

Tracing the Footsteps of Technical Debt in Microservices: A Preliminary Case Study

Roberto Verdecchia^{*[0000-0001-9206-6637]}, Kevin Maggi^{*},
Leonardo Scommegna^[0000-0002-7293-0210], and Enrico Vicario^[0000-0002-4983-4386]

Department of Information Engineering, University of Florence, Italy
roberto.verdecchia@unifi.it, kevin.maggi@stud.unifi.it,
leonardo.scommegna@unifi.it, enrico.vicario@unifi.it

Abstract. *Background:* Albeit the growing academic and industrial interest in microservice architectures and technical debt, to date no study aimed to investigate the evolution characteristics of technical debt in software-intensive systems based on such architecture. *Aims:* The goal of this study is to understand how technical debt evolves in microservice-based software-intensive systems, in terms of (i) evolution trends, and (ii) relation between technical debt and number of microservices. *Method:* We adopt a case study based on an application comprising 13 microservices, 977 commits, and 38k lines of code. The research method is based on repository mining and automated source code analysis complemented *via* manual code inspection. *Results:* While long periods of development without TD increase are observed, TD overall increases in time. TD variations can happen regardless of the number of microservices and development activity considered in a commit. TD and number of microservices are strongly correlated, albeit not always. Adding (or removing) a microservice has a similar impact on TD regardless of the number of microservices already present in a software-intensive system. *Conclusions:* Adherence to microservice architecture principles might keep technical debt compartmentalized within microservices and hence more manageable. Developers should pay keen attention to the technical debt they may introduce, regardless of the number of microservice they touch with a commit and the development activity they carry out. Keeping technical debt constant during the evolution of a microservice-based architecture is possible, but the growth of technical debt while a software-intensive systems becomes bigger and more complex might be inevitable.

Keywords: Technical Debt · Microservices · Software Evolution

1 Introduction

As companies seek to take advantage of their many benefits, microservice-based architectures are becoming more and more adopted. As often referenced, the microservice architecture style offers several advantages, such as scalability, flexibility, and the ability to develop and deploy individual components independently [11]. Albeit the many benefits microservice-based systems offer, the architectural style also comes with its own

* The first two authors contributed equally to this work.

set of challenges, including increased complexity, the need for effective management of eventual consistency, and additional effort required for integration and system testing.

In this context, we can intuitively conjecture that, in order to cope with the increased complexity and potential loss of the bigger architectural picture, developers may tend to adopt suboptimal implementation expedients. While providing temporary benefits, such expedients may tend to make future development harder or even impossible. This concept of software quality issues related to temporary expedients is commonly referred to as technical debt (TD) [5].

TD is one of the paramount factors in software development practice. If left unmanaged, TD can lead, among other consequences, to lower development speed, raise of a high number of bugs, or even completely crystallized architectures [32]. TD has been widely covered in academic literature [3, 25, 33] and is increasing in research popularity. Similarly, albeit the adoption of microservice architectural style could be considered as a relatively new phenomenon, its widespread adoption recently drew a considerable academic interest [12, 16, 30].

Surprisingly, while both TD and microservices could be considered as popular topics in current academic literature, there has been relatively little focus on the relationship between TD and microservice architectures. To date, few studies have considered how TD evolves in microservice systems and, to the best of our knowledge, none have quantitatively studied in depth the characteristics of such evolution.

With this research, through a preliminary case study on an open-source project comprising 12 microservices (see Section 3.4), we make a first step towards understanding how TD changes as microservice-based systems evolve. Our goal is to pave the way for future empirical studies that investigate the evolution of this relationship. By understanding how TD evolves in microservice architectures, and gaining insights into the characteristics of such evolution, we might be able to shed light on how TD can be effectively managed in microservice architectures, with the end goal of better supporting the long-term success of software-intensive systems based on such architectural style.

The main contributions of this research are (i) a quantitative case study reporting TD measurements through the evolution of a microservice-based software system, (ii) a thorough statistical analysis complemented by a manual code inspection and discussion of the results, and (iii) a replication package containing the entirety of the raw, intermediate, and final data, analysis traces, and code used for this study.

2 Related Work

By considering the academic literature that focuses on TD in microservice-based systems we can identify, to the best of our knowledge, only a handful of studies.

The research of Lenarduzzi *et al.* [23], where the effects of migrating from a monolithic to a microservice architecture can have on TD are investigated, might potentially be the most similar to this work. As our study, the research presents a case study based on repository mining and static code analysis. In contrast to such study however, we (i) do not focus on the effects of migrating from monolithic to microservice architecture, (ii) aim to study various characteristics affecting TD (*e.g.*, number of microservices), and (iii) consider as case study a software-intensive system which comprises 13 microservices instead of the 5 studied by Lenarduzzi *et al.* [23].

