

RESEARCH ARTICLE

Adopting microservices and DevOps in the cyber-physical systems domain: A rapid review and case study

Jonas Fritzschi¹ | Justus Bogner¹ | Markus Haug¹ | Ana Cristina Franco da Silva¹ | Carolin Rubner² | Matthias Saft² | Horst Sauer² | Stefan Wagner¹

¹Institute of Software Engineering, University of Stuttgart, Stuttgart, Germany

²Siemens AG, Siemens Technology, Munich, Germany

Correspondence

Jonas Fritzschi, Institute of Software Engineering, University of Stuttgart, Universitätsstraße 38, 70569 Stuttgart, Germany.
Email: jonas.fritzschi@iste.uni-stuttgart.de

Funding information

Siemens AG, Germany; University of Stuttgart, Germany; German Federal Ministry of Education and Research, Grant/Award Numbers: 01IS22087R, 01IS17051; Baden-Württemberg Ministry of Science, Research and the Arts; Innovation Campus Future Mobility; Artificial Intelligence Software Academy

Abstract

The domain of cyber-physical systems (CPS) has recently seen strong growth, for example, due to the rise of the Internet of Things (IoT) in industrial domains, commonly referred to as “Industry 4.0.” However, CPS challenges like the strong hardware focus can impact modern software development practices, especially in the context of modernizing legacy systems. While microservices and DevOps have been widely studied for enterprise applications, there is insufficient coverage for the CPS domain. Our goal is therefore to analyze the peculiarities of such systems regarding challenges and practices for using and migrating towards microservices and DevOps. We conducted a rapid review based on 146 scientific papers, and subsequently validated our findings in an interview-based case study with nine CPS professionals in different business units at Siemens AG. The combined results picture the specifics of microservices and DevOps in the CPS domain. While several differences were revealed that may require adapted methods, many challenges and practices are shared with typical enterprise applications. Our study supports CPS researchers and practitioners with a summary of challenges, practices to address them, and research opportunities.

KEYWORDS

case study, cyber-physical systems, DevOps, interviews, microservices, rapid review

1 | INTRODUCTION

The ongoing digital transformation of the industrial sector with a focus on Internet of Things (IoT) technologies is also referred to as “Industry 4.0.”¹ In this context, industrial software engineering (SE), that is, the systematic development of high-quality cyber-physical systems (CPS), becomes more important than ever.² It is characterized by the digital representation and management of all involved industrial processes, with the goal to increase flexibility and evolvability.³ From a SE perspective, microservices and DevOps have similar objectives regarding these qualities.⁴ Even though both originated in the context of enterprise applications and are seen as the de facto standard for modern cloud-based software systems, they are also regarded as promising approaches for the IoT and CPS domain.⁵ Within agile development,

Ana Cristina Franco da Silva is an independent researcher and participated in this research while at the Institute of Software Engineering, University of Stuttgart, Stuttgart, Germany.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Software: Practice and Experience* published by John Wiley & Sons Ltd.

combining microservices and DevOps promises higher flexibility, shorter release cycles, and better maintainability, which are highly attractive benefits for CPS manufacturers.

Fundamental changes to the software architecture are usually accompanied by operational and organizational changes. Hence, transforming a legacy system that is developed using a traditional process towards DevOps and microservices is a challenging task,⁶ which needs to be carefully considered.⁷ Various empirical studies cover challenges and practices for microservices and DevOps in the context of enterprise applications.^{8,9} However, CPS and the IoT domain differ from common enterprise systems in several aspects:

- Strong focus on hardware with a virtual representation (digital twin)²
- Distributed systems with heterogeneous communication mechanisms¹⁰
- Large amounts of generated data with real-time analysis requirements¹¹
- Strict regulations and requirements on safety and compliance¹²

Existing microservices practices for enterprise applications might therefore be inappropriate for the CPS domain. In this study, we aim to analyze differences in using microservices and DevOps in the CPS domain, especially when migrating legacy systems. The following research questions scope this objective:

RQ1: What are **challenges** related to microservices and DevOps in the context of CPS?

RQ2: What are **practices** related to microservices and DevOps in the context of CPS?

RQ3: What are **rationales** why CPS professionals want to migrate their systems toward microservices and DevOps?

To this end, we formed an industry-academia collaboration and conducted a mixed-method empirical study. First, we systematically analyzed scientific literature. Second, we validated and extended the results through an interview-based case study with industry experts from different units at Siemens AG*, a global technology enterprise with extensive CPS expertise in several domains, ranging from manufacturing, health care, or mobility to energy.

2 | BACKGROUND

In this section, we briefly describe CPS, why they evolve toward microservices and DevOps, as well as existing related work.

2.1 | Evolution of CPS toward microservices and DevOps

A major technological foundation for the ongoing Industry 4.0 transformation in the CPS domain is the IoT,¹³ or Industrial Internet of Things (IIoT). Such systems are characterized by the integration of mechanical and electrical devices that represent smart endpoints. Recent developments also utilize digital representations (digital twins), in particular for simulations in manufacturing systems.¹⁴ Gartner estimates the growth of IoT endpoints from 331.5 million in 2018 to 1.9 billion in 2028.¹⁵ This massive growth of involved devices poses new challenges for their efficient management. Mass-customization and unpredictable workloads require reconfigurable and highly automated platforms,¹⁴ which in turn has consequences for software architectures and development processes. Combemale and Wimmer express the need for more comprehensive views on the development and operation of such environments.¹⁶

Microservices are fine-grained units forming a single application that are individually scalable and deployable, and offer more flexibility than a monolith.¹⁷ Enabled by technologies like containerization, cloud-native applications heavily rely on this architectural style. DevOps comprises a variety of practices to connect development and operations in a seamless process. These practices promote a high degree of agility and automation, which is key for efficiently operating microservice-based systems. While it promises advantages for maintainability or scalability, the migration of existing systems towards microservices is a challenging task.⁸ Existing research yielded many microservices refactoring approaches to mitigate effort and partly automate tasks.^{18,19} However, their applicability is heavily dependent on several factors like,

*<https://www.siemens.com>

for example, the available input artifacts, targeted quality attributes, maturity of tool support, and its preceding evaluation. Moreover, organizational challenges like restructuring teams and establishing a DevOps culture need to be considered as well and sometimes even pose a bigger challenge, especially for large organizations.⁸ There are several studies available that collect empirical evidence on microservice migration in industry. These studies mainly cover enterprise systems like retail or consumer services. The CPS domain, however, has received little explicit attention yet for research on adopting microservices and DevOps.

2.2 | Related work

In their general CPS research roadmap, Törngren and Grogan discuss the inherent complexities of these systems, including causes, effects, and limitations of existing methodologies.²⁰ Similarly, Fahmideh et al. acknowledge the lack of research on the development lifecycle for IoT systems and conducted a mixed-method study to derive a process framework.²¹ It covers essential tasks from a SE perspective, but does not focus on DevOps and microservices.

Campeanu contributed an early mapping study on microservices architectures in the context of IoT and cloud computing.²² He cataloged 364 primary studies from 2015 to 2018, visualizing the rapid emergence of this technological combination. In a more holistic review, Joseph et al. discuss the use of microservices in specific domains.²³ They found 12 studies proposing specific solutions for the IoT domain, including related fields like health care or smart living. The authors emphasize that “the major challenge faced in developing IoT solutions is having to deal with a plethora of devices that are heterogeneous in nature.”²³ Similar findings were revealed in the review of 26 primary studies by Pereira et al. on the use of DevOps in IoT systems.²⁴ Heterogeneous device management was also found to be a main challenge.

Hasselbring et al. aim to mitigate lacking expertise in the adoption of Industry 4.0 solutions by proposing an approach to “introduce methods and culture of DevOps into industrial production environments,”²⁵ especially among small and medium-sized enterprises. Taibi et al. focus on the migration aspect toward microservices and DevOps in general in their systematic mapping.²⁶ They identified common architectural principles and patterns for microservices from 23 analyzed case studies. However, only one case was from the IoT domain.

In summary, several studies exist on either microservices and DevOps or the CPS domain. However, a comprehensive analysis of challenges and practices at the intersection of these three concepts is still missing. We aim to close this gap, and provide foundations for specific approaches for modernizing existing CPS.

3 | RESEARCH DESIGN

To answer the RQs, we formed a team of academic SE researchers and industry researchers from Siemens AG, who are experts on the CPS domain. We decided to first conduct a literature review to analyze the scientific state of the art for microservices and DevOps in the context of CPS (see Section 3.1). The results were then compared to similar publications on microservices in the enterprise application domain. Afterwards, we conducted an interview-based case study at Siemens AG to compare the literature findings with industry experiences (see Section 3.2). The results from both studies were holistically analyzed, discussed within the research team, and presented to a larger audience at Siemens AG for additional feedback. This research process is visualized in Figure 1. We share the most important study artifacts as a digital appendix via Zenodo[†].

