Market_Forecast_Commentary_Public_Final_web.pdf (visited on 03/15/2019).

— (2018b). *MRO SURVEY 2018: TACKLING INDUSTRY DISRUPTION*. Electronic Article. URL: https://www.oliverwyman.com/content/dam/oliver-wyman/v2/publications/2018/april/MRO-Survey-2018-web.pdf (visited on 03/15/2019).

Xu, X., Pautasso, C., Zhu, L., Gramoli, V., Ponomarev, A., Tran, A. B., and Chen, S. (2016). "The blockchain as a software connector". In: *13th Working IEEE/IFIP Conference on Software Architecture (WICSA)*. IEEE, pp. 182–191. ISBN: 1509021310.

Xu, X., Weber, I., Staples, M., Zhu, L., Bosch, J., Bass, L., Pautasso, C., and Rimba, P. (2017). "A taxonomy of blockchain-based systems for architecture design". In: *IEEE International Conference on Software Architecture (ICSA)*. IEEE, pp. 243–252. ISBN: 1509057293.

Zhang, Y. and Wen, J. (2017). "The IoT electric business model: Using blockchain technology for the internet of things". In: *Peer-to-Peer Networking and Applications* 10.4, pp. 983–994. ISSN: 1936-6442.

Zheng, Z., Xie, S., Dai, H., Chen, X., and Wang, H. (2017). "An overview of blockchain technology: Architecture, consensus, and future trends". In: *Big Data (BigData Congress), 2017 IEEE International Congress on*. IEEE, pp. 557–564. ISBN: 1538619962. DOI: 10.1109/BigDataCongress.2017.85.