By inspecting the other related literature, TD in microservices appears to be investigated primarily from a qualitative point of view. In a study by Toledo *et al.* [31], a multiple case study based on 25 interviews investigating architectural TD (ATD) in microservices is reported. The results of the investigation identified ATD issues, their negative impact, and the common solutions used to repay each debt type. Differently to such study, we focus on code TD [18,25], utilize a quantitative rather than qualitative research method, and focus on a case study. In a similar work by Toledo *et al.* [15], through a qualitative analysis of documents and interviews, ATD in the communication layer of microservice-based architecture is investigated. The main contribution of the paper is a list of debt types specific to the communication layer of a microservice-based architecture, as well as their associated negative impact, and solutions to repay the debt types. Regarding the differences w.r.t. our work, the same considerations previously elicited for the other study of Toledo *et al.* [31] apply.

Bogner *et al.* [10] adopted 17 semi-structured interviews to study how the sustainable evolution of 14 microservice-based systems was ensured. Albeit from the results ATD emerged as a relevant issue, differently from our study, the work of Bogner *et al.* does not explicitly focus on TD. As additional difference w.r.t. our work, while tool-based DevOps processes were often mentioned as a mean to assure evolvability, the study is based on a qualitative rather than quantitative empirical research method. Bogner *et al.*, in a different study [9], surveyed 60 software professionals *via* an online questionnaire to learn how technical debt can be limited through maintainability assurance. Results indicated that using explicit and systematic techniques benefits software maintainability. As for the previous studies, also this work by Bogner *et al.* [9] adopts a qualitative rather than quantitative research approach. In addition, albeit both the study of Bogner *et al.* [9] and this work consider TD in microservices, the primary focus of Bogner *et al.* is on studying maintainability assurance techniques, while the one of this work is on TD evolution in microservice-based software-intensive systems.

Related to the concept of TD in microservice-based systems, Pagazzini *et al.* [28] present the extension of the tool Arcan [19] to detect microservice smells. As main difference with this work, the Arcan extension focuses on architectural smells rather than focusing explicitly on TD, and does not carry out a case study on TD evolution.

A more systematic literature review on TD in microservices w.r.t. this related work section is conducted by Villa *et al.* [35]. Based on the analysis of 12 primary studies, Villa *et al.* corroborate the intuition grounding this study, namely the absence of qualitative studies focusing on the evolution of TD in microservice-based systems. From the results of Villa *et al.*, ATD and code debt result to be the most frequently reported debt types for microservices. Such finding, which reflects the general trend observed for TD in developer discussions [2,22], provides further support to the focus of this work on the evolution of code TD in microservice-based systems.

3 Study Design and Execution

In this section, we document the research design and execution of the study, in terms of research goal (Section 3.1), research questions (Section 3.2), and research process (Section 3.3).

3.1 Research Goal

The goal of this research is to conduct a preliminary investigation into the evolution of TD in software-intensive systems utilizing a microservice architecture. By using the Goal-Question-Metric framework of Basili *et al.* [8], our goal is:

Analyze software evolution

For the purpose of studying trends and characteristics

With respect to technical debt

From the viewpoint of software engineering researchers

In the context of microservice-based software-intensive systems.

In this study, we opt to focus on code TD [25], rather than other TD types (*e.g.*, ATD) guided to multiple factors, namely (i) code TD is one of the most frequent TD types appearing in microservice-based systems [35], (ii) in contrast to the other TD types, code TD is supported by a vast range of consolidated off the shelf tools, which are vastly used both in academic research and industrial practice [6], (iii) the focus on code TD allows for the natural extension of this preliminary case study to future heterogeneous case studies.

3.2 Research Questions

Based on the goal of our study, we can derive the main research question (RQ) and sub-research questions which guide our research.

The main RQ on which this study is based can be formulated as follows:

RQ: *How does code technical debt evolve in a microservice-based software-intensive system?*

With this main RQ, which encompasses the overall goal of the study, we broadly express our intent to study the evolution of code TD in microservice-based systems. To be more systematic, we decompose our main RQ into two sub-RQs, each one considering a different facet of TD evolution in microservice-based software intensive systems.

RQ₁: *What is the evolution trend of TD in a microservice-based software-intensive system?*

With *RQ₁*, we aim to understand the overall evolution trend of TD in microservice-based systems, *e.g.*, if TD is constant through the evolution of a microservice-based software-centric system, if TD showcases a growing trend in time, or if the system is characterized by a seasonal TD trend (*e.g.*, if developers are more prone to incur in TD before/after seasonal holidays).