3.1 | Literature review

The goal of a literature review is to identify and analyze existing publications for a certain topic, and to synthesize findings from the publication content. Structured approaches to literature surveys like systematic literature reviews²⁷ or systematic mapping studies²⁸ provide a rigorous and reproducible process, but also require considerable effort. Additionally, their results can be hard to integrate into industry practice. We therefore decided to conduct a Rapid Review,²⁹ a lightweight secondary study performed in close collaboration with industry and focused on providing timely and easily

[†]<https://doi.org/10.5281/zenodo.6536538>

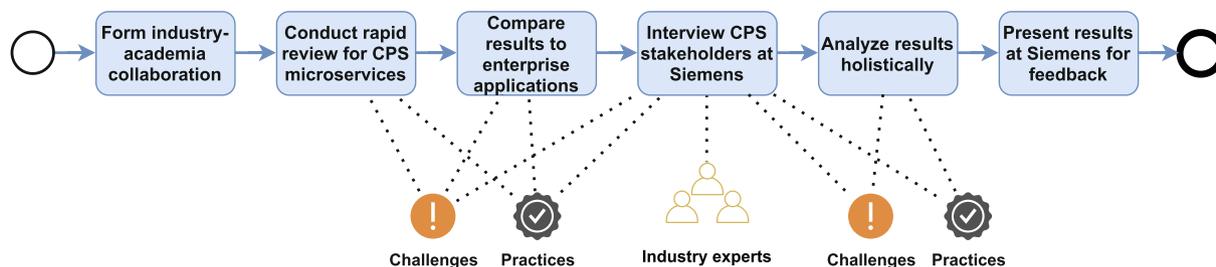


FIGURE 1 General research process

consumable results. The protocol of a rapid review is still systematic, but may consciously sacrifice rigor and extensiveness for industry relevance and efficiency. The first four authors conducted the review, with the rest engaging in discussion.

3.1.1 | Search strategy

By analyzing known relevant literature, discussing important keywords, and piloting a few alternatives, we finally decided to use the following two search strings, one for “microservices” and one for “DevOps”:

1. `microservices & (cps ∨ cyber-physical ∨ manufacturing ∨ industrial ∨ edge)`
2. `devops & (cps ∨ cyber-physical ∨ manufacturing ∨ industrial ∨ edge)`

An AND relation was used to combine each main term with CPS keywords, of which one had to match as well. Both search strings were executed independently. To reduce false positives, the query was only executed for the title field. We selected the following five common digital libraries/academic search engines:

- Google Scholar
- ACM Digital Library
- IEEE Xplore
- ScienceDirect
- SpringerLink

To ensure a manageable number of results, we defined a stopping criterion: only the first 50 publications per search string and source would be considered, that is, a maximum of 500 publications (2 search strings × 5 sources × 50 hits). After data extraction, one round of backward and forward snowballing³⁰ (reference and citation search) complemented this initial search. The rationale for this strategy was to combine sufficiently broad coverage with manageable effort.

3.1.2 | Study selection

We used the web-based tool Rayyan,[‡] which offers convenient features to include or exclude studies in a multi-user mode. Each paper was assigned to two researchers, who performed the selection independently and compared the results. Disagreements were discussed within the whole group until consensus was reached. Papers were included based on at least one of the following criteria:

- Paper describes CPS challenges associated with microservices or DevOps

[‡]<https://rayyan.ai>

- Paper describes a practice to develop CPS using microservices or DevOps
- Paper describes a practice to migrate CPS towards microservices or DevOps

Practices include, for example, processes, techniques (e.g., for system migration), platforms, reference architectures, algorithms, quality assurance practices, etc. Additionally, a paper could also be excluded for one of the following reasons:

- Short paper with four pages or fewer
- Full-text not available to us
- Not published in English
- Published before 2015
- Not peer-reviewed scientific literature (e.g., books, student theses, etc.)
- Duplicate or extension (the more recent paper is kept)

3.1.3 | Data extraction

Included papers were exported into a spreadsheet with columns for challenges and practices plus some additional columns for categorization, for example, what part of the CPS domain the paper focused on. To validate the sheet, all four researchers independently extracted the first five papers. Afterwards, results were compared, which further improved the sheet and led to more concrete extraction guidelines. We repeated this process a second time, after which we felt confident that our shared understanding leads to reasonably consistent extractions. For efficiency, we therefore assigned each remaining paper to a single researcher for extraction. If it was discovered during extraction that a paper did not completely fulfill the selection criteria, it was excluded.

3.1.4 | Snowballing

All successfully extracted papers were used in one round of snowballing. For backwards snowballing, the references of all papers were automatically extracted from the PDF via a tool.[§] The references were de-duplicated, and equally distributed among the four researchers for selection. For forward snowballing, the included papers were again split up, and each researcher examined their citations via Google Scholar up to a maximum of 50 citations per paper. All newly proposed papers were examined by the other researchers and only kept if nobody objected. Afterwards, we extracted the snowballed papers following the same process as described above.

3.1.5 | Data synthesis

The final extractions were then holistically analyzed to synthesize answers to our RQs. For basic columns like the CPS domain or targeted SE activities, not much effort was required to consistently aggregate them. However, since challenges and practices were extracted via direct paper quotes, we had to apply thematic synthesis.³¹ We first assigned codes to each individual extraction (*open coding*), and afterwards synthesized these codes into higher-level themes (*axial coding*). This was a very iterative process, with lots of discussion and the frequent merging or splitting of codes and themes. The final codes and themes were aggregated and visualized using charts.

3.2 | Interview-based case study

To validate and enrich the findings from the rapid review, we conducted an interview-based case study at Siemens AG. A case study³² is a research method to deeply analyze something in a real-world setting (a case), usually following a

[§]<https://github.com/helenocampos/PDFReferencesExtractor>

qualitative and exploratory methodology. While a case study can provide rich insights into a topic, the generalizability of the produced results may be limited. We conducted an *embedded case study*,³² that is, our case was Siemens AG as a CPS manufacturer and our units of analysis were individual CPS in different organizations at Siemens. These systems were studied through the experiences and opinions of their stakeholders, that is, interview participants in various roles with sufficient industry experience (at least 3 years). Ideally, their systems should also be in the process of migrating toward microservices and DevOps, or such a migration was considered for the future. We used convenience sampling via our personal networks at Siemens AG to find matching participants. For flexibility, we chose semi-structured interviews.³³ The interviews were conducted by the first three authors.

3.2.1 | Study preparation

We created several artifacts for the interviews. An *interview preamble*³² explaining the process was used to recruit participants and to make them familiar with the study. The preamble also described ethical considerations and requested consent for audio recordings. Additionally, we created an interview guide³³ with the most important questions. This guide was used to loosely structure the interviews. We also prepared a small slide set with additional information for certain questions. Lastly, we created a case characterization matrix³³ with basic attributes for the analysis.

3.2.2 | Evidence collection

We conducted nine individual interviews, all of them via remote communication software. Three interviews were held in English and six in German. Interviews took 45–70 min, and all participants agreed to a recording. At least two researchers were present during each interview, and swapped between moderator and note writer. After a brief introduction, we showed a slide with microservices characteristics to ensure the same basic understanding. We then inquired basic demographics, and let participants describe their system from the CPS domain. Afterward, we asked about challenges and practices related to microservices and DevOps, specifically in a migration context. Additionally, we asked participants about their rationales for migrating their CPS toward microservices and DevOps. To support participant reflections, we briefly introduced a migration methodology created by us,[¶] where participants could describe their experience in each step.

We manually transcribed each recording into a text document. These documents were then sent to interviewees for review and approval. Participants could delete sensitive content or adapt unclear statements. We then used the approved transcripts for qualitative analysis.

3.2.3 | Data analysis

To synthesize the interview findings, we used card-sorting,³⁵ a lightweight, collaborative analysis technique. Instead of paper cards, we used a web-based tool with feedback cards for agile retrospectives.[#] Each interview block was assigned to one researcher, who then extracted relevant quotes per transcript into individual cards. In an iterative process, these cards were then sorted according to their similarity, and preliminary labels were assigned. Lastly, the results were presented to the other researchers. Together, we refined the labels and formed higher-level categories. Finally, important findings were extracted into spreadsheets, aggregated, and visualized.

3.3 | Data availability

For transparency and reproducibility, we share all important study artifacts for both the rapid review and the interviews on Zenodo.^{||} This includes, for example, the list of primary studies (CSV), the extracted practices and challenges (CSV), the interview guide, and the aggregated interview results (CSV).

[¶]The details of the methodology can be found in another paper from us.³⁴

[#]<https://metroretro.io>

^{||}<https://doi.org/10.5281/zenodo.6536538>

4 | RESULTS

In this section, we first present the results for the rapid review (challenges and practices) and then for the interviews (rationales, challenges, and practices).

4.1 | Rapid review

In total, we included 146 primary studies in the review, 87 of which were discovered in the initial search and another 59 during snowballing (see Figure 2). The papers were published between 2015 and 2021, and covered all phases of the SE lifecycle, with design being the major focus (mentions in 127 papers), followed by deployment (44) and operation (31).

4.1.1 | Challenges (RQ1)

From all 146 primary studies, we extracted 405 statements, which were converted into 46 challenge labels in 11 categories during the analysis. The 46 challenges were mentioned a total of 490 times among all statements. A challenge was counted at most once per paper. Figure 3 shows the 11 challenge categories with their number of mentions as orange bars, while Table 1 shows the 10 most frequent challenge labels.

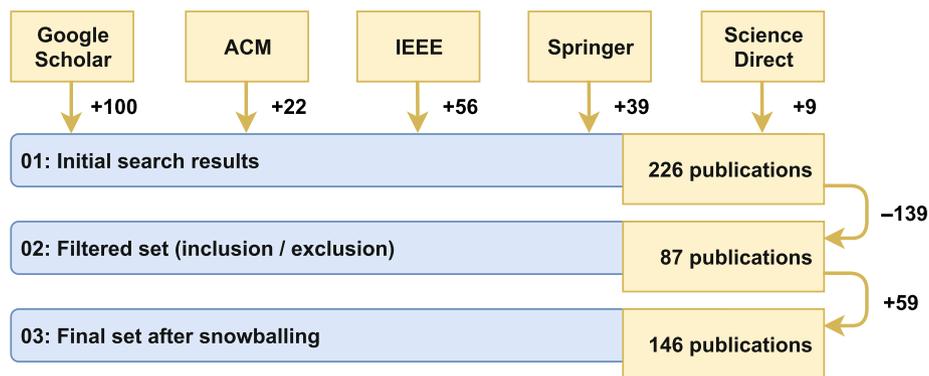


FIGURE 2 Literature review: number of publications per stage

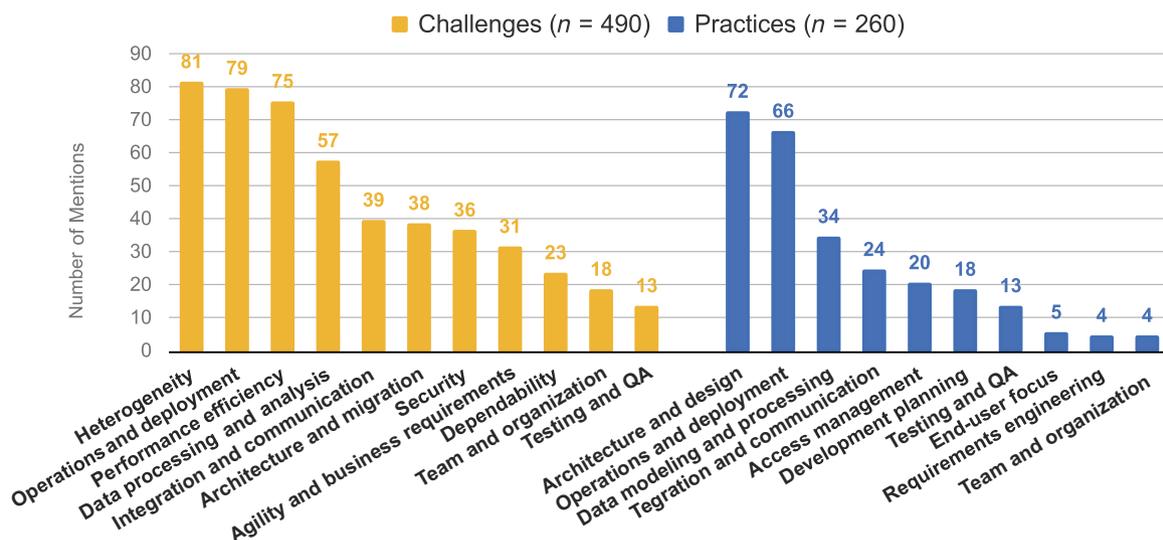


FIGURE 3 Literature review: challenge and practice categories

TABLE 1 Most mentioned challenges and practices (by number of papers)

Challenge	Number	Category
Technological heterogeneity of devices	48	Heterogeneity
Real-time processing requirements (response time of cloud microservices too high)	31	Performance efficiency
Limited hardware resources of devices and edge nodes	29	Performance efficiency
Large amount of data to process	25	Data processing and analysis
Variety of data formats and communication protocols	24	Heterogeneity
Evolvability is limited by traditional architectures (monoliths, SOA)	21	Architecture and design
Efficient management of a large number of components (microservices, devices, etc.)	20	Operations and deployment
Establishing connectivity and interoperability	19	Integration and communication
Complex data storing and processing requirements	18	Data processing and analysis
Efficient microservices placement, scheduling, and load balancing	18	Operations and deployment
Practice		
Architecture proposals for various application contexts (ICS, IIoT, cloud robotics, health care, manufacturing, etc.)	55	Architecture and design
Techniques for placement, scheduling, allocation, deployment in various domains	35	Operations and deployment
Microservices for data processing and analytics on IoT/edge nodes	18	Data modeling and processing
Conceptual modeling and model-driven engineering	16	Data modeling and processing
Security techniques for IoT/edge systems	16	Access management
DevOps adoption in various application contexts (CPS, ICS, IoT)	15	DevOps process
Techniques and patterns for monitoring of microservice-based systems	12	Operations and deployment
Techniques, tools, patterns for efficient cloud/edge deployment	8	Operations and deployment
Techniques and patterns for microservice registration and discovery	7	Integration and communication
Guidelines for determining microservices boundaries	6	Architecture and design
Microservices architectural style to meet requirements of IoT/edge gateways	6	Architecture and design
Orchestration of services / containers for automated deployment in IoT/edge solutions	6	Operations and deployment
Simulation of IoT environments / applications to facilitate testing and evaluation	6	Testing and QA

Abbreviations: CPS, cyber-physical systems; ICS, industrial control systems; IoT, Internet of Things; IIoT, industrial Internet of Things.

The dominant category was *heterogeneity*, with *technological heterogeneity of devices* (48) and *variety of data formats and communication protocols* (24) being the most mentioned challenges in it. Numerous devices also pose challenges for *operations and deployment*. In this category, the management of components like the *efficient placement, scheduling, and load balancing of microservices* (18) were prevalent issues. The third most popular category was *performance efficiency*, which contained the second and third most mentioned challenges overall: *real-time processing requirements* (31) and *limited hardware resources of devices and edge nodes* (29).

The large amount of data to process (25) together with complex data storing and processing requirements (18) highlight prevalent challenges for data management and processing. In general, we observed a strong tendency toward hardware-related aspects and the management and operation of a multitude of distributed components. One of the major aspects here was establishing connectivity and interoperability (19), which was the most mentioned challenge in the integration and communication category. The *architecture and design* category was mostly driven by the limited evolvability of traditional architectures like monoliths and SOA (21), which expresses the need for more suitable architectural patterns.

The *security* category with issues around data privacy and authorization for a multitude of components (12) and strong requirements for security (9) also received considerable attention. *Agility and business requirements* pose challenges on CPS as well, as the two most frequent challenges in this category reflect: highly dynamic and uncertain environments (10) and quickly emerging technological advancements require more flexibility (8). The dependability category was dominated by challenges around strict safety and compliance regulations (10) and strong requirements for reliability (8). Organizational and process-related challenges, however, played a minor role. The most mentioned challenge in this category was adapting IoT/CPS processes and philosophy to DevOps (6).

4.1.2 | Practices (RQ2)

For practices, we extracted 283 statements, which were condensed into 32 labels in 10 categories. The 32 practices were mentioned a total of 260 times among all statements.

As with challenges, a practice was counted at most once per paper. The first of two dominant categories was architecture and design, with specific architecture proposals for various application contexts (55) being the most frequent label, for example, microservice architectures for industrial control systems (ICS) or cloud robotics. Other practices in this category were rather insignificant. The second dominant practice category was *operations and deployment*, with the top label being techniques for placement, scheduling, allocation, deployment in various domains (35), often for edge-cloud continuums. Other labels received considerably fewer mentions: the second most addressed practice was targeting techniques or patterns for monitoring of microservice-based systems (12).

All other categories received notably less attention. Data modeling and processing was the third most popular category, with practices centered around microservices for data processing and analytics on IoT/edge nodes (18) or conceptual modeling and model-driven engineering (16). Several papers also proposed practices for integration and communication aspects, such as techniques and patterns for microservice registration and discovery in IoT/IIoT environments (7), followed by techniques for autonomous/automated device management and reconfiguration (5). Access management centered around various security techniques for IoT/edge systems (16). Development planning was dominated by practices for DevOps adoption in various application contexts (15), such as ICS or IoT. While several DevOps-related aspects are also covered by other categories and labels, this label focuses on more general, process-related facets. Finally, testing and QA solutions received relatively little attention (13), similar as with challenges in this area. Most practices in this category were targeting simulation and performance testing.

In general, practices were more condensed within a few categories, dominated by concrete architecture proposals and techniques for the dynamic scheduling of microservices within edge/fog/cloud computing. Organizational and process-related solutions received noticeably little attention in our review.