RQ₂: *Is there a relation between TD evolution and number of microservices?*

With *RQ₂*, we aim to understand if a relation exists between the evolution of TD and the number of microservices composing a software-intensive system. As example, we could conjecture that, due to suboptimal practices, as the number of microservices grows, TD grows at a higher rate (*e.g.*, TD is in superlinear or even exponential relation with the number of microservices).

3.3 Research Process

An overview of the process followed in this study is depicted in Figure 1, and is further documented below.

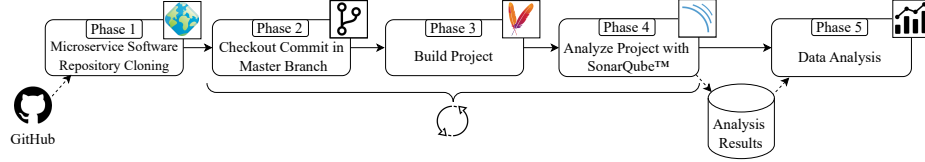


Fig. 1. Research process overview.

The research process consists of five phases, from cloning the case study software repository containing a microservice-based software project, to the static analysis of its source code, and the concluding statistical and manual analysis of the collected data. Each research phase is described in detail in the following.

3.4 Phase 1: Microservice Software Repository Cloning

The first step of our research process consists in cloning a repository containing a microservice-based software project. For this preliminary case study, we select the software repository **Cloud Native GeoServer**.¹ The project is a microservice implementation of GeoServer, an open source server for sharing geospatial data. **Cloud Native Geoserver** splits the original GeoServer geospatial services and API offerings into individually deployable components of a microservices based architecture.

The case study is selected starting from the manually validated list of microservice-based open-source projects hosted on GitHub elicited by Baresi *et al.* [7]. The project is selected from the list by considering as selection criteria (SC_1) the real use of the application, (SC_2) the number of times the repository is forked and starred, (SC_3) the number of repository commits, and (SC_4) the number of microservices the project comprises. We use SC_1 to exclude as potential case study a toy project and demo. SC_2 instead provides us assurance on the quality and popularity of the repository. Finally, SC_3 and SC_4 guarantee us that the project is representative of a long-lived, complex, software application based on a microservice-based architecture. **Cloud Native GeoServer** to date is forked a total of 52 times and starred 176 times. The repository currently counts a total of 985 commits, comprises 13 distinct microservices, and is composed of 38k lines of code (NLOC), and is contributed to by 10 developers.²

3.5 Phase 2: Checkout Commit in Main Branch

The second phase of our research consists in checking out in temporal order the commits of the selected repository. For this process, we opt to consider the commits

¹ <http://geoserver.org/geoserver-cloud>. Accessed 4th July 2023.

² <https://github.com/geoserver/geoserver-cloud>. Accessed 4th July 2023.

present in the main branch of the selected repository. While we are aware of the potential pitfalls implied by considering exclusively the main branch during repository mining processes [21], we deem analyzing also other branches as out of scope for this preliminary investigation. Related threats to validity are discussed in Section 5.

3.6 Phase 3: Build Project

After checking out a commit, the project is built by using the build automation tool used by the case study software-intensive system, namely Maven³. This step is a prerequisite for the analysis of the project *via* SonarQube (Phase 4, see Section 3.7), as the tool requires the compiled code of the software project in order to analyze it.

During this step, 7 out of 985 builds result to fail (0.7% of all builds). Upon inspection, we identify the failure to be caused by issues related to the Project Object Model (POM) of the Maven build. Rather than using a subjective heuristic to fix the issue, we opt to discard the commits associated to these failing builds. A single commit results instead to be characterized by an erroneous date in the versioning system. In order to avoid to independently estimate the correct commit date *via* some *ad hoc* heuristic, we opt to discard such commit from our analysis. Given the relatively low number of commits skipped due to build failures or dating issues (8/985 in total), we do not deem this factor to noticeably influence our results. Further considerations are reported in the threats to validity section (Section 5). For scrutiny and traceability purposes, the metadata of the 8 commits omitted from our analysis are documented in the replication package of this study.

To iterate and avoid possible confusion, due to the failing builds and an ill-dated commit, 977 out of the 985 total commits are considered for analysis of this study.

3.7 Phase 4: Analyze Project with SonarQube

After obtaining the compiled code of the project, the code is analyzed by utilizing the SonarQube tool. All commits are analyzed by using SonarQube version 9.9 LTS with SonarScanner for Maven version 3.9.1. During this process, in addition to the SQUALE metric measuring TD [24], other metrics and metadata of the project is collected, *e.g.*, the project size in terms of NCLOC, number of files, cognitive complexity, and committer name. To measure TD, the standard out of the box SonarQube rules configuration is used, in order to avoid subjective tempering of the tool settings.