4.2 | Interviews

In total, we conducted nine interviews with participants from different business units across Siemens AG. Table 2 lists their demographics. The experience column gives their professional experience in years, with the amount of microservice experience in parentheses. Except for P7, all participants had more than 10 years of professional experience. However, their prior experience with microservices and DevOps varies considerably, with two participants with no prior experience,

TABLE 2 Participant demographics (years of microservices experience in parentheses)

ID	Role	Experience	System
P1	Architect, researcher	18 (3)	S1
P2	Architect	19 (3)	S2
P3	Software development lead	20 (6)	S3
P4	Architect	15 (2)	S4
P5	Architect	25 (0)	S5
P6	Architect	16 (0)	S4
P7	Researcher	6 (5)	S6
P8	Project manager	14 (4)	S6
P9	Project manager	15 (5)	S7

TABLE 3 System and migration properties

ID	Purpose	Size	Staffing	Timeframe	Targeted architecture
S1	Manufacturing management system	n/a	50 (5 teams)	2020–	Aspects of microservices
S2	Engineering of automation systems	70 MLOC	300	2018–2026	Aspects of microservices
S3	Digital health care services platform	100 MS	250 (25 teams)	2017–2023	Full microservice architecture
S4	Digital health care services platform	15 MLOC	300 (30 teams)	2019–	Full microservice architecture
S5	Control software for electric motors	5 MLOC	100	2011–	Modularization
S6	Digital twin platform for power distribution equipment	150 kLOC	10	2021–	Full microservice architecture
S7	Process control engineering	600 kLOC	250 (12 teams)	2021–2023	Aspects of microservices

and 2–6 years for the rest. The interviewees described seven different systems across a variety of fields. Table 3 shows their purpose, size, staffing, and the projected migration timeframe. As we can see, even the “smallest” system (S6) still had over 100,000 lines of source code.

All systems were in the process of a migration. However, the expected migration timeframes differed considerably, ranging from two years (S7) to over a decade (S5). Often, no distinct end date was mentioned, for example, for S1, S4, or S6. For some systems, the reason was an ongoing evaluation of the feasibility of a full migration. In other cases, the migration followed a continuous evolution strategy, where components were slowly migrated toward microservices when they had to be changed in the course of day-to-day development. We also discussed the architecture that should be achieved with the migration, which varied regarding the commonly understood characteristics of microservices,¹⁷ that is, the characteristics we introduced at the beginning of the interviews. For three systems, the targeted architecture was a full microservice architecture that relied on almost all characteristics. For another three systems (denoted with aspects of microservices in Table 3), the targeted architecture fulfilled some microservices characteristics, but not enough to be considered a full microservice architecture. That means S5 was the only system where the team mainly aimed to achieve a higher modularization. All other systems set out for a partial (i.e., some characteristics) or full microservice architecture (i.e., nearly all characteristics).

4.2.1 | Migration rationales (RQ3)

When asked for their migration rationales, that is, why they modernized their systems toward microservices and DevOps, our participants mentioned a total of 10 different reasons, most of them quality attribute goals (see Figure 4). For six of the seven systems, increased development productivity was a main driver (all systems except S1). Decreasing cycle time, establishing faster feedback loops, or providing new functionality faster were prevalent themes here. Related to this was the goal

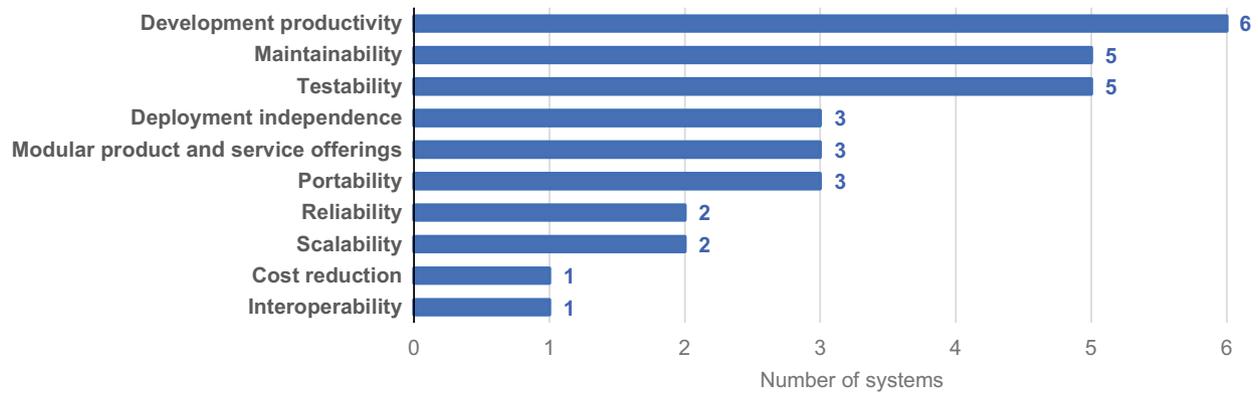


FIGURE 4 Distribution of interview migration rationales

to improve maintainability, which was mentioned for five systems (S2, S4, S5, S6, and S7). While better maintainability ultimately should lead to better productivity, the mentioned reasons here were more technical and structural, for example, improving modularity, reducing coupling, improving code quality, reducing complexity, or improving reusability. Five systems (S2, S3, S4, S5, and S7) also had the goal to improve testability, that is, to be able to test smaller components in isolation and in less time.

Other rationales occurred less frequently. For three systems (S3, S4, and S6), deployment independence was explicitly mentioned, that is, each team should be able to release and deploy their services independently of all others. More a business than a technical driver was the need for modular product and service offerings. This was named for S1, S2, and S4, where fine-grained and more specialized offerings to customers and avoiding redundant offerings were very important. In three cases (S1, S2, and S6), portability was mentioned as a migration rationale, that is, supporting various deployment environments like cloud versus edge, Windows versus Linux, digital twin platform versus real hardware, etc. Less prevalent rationales were improving reliability (S4, S6) or scalability (S4, S6), even though these two are typical quality attributes associated with microservices. Lastly, cost reduction (S6) and interoperability (S2) were only mentioned once.

4.2.2 | Challenges (RQ1)

Our interview analysis revealed 26 unique challenge labels, from which we could reuse 16 from the literature review, that is, 10 were newly created (38%). These individual challenges were grouped into 10 categories, all of them reused from the review. The most prominent categories were architecture and design (nine mentions), team and organization (8), integration and communication (5), and performance efficiency (5). Categories like *operations and deployment* (3), *MDE and digital twins* (2), or *security* (1) were less prominent. Furthermore, individual labels were pretty spread out, with 13 challenges being mentioned only once and no challenge being mentioned for more than three of the seven systems, that is, there were no “top” challenges affecting most CPS in our sample. The category distributions are shown in Figure 5, while the most mentioned challenges are presented in Table 4. For the top categories, we present further details below.

Architecture and design

All systems except S1 reported at least one challenge from this category. The most mentioned label was reducing coupling or hidden dependencies (S3, S4, S6), that is, participants found it challenging to untangle elements in preparation for the migration or service decomposition. Another challenge in this area was technical debt of legacy systems (S2, S5), which would slow down a migration considerably. All other challenges like lack of architecture documentation (S6), massive system size (S2), or incorporating legacy components (S6) were only mentioned once. Lastly, only one label (limited evolvability due to traditional architectures like monoliths or SOA from S7) was reused from the review, as most mentioned challenges here were not strongly specific to CPS.

Team and organization

Challenges related to human factors were also very prominent. For S4, S5, and S7, the lack of personnel with modern technology experience was problematic, that is, many CPS engineers would know little about microservices or DevOps.

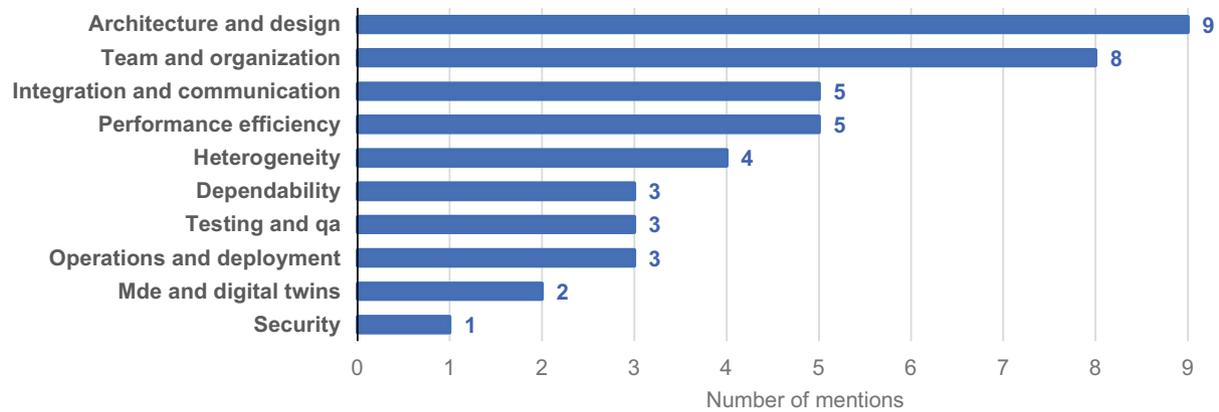


FIGURE 5 Distribution of interview challenge categories

TABLE 4 Challenges affecting two or three systems

Challenge	Number of systems	Category	In LR
Reducing coupling / hidden dependencies	3	Architecture and design	No
Establishing connectivity and interoperability	3	Integration and communication	Yes
Real-time processing requirements	3	Performance efficiency	Yes
Lack of personnel with modern technology experience	3	Team and Organization	Yes
Technical debt of legacy system	2	Architecture and design	No
Strong requirements for reliability	2	Dependability	Yes
Technological heterogeneity of devices	2	Heterogeneity	Yes
Variety of data formats and communication protocols	2	Heterogeneity	Yes
Predictability and consistency in distributed processes	2	Integration and communication	Yes
Limited hardware resources of devices and edge nodes	2	Performance efficiency	Yes
Building developer confidence in microservices and DevOps	2	Team and organization	Yes
Creating shared system understanding with all teams	2	Team and organization	No
Creating new decoupled test suites	2	Testing and QA	No

Project manager P9 remarked that “architects with the required knowledge are currently most wanted in the project and are therefore somewhat of a bottleneck.” Related to that, building developer confidence in microservices and DevOps was sometimes challenging (S5, S6), and could lead to “trust issues” (researcher P7) if not properly managed. For very large systems, a challenge not mentioned in the literature was creating a shared system understanding with all teams (S2, S4). Lastly, adapting the CPS processes and philosophy to microservices and DevOps was problematic for S7, for example, P9 saw the need for a cultural change while moving to DevOps and cross-functional teams in his department.