The number of microservices appearing in each commit version is instead collected by following the method first introduced by Baresi *et al.* [7]. Such method relies on the analysis of Docker Compose files, in order to identify *via* parsing the microservices composing a software-intensive system.

Phases 2-4 are repeated for each of the 977 commits of the software project considered for this case study.

³ <https://maven.apache.org>. Accessed 4th July 2023

3.8 Phase 5: Data Analysis

As final step of the research process, the data collected through Steps 1-4 is analyzed to answer our RQs.

To answer RQ_1 , we decompose the TD evolution trend into its seasonal, trend, and irregular components [17] by utilizing on the STL algorithm [14]. We adopt the STL algorithm as it does not assume a time series distribution, it was successfully used in previous software engineering studies [4, 26], and an open-source implementation is available as an R library.⁴ The resulting trend is then inspected qualitatively by graphical means. To gain further insights into the “TD hotspots”, *i.e.*, commits showcasing the most outlier values in TD measurements, the content of the outlier commits are manually scrutinized. To identify outlier values, we leverage the STL decomposition, by first removing any seasonality and trend in the TD time series, and subsequently selecting the 10 most anomalous outliers identified in the STL irregular component series for manual scrutiny.

To answer RQ_2 , we first study the potential correlation between the number of microservices and TD time series. Afterwards, we analyze the potential correlation between the derivatives of such series, to understand the relation between the growth speed of TD w.r.t. microservice number. For both cases, we test the correlation by using the Multivariate Granger causality analysis [20]. To calculate the optimal lag order for the Granger analysis, we adopt the Akaike Information Criterion [1]. As the Granger test assumes the time series to be stationary, we test such assumption *via* the Augmented Dickey-Fuller test [13]. In case the time series result to be non-stationary, we make them stationary by differencing the data, *i.e.*, by subtracting the value of each observation from the value of the previous observation in the time series.

4 Results

In this section we report and discuss the data gathered to answer our RQs. Specifically, in the next section (Section 4.1) we consider the results of RQ_1 , while in Section 4.2 we take into account the results of RQ_2 .

4.1 Results RQ_1 : Evolution of TD in a microservice-based software-intensive system

As described in Section 3.8, in order to study the TD evolution of our case study, namely the `Cloud Native GeoServer` application, we consider three different components, namely the TD evolution trend, seasonality, and irregularities. An overview of such decomposition is depicted in Figure 2.

As we can observe from the leftmost diagram of Figure 2, the TD evolution showcases an overall growing trend. Interestingly, two outstanding jumps, *i.e.*, sudden increases in TD values, can be noticed in the plot. By comparing the trend figure

⁴ <https://stat.ethz.ch/R-manual/R-devel/library/stats/html/stl.html>. Accessed 4th July 2023.

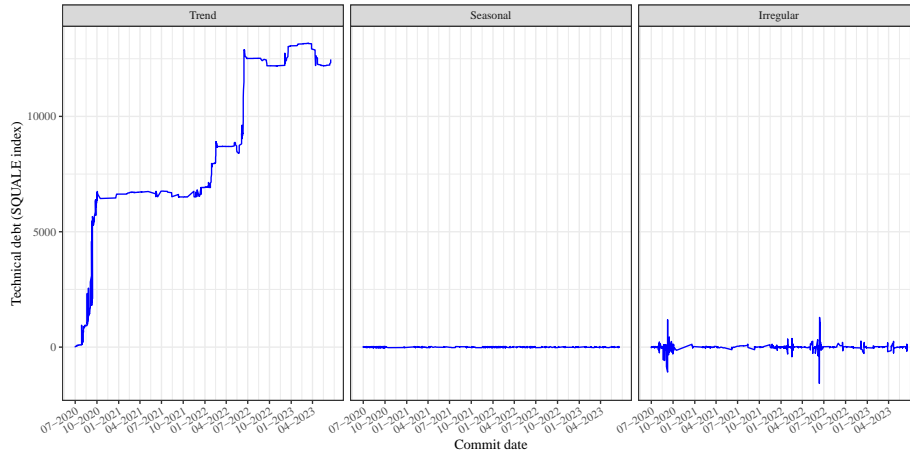


Fig. 2. Decomposition of the **Cloud Native GeoServer** application technical debt evolution *via* the STL algorithm.

with the one reporting the irregularities in TD evolution (rightmost diagram of Figure 2), we note that such outliers are captured by the STL algorithm decomposition. The commits corresponding to such jumps are further analyzed in the second data analysis carried out to answer RQ_1 , namely the manual scrutiny of the “TD hotspots” (see Section 3.8). Overall, as could be expected, TD tends to naturally increase during time as the application becomes bigger and more complex.