Integration and communication

This category was only comprised of two labels. However, they were mentioned for two and three systems respectively, and both appeared in the literature review. For S1, S4, and S6, establishing connectivity and interoperability between the newly

distributed services was difficult. Architect P1 saw “the challenging aspects of communication via message-passing” as the main problem in the migration. Related to this was the challenge to maintain predictability and data consistency within distributed processes (S1, S4). Eventual consistency was not acceptable in some use cases, so creating “several services that are independent yet for which we also need to guarantee data consistency” (P6) was difficult and therefore sometimes completely avoided.

Performance efficiency

In this category, two challenges were mentioned, both of which could be reused from the literature review. As expected, several CPS (S1, S4, S5) had strong real-time processing requirements, which needed to be kept intact during and after the migration. This meant that the response time of cloud-based microservices was often not sufficient, for example, for “functionalities that need to be close to the production line” (P1). For some use cases, interprocess communication between microservices was not feasible at all, or as described by architect P5: “we need to calculate cycles for electric current regulators within a two-digit microsecond window.” A related challenge were the limited hardware resources of devices and edge nodes (S4, S5). This could also impact latency or scalability, but mainly limited using resource-intensive technologies like containerization.

4.2.3 | Practices (RQ2)

We identified practices in the following categories, each listed with the number of mentioned practices: *team organization* (12 mentions), *process strategy* (12), *tool support* (11), *service identification* (10), *DevOps* (8), and *architecture and infrastructure* (7). As we provided the interviewees with a stepwise migration methodology created by us, the practice categories were prescribed to some extent. Figure 6 shows the categories we queried about with the number of mentions, respectively. We discuss them in detail below.

Team organization and DevOps

Six participants reported an associated team re-organization during their migration. The prevalent motivation was promoting team independence and fostering collaboration between teams, mentioned four times. Two participants reported no organizational changes (P5, P6), and also stated that establishing DevOps practices was not a priority yet. Reported DevOps practices were establishing CI/CD pipelines (S1, S2, S3, S6) and using cloud platform services (S1, S3). Four interviewees had already established a DevOps culture. Two aimed to do so (S4, S7), but were hindered by technical or organizational obstacles. One participant remarked that team re-organization was a sensitive topic in Germany, and related it to the difficulty of DevOps adoption in large enterprises.

Process strategy

The migration of a legacy system can follow different strategies. For brownfield developments, we distinguish between a *rebuild* and *refactor* type. The former is often used when the system is rebuilt using newer technologies and was applied for systems S1, S2, and S4. Sometimes, even a mix of several strategies can be used for different parts or subsystems, as was the case for S2, S4, and S6. We can further distinguish between a *big bang* migration that aims to minimize the duration,

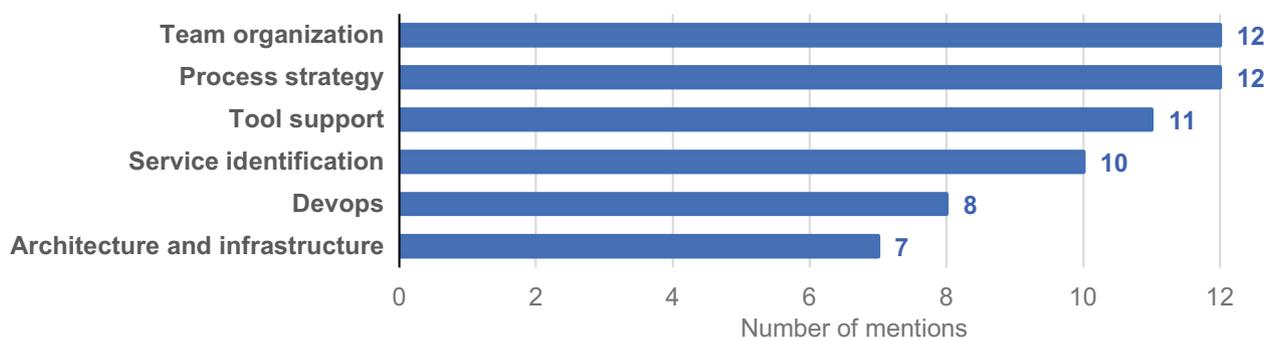


FIGURE 6 Distribution of interview practice categories

and a continuous evolution approach that tries to minimize the needed resources. The latter was used for four systems (S4, S5, S6, S7), and often based on the popular *Strangler* pattern³⁶ (S2, S3, S4). Gradually building the new system with parallel operation next to the existing one was reported for S4 and S6. As rationales, participants stated limited resources, time pressure, and strong requirements for stability and compliance.

Service Identification and Tool Support

Existing literature identifies the decomposition task as the most complex one in a migration scenario.⁸ In a similar notion, P1 stated that “success or failure will depend on the services having good granularity or not.” Five of our participants described service decomposition as a manual task. P1 and P3 reported Domain-Driven Design as the used technique, others mentioned static code and metadata analysis (P2, P9), as well as dynamic analysis (P5) of the existing system. Our participants used a variety of mostly commercial tools that support the system comprehension via static code analysis (P2, P3, P6, P8). However, a subsequent automatic decomposition is often beyond their scope. For this purpose, P5, P6, and P9 relied on self-developed tools, two involved external consultants to compensate missing in-house expertise (P1, P5), and another two did not use any tools for this purpose (P4, P7).

Architecture and infrastructure

Three participants mentioned workshops as a practice to support the architecture design and infrastructure choices (P2, P3, P4). Conducted cross-team and involving all system stakeholders, they facilitate the build-up of a common understanding and the identification of requirements. P6 expressed his interpretation of agility in the following way: “Big upfront design is a no-go. Just as much architecture as necessary.” For the preferred cloud model, platform as a service (PaaS) hosting was mentioned twice. On-premise hosting was a strong requirement for S7 due to its tight integration with physical devices.

5 | DISCUSSION

The review and interview results show many similarities, but also reveal certain discrepancies. In the following, we first compare the CPS results of our rapid review to similar literature studies or surveys for enterprise applications. Afterwards, we contrast the rapid review results with the interview results.

5.1 | Comparison with enterprise application microservices

To highlight the peculiarities of the CPS domain and show differences to common enterprise applications, we extracted challenges and practices from existing studies on enterprise application microservices. For this purpose, we relied on our extensive experience with enterprise application microservices in industry^{8,37-40} to select and analyze nine publications, primarily recent literature studies and two interview studies, that focus on general practices and challenges in the microservices and DevOps migration context. Table 5 lists these nine publications along with the number of included primary sources or study participants. The last two columns show the number of extracted challenges and practices per paper. Altogether, we extracted 513 challenges and 114 practices from all sources. For reviews, we counted a challenge or practice per primary source that stated it, for the interview and survey paper per participant that stated it. To ease comparability, the labeling was harmonized with our preceding literature review by using the same categories, as can be seen in Figures 7 and 8.

5.1.1 | Challenges

Heterogeneity reflects the most mentioned challenge with CPS systems. It goes along with achieving sufficient system performance as well as efficient data processing in such environments, which are the third and fourth most mentioned challenges for CPS. We attribute these aspects to typical IoT/IIoT systems, which are characterized by operating a large number of devices with limited hardware resources. Compared to enterprise applications, these three challenge categories, plus *integration and communication* aspects, are seen as more important by the CPS literature. Challenges with enterprise systems, on the other hand, are strongly centered around *organizational aspects*. This category seems to have

TABLE 5 Empirical studies and reviews on enterprise application microservices

ID	Title	Year	Type	# Sources	# Chal.	# Prac.
1	Deployment and communication patterns in microservice architectures: A systematic literature review	2021	SLR	38	57	14
2	Monoliths to microservices - Migration Problems and Challenges: A SMS	2021	SMS	37	63	2
3	Understanding and addressing quality attributes of microservices architecture: A Systematic literature review	2021	SLR	72	0	20
4	A Systematic Mapping Study on Microservices Architecture in DevOps	2020	SMS	47	72	37
5	Promises and challenges of microservices: an exploratory study	2020	Interview + Survey	21+37	111	21
6	Architecting with microservices: A systematic mapping study	2019	SMS	103	0	7
7	Microservices Migration in Industry: Intentions, Strategies, and Challenges	2019	Interview	16	70	0
8	Continuous Integration, Delivery and Deployment: A Systematic Review on Approaches, Tools, Challenges and Practices	2017	SLR	69	140	13
Total				440	513	114

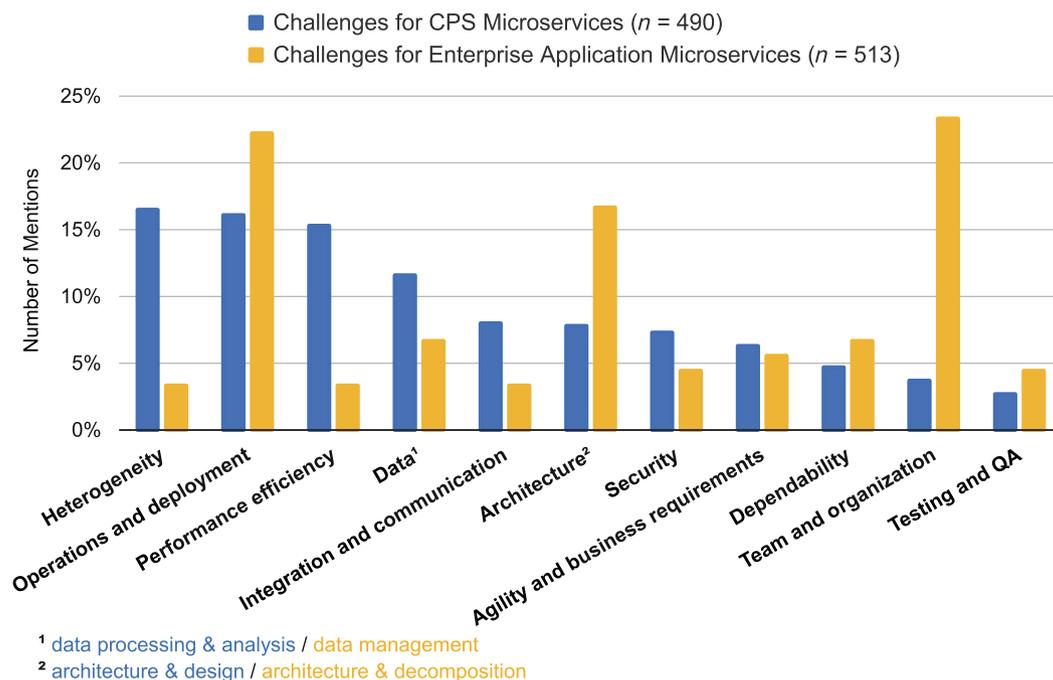


FIGURE 7 Differences in reported microservices challenges between cyber-physical systems and enterprise applications domains

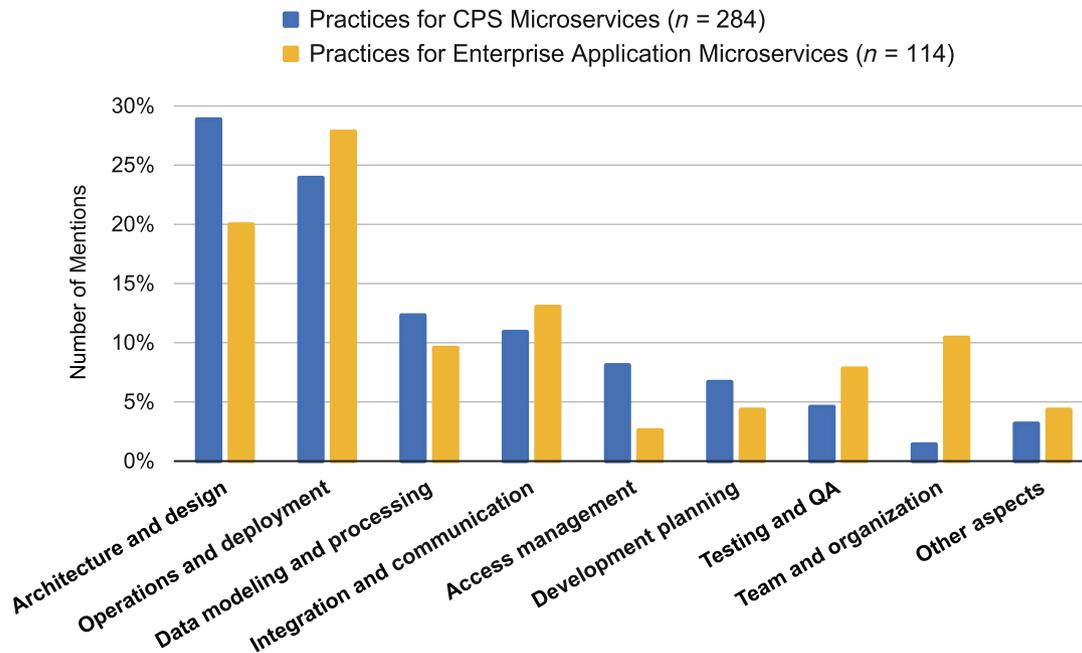


FIGURE 8 Differences in reported microservices practices between cyber-physical systems and enterprise applications domains

received very little attention in the CPS literature, as the big gap in mentions reveals. Difficulties with efficient *operations and deployment* follow, which is a very relevant category also for CPS systems. When it comes to *architectural aspects*, the gap again reveals a higher importance for typical enterprise systems. Here, challenges are mainly centered around finding a suitable decomposition into microservices. As many CPS are naturally designed as distributed systems, the modularization aspect might therefore be perceived as less challenging compared to other design aspects. *Security concerns* showed a little higher relevance for CPS systems, while the remaining categories draw a largely balanced picture. Interestingly, *dependability-related challenges*, which one would intuitively attribute to CPS microservices, played no major role in both domains, and received even slightly more mentions for enterprise application microservices.

5.1.2 | Practices

The comparison of practices shows a more balanced picture. We again see a big gap related to *organizational matters* and for *architecture and design* practices. A closer look into the latter reveals that architectural solutions proposed for CPS account for distributed infrastructure, heterogeneity, and performance aspects. Solutions for the enterprise domain, on the contrary, focus mainly on decomposing monolithic architectures into a suitable set of services, which correlates with the identified challenges in this category. A similar difference shows up for the *operations and deployment* category. Deployment solutions for CPS mainly address challenges of automation, placement of services, load balancing and monitoring. Solutions for enterprise systems prioritize the automation aspect as well, but focus less on performance-related aspects than on integrated CI/CD pipelines that foster agility and short release cycles. The inherently more challenging *security* aspect for CPS systems is reflected by a higher number of practices around necessary access management. Other aspects include techniques and tools to improve autonomy or automation, managing product variants, and customer involvement.

5.2 | Comparing rapid review and interview results

The primary studies in our rapid review reflect a strong focus on hardware-related aspects around the management and operation of heterogeneous environments, especially the identified challenges. This is augmented with the complexity of big data processing and real-time conditions. To address unsuitable CPS architectures that fail to address evolvability,

operability, or performance efficiency, researchers proposed many concrete microservice-based architectures for various CPS contexts. They predominantly offer solutions for efficient operations and deployment in the cloud-edge continuum, but also for data processing and modeling.

However, our interviewees from Siemens AG perceived most challenges with different importance. The top two review categories (*heterogeneity* and *operations and deployment*) played no major role for most analyzed systems, with only a few mentions. Standardization in domains like health care (partly) addressed the former, while the latter was more relaxed because none of the analyzed systems needed dynamic edge-cloud deployment and scheduling at runtime. Another reason for the lack of operational challenges might be that most systems were still early in the migration. Instead, *architecture and design* was the dominant challenge category, with a focus on the difficulty of modularizing and decomposing monolithic systems. Mentioned issues were commonly associated with legacy systems, for example, huge size, hidden dependencies, or accumulated technical debt. Consequently, the decomposition into services was a major challenge among our interviewees that was mostly described as a manual task with a lack of adequate tool support.

Another difference was the importance of *team and organizational challenges* in the interviews, while these seemed neglected in the review. Establishing a DevOps culture and building developer confidence were seen as challenging, and the lack of employees experienced with current technologies could present a bottleneck. Consequently, our interviewees reported compensatory practices around promoting team independence and collaboration, for example, through conducting workshops. Moreover, we also saw a preference for continuously evolving a system towards microservices and DevOps rather than a swift shift, which was, however, also due to strict quality and compliance requirements.