The TD evolution trend also presents noticeable plateaus, *i.e.*, periods where the TD values remain approximately stable. The first plateau, starting from October 2020, lasts approximately one year and three months of development. By inspecting the commit dates, we note that none of the plateaus reported in the trend of Figure 2 is due to periods of development inactivity. Therefore, we conjecture that the development periods associated to the plateaus correspond to development periods where deliberate efforts might be made to prevent a TD increase.

By considering the seasonality of the time series (center diagram of Figure 2), we can intuitively observe that no seasonal behavior is present in the TD evolution of the **Cloud Native GeoServer** application. This implies that TD is not more likely to be introduced during a certain period of the year.

As second data analysis process carried out to answer RQ_1 , we manually inspect the potential “TD hotspots” (see Section 3.8). The most noticeable “TD hotspot” corresponds to a sudden increase in TD values recorded on June 2022 (see Figure 2, rightmost plot). From manual scrutiny, this sudden TD variation results to be due to the upgrade of the JUnit testing framework⁵. The commit also includes the cross-microservice refactoring of test files according to the upgrade.

Other seven “hotspots”, which are not graphically appreciable from the TD evolution irregularities depicted in the rightmost plot of Figure 2, happen on the

⁵ <https://junit.org/junit5/>. Accessed 5th July 2023.

same day as the JUnit upgrade commit. Upon manual inspection, we note that the commits corresponding to these additional seven hotspots involve many microservices of the **Cloud Native GeoServer** application. The commits result to either focus on (i) further refactoring of testing artifacts, (ii) bug fixing, (iii) implementing logging mechanisms, and (iv) introducing automation processes.

The sudden increase of TD values in October 2020 results instead from manual inspection caused by the addition of 33 new files in a microservice. The commit involves the extension of the **Cloud Native GeoServer** features *via* the binding to a new JSON parser.

The last of the 10 “TD hotspots” considered for manual analysis instead corresponds to a sudden increase of TD values happening in September 2020. In this case, the TD increase results to be due to a refactoring activity carried out across 70 files, which involved a considerable number of NCLOC (1.5k NCLOC additions, and 589 NCLOC deletions).

As conclusion of the manual scrutiny, we highlight that both working on a single microservice, or multiple ones at the same time, can drastically influence TD. Considerable TD variations can happen independently of the developer activity conducted, *e.g.*, a framework upgrade can have an unforeseen cascading impact on TD, or a refactoring activity could lead to a considerable TD increase.

RQ₁ answer (*TD Evolution in a Microservice-based Architecture*)

TD displays an overall increasing trend in time, albeit long periods of continued development without TD increase are noticeable. Considerable TD variations can happen by working on one or multiple microservices, and may occur regardless of the development activity conducted.

4.2 Results *RQ₂*: Relation between TD evolution and number of microservices

In order to study the potential correlation between TD and number of microservices, we start by graphically inspecting the time series of the two metrics. As can be seen in Figure 3, both TD and microservices seem to display an overall similar growth trend. However, a correlation does not appear always to be present in all commits. As example, by considering Figure 3, we can observe that the removal of one microservice in April 2021 did not correspond to any noticeable TD decrease. In some cases, *e.g.*, April 2023, the addition of a microservice is even associated with a decrease in TD (*i.e.*, microservice number and TD are inversely proportional). This would imply that, while number of microservices could display a strong correlation of directly proportional nature, this might not always be the case.

In order to gain statistical insight into the relation, we follow the analysis process presented in Section 3.8. From the results of the Granger causality test we confidently reject the null hypothesis and conclude that the evolution of number of microservices and TD are strongly correlated ($p\text{-value} < 4.593^{-6}$). This implies that, albeit seldom irregularities, a growth (or decrease) of microservice number corresponds to a similar change in TD. As additional remark, this indicates that the number of microservices could be used to predict TD values.

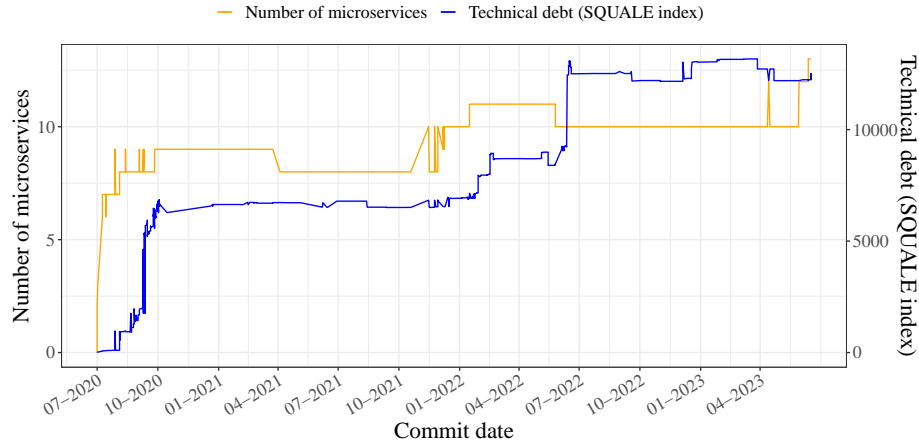


Fig. 3. Overview of the evolution of technical debt and number of microservices of the Cloud Native GeoServer application.