Contrary to the review, the interviews mostly reflected typical challenges and practices of enterprise microservices, as described by, for example, Baškarada et al.⁷ and Fritzsche et al.⁸ Except for performance challenges (e.g., real-time processing requirements), typical CPS challenges like heterogeneity, dependability, or security were not seen as problematic for using microservices. For several systems, that meant that not all microservices characteristics were desired for every part of the system. However, this tendency has also been observed for enterprise applications.^{37,41} The good news about these findings is that many general approaches for microservices and DevOps are also valuable for the CPS domain. Nonetheless, slight adaptations or tradeoffs are needed in some cases, as several CPS characteristics still seem to have an influence, albeit not a major one.

5.3 | Main takeaways

In the following paragraphs, we briefly synthesize the main takeaways of the combined study results and their implications.

5.3.1 | The CPS and enterprise application literature differ in the importance of several microservice-related challenges

For example, *heterogeneity* and *performance efficiency* were perceived as much more challenging for CPS microservices than for enterprise application microservices. Conversely, challenges related to *team and organization* were barely discussed in the CPS literature, while this was one of the top challenge categories for enterprise application microservices. This may imply that the different usage scenarios and quality requirement prioritization between these two types of systems lead to a slightly different research focus. Since the introduction of microservices in the CPS domain happened later than for enterprise applications, it might also mean that some industry-relevant challenges are not yet discussed frequently, but will appear more in future publications, once the field matures a bit more.

5.3.2 | Osmotic computing⁴² was a prevalent microservices theme in the CPS literature, but not at Siemens

As a concrete example of the different research focus, many proposed CPS practices in the literature centered around the dynamic scheduling and deployment of microservices at runtime within the cloud-fog-edge continuum, for example,

based on the origin of the request or the current load of the system. However, not a single one of the studied CPS at Siemens had this requirement. It may be that osmotic computing is only relevant for CPS outside of Siemens's scope, for example, very large, geographically distributed IoT environments, or that research in this space is still too early to broadly carry over to industry.

5.3.3 | Improved development productivity and maintainability were the primary migration rationales at Siemens

CPS stakeholders predominantly wanted an architecture that would enable them to deliver features faster and more sustainably in the long run. While this is also one of the main drivers in the enterprise application literature,³⁷ other prominent microservices quality attributes like scalability, elasticity, or interoperability were rarely mentioned at Siemens. A potential reason could be that the average CPS may not experience as frequent or rapid changes in the required throughput as typical internet-facing applications.

5.3.4 | Typical CPS requirements like performance, reliability, security, or privacy were not perceived as major challenges regarding microservice migration at Siemens, but were mentioned for selected systems

Three systems named performance as a challenge with microservices (S1, S4, S5), two systems reliability (S1, S4), and one system each did the same for compliance (S3) and security (S6). When explicitly queried for such quality attributes, other participants described that, while these requirements existed to a certain degree, the teams had experience in assuring these qualities, as they had to take care of them anyway in their domains. Using microservices would not fundamentally change this, that is, it would not be a migration-related challenge. This implies that it is highly situational if assuring these quality attributes is challenging with CPS microservices and if it requires special methods or tool support.

5.3.5 | Several studied CPS at Siemens did not target a full microservice architecture with all characteristics

Apart from S3, S4, and S6, the systems in our sample consciously decided to only aim for certain microservice characteristics and to leave out the rest, which was sometimes motivated by special CPS requirements. For example, the "no shared databases" principle was consciously ignored for S4 because data consistency was just too important for their regulated domain. This is important to keep in mind for researchers that work on methods or tool support for microservices in the CPS domain.

5.3.6 | Most frequently mentioned challenges and practices at Siemens were not highly specific to CPS

Similar to the disconnect between the literature on CPS versus enterprise application microservices, there are also several differences between the CPS literature and our case study findings at Siemens. While the above sections mentioned several facets where the CPS domain had an influence on the usage of microservices, the overall sentiment regarding microservices challenges and practices was much closer to the enterprise application literature than to the CPS literature. A potential reason could be that our rapid review contained many papers on large, decentralized IoT environments, which have some differences to the typical CPS at Siemens. Additionally, it could be argued that a few of our studied systems are more platforms or applications from the CPS domain rather than actual CPS. Nonetheless, it seems plausible that a large proportion of methods and tools for enterprise application microservices will at least be partly of value for CPS microservices. Some flexibility for the pointed out differences will be helpful, but the extensiveness of the required adaptations should be subject to future research.

5.4 | Threats to validity

We generally aimed for a rigorous study design, but also made some tradeoffs regarding efficiency, as is common in rapid reviews. The mixed-method approach and the collaboration between industry and academia could have partially mitigated this. Still, several limitations need to be mentioned for our results, which we structure according to three common notions of validity described by Wohlin et al.⁴³

Construct validity refers to a potential gap between operationalized measurements and the constructs targeted in the study, that is, if the collected data is suitable for the intended analysis purpose. In our study, data collection was text-based and qualitative, namely extractions from papers and interview transcripts with subsequent thematic analysis. To combat potential threats in this area, we spent great care to arrive at a unified understanding of our central constructs in the research team, for example, microservices, practices, challenges, rationales, etc. Additionally, we introduced important constructs to interview participants, posed clarifying questions when encountering ambiguous terms, and sometimes briefly reflected our summarized understanding back to interviewees. Nonetheless, there is a slight possibility that we interpreted a few instances differently than the authors of our primary studies or our participants. Furthermore, interviews rely on the described experiences and opinions of participants, which are subjective. A few interview statements could even be incorrect without us realizing. In general, however, we estimate the impact of threats in this area to be very low.

Internal validity is concerned with potential confounding factors and the rigor and consistency of the study design. While we consciously chose a more lightweight research method for the literature review, we still followed a very systematic process. At least two researchers decided if a study should be included or excluded. While only a single researcher performed the data extraction for most papers, we only split up the work after having piloted the extraction sheet and improved our shared understanding of it with 10 papers. Additionally, all important findings were discussed among several researchers to counteract potential subjective bias. For the interviews, confidentiality and anonymity were provided to enable participants to talk freely about their experiences. Participants also seemed not to be afraid to describe negative or more sensitive areas of their systems. We also tried to limit our influence as interviewers on our participants, for example, by avoiding leading questions or any judgmental statements. However, despite these countermeasures, there is still a chance for nonmitigated influences on data collection, which we nonetheless perceive as noncritical.

Lastly, *external validity* is concerned with how well the findings can be generalized to other settings or populations. From the discussed facets of validity, this is the weakest one in our study. A rapid review usually covers a lower percentage of the relevant literature, in our case via a stopping criterion (a maximum of 50 hits per search string and source). We counteracted this via five different data sources and one round of snowballing, which should have led to sufficient study diversity. Nonetheless, it is possible that the distributions for extractions could be different with a larger number of primary studies. We also need to be careful to generalize the review findings to industry settings. However, the more important limitation in this area is that we conducted a case study, a method that trades off external validity for result richness and depth. Our interview-based findings at Siemens AG are valuable, but not all of them may be generalizable to other companies. While Siemens is a global technology enterprise with diverse business units and projects in various domains, a certain degree of homogeneity between our units of analysis is to be expected. Parts of our case study findings may therefore not be transferable to other CPS companies, especially to smaller or less experienced ones. Even though we touched several CPS domains, our sample only contains seven systems that were described by nine CPS professionals. Additional studies with other companies are therefore necessary to broaden the understanding of microservices and DevOps for CPS, which we try to enable by sharing our study artifacts for replications.

6 | CONCLUSION

While challenges and practices regarding microservices and DevOps are well studied for enterprise applications, an overview of the CPS domain is missing in this regard. To address this, we formed an industry-academia collaboration, and conducted a rapid review (146 papers) followed by an interview-based case study with 9 different CPS experts at Siemens AG. We found that, with a few exceptions, our interviewees' perception of challenges and their applied practices were much closer to the enterprise application domain than to the CPS literature. Researchers and practitioners should take these findings into account when studying, developing, or modernizing CPS with microservices and DevOps. We especially see potential for adaptation mechanisms in existing microservices methods and tools to support the mentioned

challenges and peculiarities of CPS microservices. Additionally, more research is needed to see to what degree our case study findings at Siemens are generalizable to other CPS companies.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available in Adopting Microservices and DevOps in the CPS Domain at <https://doi.org/10.5281/zenodo.6536538>.

ACKNOWLEDGMENT

Open Access funding enabled and organized by Projekt DEAL.