As further analysis conducted to answer RQ_2 , we study the derivatives of the two time series. This allows us to understand if the two series grow at similar rates, or if a growth in the microservice series correspond to the growth at a higher rate in the TD series. Intuitively, we could expect that, as the number of microservices increases, the software-intensive system becomes more complex, and hence TD grows at a higher rate as the system becomes bigger. From the Granger test results however, we understand that this is not the case. In fact, we observe that also the derivatives of the time series are strongly correlated ($p\text{-value} < 7.896^{-6}$), *i.e.*, we can discard with statistical confidence the null hypothesis. This implies that the impact of adding (or removing) a microservice on TD is similar regardless of the number of microservices already present in a software-intensive system.

As subjective interpretation of this latter result, we can conjecture that this pattern indicates an appropriate adherence to the microservice architectural principles, through which microservices are developed independently by following a loosely coupled and highly cohesive architecture.

RQ₂ answer (Relation between TD and number of microservices)

TD and number of microservices are strongly correlated, albeit in seldom cases such relation does not persist. The impact of adding (or removing) a microservice on TD is similar, regardless of the number of microservices already present in a software-intensive system.

5 Threats to Validity

The presented results have to be interpreted in light of potential threats to validity. By following the categorization of Runeson *et al.* [29], we consider four aspects.

5.1 Construct Validity

To answer our RQs, we measured code TD by adopting the SQUALE index, a metric widely used in the literature to study TD [6, 23, 34]. The number of microservices was measured by utilizing the heuristic first introduced by Baresi [7] (see also Section 3.7). The use of the heuristic might have marginally affected our results, as it relies on the analysis of the Docker Compose file. Therefore, a service could be identified at its insertion in the Docker Compose file, which does not necessarily imply the start of its actual development. At most, this threat could have introduced a lag in the TD timeseries w.r.t. the number of microservice one (corresponding to the time elapsed from the insertion of a microservice in the Docker Compose file, and the start of its actual development). As the potential effects of this marginal threat were not noticeable in our data analysis, we do not deem the threat considerably influenced our results. The threat would at most imply a stronger correlation between microservice number and TD than the one observed. Regarding the focus on the main development branch of `Cloud Native Geoserver` (see Section 3.5), we note that the application possesses other two branches, which are characterized by 2 and 48 commits ahead, and 256 and 28 commits behind the master branch respectively. Given the low number of commits in such branches w.r.t. the master branch (2/985 and 48/985), we do not deem that this research design choice could have drastically influenced our results.

5.2 Internal Validity

To avoid potential confounding factors, we (i) discarded all failing builds, and a commit associated with an incorrect date, (ii) manually scrutinized a set of commits presenting anomalous TD values, (iii) conducted a rigorous statistical analysis on the collected data.

5.3 External validity

As any case study, we do not claim the complete generalizability of the obtained results. While comparable results might be observed in software-intensive systems of similar development context and characteristics as `Cloud Native GeoServer`, this could also not be the case. To mitigate this threat in the future, we plan to contact the developers of the `Cloud Native GeoServer` application to qualitatively complete the results of primarily quantitative nature presented in this study. This would allow us to gain a clearer picture of the phenomenon under study, and better understand to which extent the results can apply to other software-intensive system and contexts.

5.4 Reliability and Replication Package

If and to what extent the results of the study can be independently reproduced by other researchers. With exception of the manual scrutiny conducted to analyze the commits presenting the most anomalous TD values, the present study is completely based on the execution of mining and data analysis scripts. As we make all data, scripts, and

settings available in a companion replication package⁶, and given the almost purely quantitative nature of the study, we deem the reliability of the study as very high.

6 Conclusion and Future Work

In this study, we present a preliminary case study investigating the evolution of technical debt in microservice architectures. The investigation utilizes as case study the application `Cloud Native GeoServer`, which comprises a total of 13 distinct microservices, 977 commits, and 38k NCLOC. The study is primarily based on repository mining and source code analysis. The results show that TD evolution displays a growing trend interspersed with moments of TD stability. TD variation are independent of the number of microservices and development activities considered in a commit. TD and number of microservices are correlated, and adding or removing a microservice has the same impact on TD regardless the number of microservices already present.