ORCID

Jonas Fritzschn  <https://orcid.org/0000-0002-6121-2731>

Justus Bogner  <https://orcid.org/0000-0001-5788-0991>

Markus Haug  <https://orcid.org/0000-0001-9377-0677>

Ana Cristina Franco da Silva  <https://orcid.org/0000-0001-8549-350X>

Stefan Wagner  <https://orcid.org/0000-0002-5256-8429>

REFERENCES

1. Lasi H, Fettke P, Kemper HG, Feld T, Hoffmann M. Industry 4.0. business & information. *Syst Eng*. 2014;6(4):239-242. doi:10.1007/s12599-014-0334-4
2. Broy M, Schmidt A. Challenges in engineering cyber-physical systems. *IEEE Comput*. 2014;47(2):70-72. doi:10.1109/MC.2014.30
3. Ibarra-Junquera V, Gonzalez-Potes A, Paredes C, Martinez-Castro D, Nunez-Vizcaino RA. Component-based microservices for flexible and scalable automation of industrial bioprocesses. *IEEE Access*. 2021;9:1. doi:10.1109/ACCESS.2021.3072040
4. Waseem M, Liang P, Shahin M. A systematic mapping study on microservices architecture in DevOps. *J Syst Softw*. 2020;170:110798. doi:10.1016/j.jss.2020.110798
5. Banica L, Polychronidou P, Radulescu M, Stefan C. When IoT meets DevOps: fostering business opportunities. *Econom Balkan Eastern Eur Countr Chang World (Ebeec 2018)*. 2018;2018:250-264. doi:10.18502/kss.v3i10.3542
6. Knoche H, Hasselbring W. Drivers and Barriers for microservice adoption-a survey among professionals in Germany. *Enterpr Model Inf Syst Arch (EMISA)*. 2019;14(1):1-35. doi:10.18417/emisa.14.1
7. Baškarada S, Nguyen V, Koronios A. Architecting microservices: practical opportunities and challenges. *J Comput Inf Syst*. 2018;60:428-436. doi:10.1080/08874417.2018.1520056
8. Fritzschn J, Bogner J, Wagner S, Zimmermann A. Microservices migration in industry: intentions, strategies, and challenges. Proceedings of the 2019 IEEE International Conference on Software Maintenance and Evolution (ICSME); 2019; IEEE.
9. Taibi D, Lenarduzzi V, Pahl C. Processes, motivations, and issues for migrating to microservices architectures: an empirical investigation. *IEEE Cloud Comput*. 2017;4(5):22-32. doi:10.1109/MCC.2017.4250931
10. Törngren M, Sellgren U. Complexity challenges in development of cyber-physical systems. LNCS of Lecture Notes in Computer Science; Vol. 10760, 2018; Springer.
11. Jiang Y, Huang Z, Tsang DHK. Challenges and solutions in fog computing orchestration. *IEEE Netw*. 2018;32(3):122-129. doi:10.1109/MNET.2017.1700271
12. Li F, Gelbke L. Microservice architecture in industrial software delivery on edge devices. Proceedings of the 19th International Conference on Agile Software Development: Companion; 2018; ACM, New York.
13. Mizutani I, Ramanathan G, Mayer S. Integrating multi-disciplinary offline and online engineering in industrial cyber-physical systems through DevOps. Proceedings of the 11th International Conference on the Internet of Things (IoT'21); 2021; ACM.
14. Ciavotta M, Maso GD, Rovere D, Tsvetanov R, Menato S. Towards the digital factory: a microservices-based middleware for real-to-digital synchronization. *Microservices*. Springer International Publishing; 2020.
15. Gartner. Gartner 2020 hype cycle for supply chain. Online; 2020. <https://www.gartner.com/en/newsroom/press-releases/2020-09-09-gartner-2020-hype-cycle-for-supply-chain-strategy-shows-internet-of-things-is-two-to-five-years-away-from-transformational-impact>.
16. Combemale B, Wimmer M. Towards a model-based DevOps for cyber-physical systems. Proceedings of the LNCS of Lect Notes in Computer Science. Vol. 12055, 2020; Springer International Publishing.
17. Fowler M. Microservices resource guide. online; 2015. <http://martinfowler.com/microservices>.
18. Fritzschn J, Bogner J, Zimmermann A, Wagner S. From mornolith to microservices: a classification of refactoring approaches. Lecture Notes in Computer Science; 11350, 2019:128-141; Cham, Springer International Publishing.
19. Ponce F, Márquez G, Astudillo H. Migrating from monolithic architecture to microservices: a rapid review. Proceedings of 38th International Conference of the Chilean Computer Science Society (SCCC 2019); 2019; Chile, IEEE.
20. Törngren M, Grogan P. How to deal with the complexity of future cyber-physical systems? *Designs*. 2018;2(4):40. doi:10.3390/designs2040040

21. Fahmideh M, Abbasi AA, Behnaz A, Grundy J, Susilo W. Software engineering for Internet of Things. *IEEE Trans Softw Eng.* 2022;48:2857-2878. doi:10.1109/TSE.2021.3070692
22. Campeanu G. A mapping study on microservice architectures of Internet of Things and cloud computing solutions. Proceedings of the 7th Mediterranean Conference on Embedded Computing (MECO); June 2018:1-4; IEEE.
23. Joseph CT, Chandrasekaran K. Straddling the crevasse: a review of microservice software architecture foundations and recent advancements. *Softw Pract Exp.* 2019;49(10):1448-1484. doi:10.1002/spe.2729
24. Pereira IM, Carneiro T, Figueiredo E. A systematic review on the use of DevOps in internet of things software systems. *Proceedings of the 36th Annual ACM Symposium on Applied Computing*; 2021; ACM, New York, NY.
25. Hasselbring W, Henning S, Latte B, et al. Industrial DevOps. Proceedings of the IEEE International Conference on Software Architecture - Companion, ICSA-C 2019; March 2019:123-126; IEEE.
26. Taibi D, Lenarduzzi V, Pahl C. Continuous architecting with microservices and DevOps: a systematic mapping study. Proceedings of the Cloud Computing and Services Science. CLOSER 2018. Communications in Computer and Information Science; 2019:126-151; Springer.
27. Kitchenham B, Charters S. Guidelines for performing systematic literature reviews in software engineering. Technical report, School of Computer Science and Mathematics, Keele University; Keele, UK; 2007.
28. Petersen K, Feldt R, Mujtaba S, Mattsson M. Systematic mapping studies in software engineering. Proceedings of the 12th international conference on Evaluation and Assessment in Software Engineering EASE'08; 2008:68-77.
29. Cartaxo B, Pinto G, Soares S. Rapid reviews in software engineering. *Contemporary Empirical Methods in Software Engineering*. Springer International Publishing; 2020:357-384.
30. Wohlin C. Guidelines for snowballing in systematic literature studies and a replication in software engineering. *Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering - EASE'14*. ACM Press, New York; 2014:1-10.
31. Cruzes DS, Dyba T. Recommended steps for thematic synthesis in software engineering. Proceedings of the 2011 International Symposium on Empirical Software Engineering and Measurement; 2011:275-284; IEEE.
32. Runeson P, Höst M. Guidelines for conducting and reporting case study research in software engineering. *Empir Softw Eng.* 2009;14(2):131-164. doi:10.1007/s10664-008-9102-8
33. Seaman CB. Qualitative methods. *Guide to Advanced Empirical Software Engineering*. Springer; 2008:35-62.
34. Fritzsich J, Bogner J, Haug M, Wagner S, Zimmermann A. Towards an architecture-centric methodology for migrating to microservices. arXiv 2022. doi: 10.48550/ARXIV.2207.00507
35. Zimmermann T. Card-sorting: from text to themes. *Perspectives on Data Science for Software Engineering*. Elsevier; 2016:137-141.
36. Fowler M. Strangler application. Online; 2004. <https://www.martinfowler.com/bliki/StranglerApplication.html>.
37. Bogner J, Fritzsich J, Wagner S, Zimmermann A. Microservices in industry: insights into technologies, characteristics, and software quality. Proceedings of the International Conference on Software Architecture Companion (ICSA-C); 2019; IEEE.
38. Matias T, Correia FF, Fritzsich J, Bogner J, Ferreira HS, Restivo A. Determining microservice boundaries: a case study using static and dynamic software analysis. In: Jansen A, Malavolta I, Muccini H, Ozkaya I, Zimmermann O, eds. *Software Architecture ECSA 2020 Lecture Notes in Computer Science 12292 of Lecture Notes in Computer Science*. Springer; 2020:315-332.
39. Bogner J, Fritzsich J, Wagner S, Zimmermann A. Industry practices and challenges for the evolvability assurance of microservices. *Empir Softw Eng.* 2021;26(5):104. doi:10.1007/s10664-021-09999-9
40. Vale G, Correia FF, Guerra EM, Oliveira Rosa DT, Fritzsich J, Bogner J. Designing microservice systems using patterns: an empirical study on quality trade-offs. Proceedings of the 2022 IEEE 19th International Conference on Software Architecture (ICSA); 2022:69-79; IEEE.
41. Zhang H, Li S, Jia Z, Zhong C, Zhang C. Microservice architecture in reality: an industrial inquiry. Proceedings of the 2019 IEEE International Conference on Software Architecture (ICSA); 2019:51-60; IEEE.
42. Villari M, Fazio M, Dustdar S, Rana O, Ranjan R. Osmotic computing: a new paradigm for edge/cloud integration. *IEEE Cloud Comput.* 2016;3(6):76-83. doi:10.1109/MCC.2016.124
43. Wohlin C, Runeson P, Höst M, Ohlsson MC, Regnell B, Wesslén A. *Experimentation in Software Engineering*. Springer; 2012.

How to cite this article: Fritzsich J, Bogner J, Haug M, et al. Adopting microservices and DevOps in the cyber-physical systems domain: A rapid review and case study. *Softw: Pract Exper.* 2023;53(3):790-810. doi: 10.1002/spe.3169