As concluding takeaways, we note that adhering to microservice architecture principles might keep technical debt compartmentalized within microservices, and therefore make TD more manageable w.r.t. other types of architectures (*e.g.*, monolithic ones). It is crucial for developers to remain aware of the potential TD they may incur in, irrespective of the quantity of microservices they modify or the nature of the development activity they undertake. An intuitively trivial change, such as the upgrade of a testing framework, could have a massive cascading effect on the TD of a microservice-based software-intensive system. While it is feasible to maintain a consistent level of TD during the evolution of a microservice-based application, an increase in technical debt may be inevitable as the software-intensive system grows in size and complexity.

The presented case study has to be considered as preliminary in nature. Several facets which could provide us more information on the investigated phenomenon are not considered in the current version of the study. As future work, we plan to (i) study the individual contribution of each microservice to the TD measured at system level, (ii) conduct a more in-depth analysis of “TD hotspots”, (iii) utilize dedicated tools to measure other types of TD, *e.g.*, ATD *via* the ATDx tool [27], (iv) interview the developers of the software-intensive system to gain further insights on trends and “TD hotspots”, and (v) extending the research to a multiple-case study.

Acknowledgments

This work was partially supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on “Telecommunications of the Future” (PE0000001 - program “RESTART”).

Roberto Verdecchia would like to thank Curzio Checcucci for his lighthearted yet insightful feedback on the data analysis process.

⁶ **Replication package of this study:** <https://github.com/STLab-UniFI/QUALIFIER-2023-TD-microservices-rep-pkg>. Accessed 6th July 2023.

References

1. Akaike, H.: Information theory and an extension of the maximum likelihood principle. In: Selected papers of hirotugu akaike, pp. 199–213. Springer (1998)
2. Aldrich Edbert, J., Jannat Oishwee, S., Karmakar, S., Codabux, Z., Verdecchia, R.: Exploring technical debt in security questions on stack overflow. International Symposium on Empirical Software Engineering and Measurement (2023)
3. Alves, N.S., Mendes, T.S., De Mendonça, M.G., Spínola, R.O., Shull, F., Seaman, C.: Identification and management of technical debt: A systematic mapping study. *Information and Software Technology* **70**, 100–121 (2016)
4. Atchison, A., Berardi, C., Best, N., Stevens, E., Linstead, E.: A time series analysis of travistorrent builds: to everything there is a season. In: 2017 IEEE/ACM 14th International Conference on Mining Software Repositories. pp. 463–466. IEEE (2017)
5. Avgeriou, P., Kruchten, P., Ozkaya, I., Seaman, C.: Managing technical debt in software engineering (dagstuhl seminar 16162). *Dagstuhl Reports* **6** (01 2016)
6. Avgeriou, P.C., Taibi, D., Ampatzoglou, A., Fontana, F.A., Besker, T., Chatzigeorgiou, A., Lenarduzzi, V., Martini, A., Moschou, A., Pigazzini, I., et al.: An overview and comparison of technical debt measurement tools. *Ieee software* **38**(3), 61–71 (2020)
7. Baresi, L., Quattrocchi, G., Tamburri, D.A.: Microservice architecture practices and experience: a focused look on docker configuration files. arXiv preprint arXiv:2212.03107 (2022)
8. Basili, V.R., Caldiera, G., Rombach, D.: The Goal Question Metric Approach. In: *Encyclopedia of Software Engineering*, pp. 528–532. Wiley (1994)
9. Bogner, J., Fritzsich, J., Wagner, S., Zimmermann, A.: Limiting technical debt with maintainability assurance: An industry survey on used techniques and differences with service-and microservice-based systems. In: *Proceedings of the 2018 International Conference on Technical Debt*. pp. 125–133 (2018)
10. Bogner, J., Fritzsich, J., Wagner, S., Zimmermann, A.: Assuring the evolvability of microservices: insights into industry practices and challenges. In: 2019 IEEE International Conference on Software Maintenance and Evolution. pp. 546–556. IEEE (2019)
11. Bogner, J., Fritzsich, J., Wagner, S., Zimmermann, A.: Microservices in industry: insights into technologies, characteristics, and software quality. In: 2019 IEEE international conference on software architecture companion. pp. 187–195. IEEE (2019)
12. Bogner, J., Wagner, S., Zimmermann, A.: Automatically measuring the maintainability of service-and microservice-based systems: a literature review. In: *Proceedings of the 27th International Workshop on Software Measurement and 12th International Conference on Software Process and Product Measurement*. pp. 107–115 (2017)
13. Cheung, Y.W., Lai, K.S.: Lag order and critical values of the augmented dickey–fuller test. *Journal of Business & Economic Statistics* **13**(3), 277–280 (1995)
14. Cleveland, R.B., Cleveland, W.S., McRae, J.E., Terpenning, I.: Stl: A seasonal-trend decomposition. *J. Off. Stat* **6**(1), 3–73 (1990)
15. De Toledo, S.S., Martini, A., Przybyszewska, A., Sjøberg, D.I.: Architectural technical debt in microservices: a case study in a large company. In: 2019 IEEE/ACM International Conference on Technical Debt. pp. 78–87. IEEE (2019)
16. Di Francesco, P., Lago, P., Malavolta, I.: Architecting with microservices: A systematic mapping study. *Journal of Systems and Software* **150**, 77–97 (2019)
17. Dickey, D.A., Fuller, W.A.: Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American statistical association* **74**(366a), 427–431 (1979)
18. d’Aragona, D.A., Pecorelli, F., Baldassarre, M.T., Taibi, D., Lenarduzzi, V.: Technical debt diffuseness in the apache ecosystem: A differentiated replication. In: 2023 IEEE

- International Conference on Software Analysis, Evolution and Reengineering (SANER). pp. 825–833. IEEE (2023)
19. Fontana, F.A., Pigazzini, I., Roveda, R., Tamburri, D., Zanoni, M., Di Nitto, E.: Arcan: A tool for architectural smells detection. In: 2017 IEEE International Conference on Software Architecture Workshops. pp. 282–285. IEEE (2017)
 20. Granger, C.W.: Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society* pp. 424–438 (1969)
 21. Kovalenko, V., Palomba, F., Bacchelli, A.: Mining file histories: Should we consider branches? In: Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering. pp. 202–213 (2018)
 22. Kozanidis, N., Verdecchia, R., Guzmán, E.: Asking about technical debt: Characteristics and automatic identification of technical debt questions on stack overflow. In: Proceedings of the 16th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement. pp. 45–56 (2022)
 23. Lenarduzzi, V., Lomio, F., Saarimäki, N., Taibi, D.: Does migrating a monolithic system to microservices decrease the technical debt? *Journal of Systems and Software* **169**, 110710 (2020)
 24. Letouzey, J.L.: The sqale method for evaluating technical debt. In: 2012 Third International Workshop on Managing Technical Debt. pp. 31–36. IEEE (2012)
 25. Li, Z., Avgeriou, P., Liang, P.: A systematic mapping study on technical debt and its management. *Journal of Systems and Software* **101**, 193–220 (2015)
 26. Malavolta, I., Verdecchia, R., Filipovic, B., Bruntink, M., Lago, P.: How maintainability issues of android apps evolve. In: 2018 IEEE international conference on software maintenance and evolution. pp. 334–344. IEEE (2018)
 27. Ospina, S., Verdecchia, R., Malavolta, I., Lago, P.: Atdx: A tool for providing a data-driven overview of architectural technical debt in software-intensive systems. In: European Conference on Software Architecture (2021)
 28. Pigazzini, I., Fontana, F.A., Lenarduzzi, V., Taibi, D.: Towards microservice smells detection. In: International Conference on Technical Debt. p. 6. ACM, Seoul, Republic of Korea (2020). <https://doi.org/10.1145/3387906.3388625>
 29. Runeson, P., Höst, M.: Guidelines for conducting and reporting case study research in software engineering. *Empirical software engineering* **14**, 131–164 (2009)
 30. Soldani, J., Tamburri, D.A., Van Den Heuvel, W.J.: The pains and gains of microservices: A systematic grey literature review. *Journal of Systems and Software* **146**, 215–232 (2018)
 31. de Toledo, S.S., Martini, A., Sjøberg, D.I.: Identifying architectural technical debt, principal, and interest in microservices: A multiple-case study. *Journal of Systems and Software* **177**, 110968 (2021)
 32. Verdecchia, R., Kruchten, P., Lago, P., Malavolta, I.: Building and evaluating a theory of architectural technical debt in software-intensive systems. *Journal of Systems and Software* **176**, 110925 (2021)
 33. Verdecchia, R., Malavolta, I., Lago, P.: Architectural technical debt identification: The research landscape. In: Proceedings of the 2018 International Conference on Technical Debt. pp. 11–20 (2018)
 34. Verdecchia, R., Malavolta, I., Lago, P., Ozkaya, I.: Empirical evaluation of an architectural technical debt index in the context of the apache and onap ecosystems. *PeerJ Computer Science* **8**, e833 (2022)
 35. Villa, A., Ocharan-Hernandez, J.O., Perez-Arriaga, J.C., Limon, X.: A systematic mapping study on technical debt in microservices. In: 2022 10th International Conference in Software Engineering Research and Innovation. IEEE (2022)