

A Decomposition and Metric-Based Evaluation Framework for Microservices

Davide Taibi and Kari Systä*

TASE - Tampere Software Engineering Research Group,
Tampere University, Tampere, Finland
{davide.taibi,kari.systa}@tuni.fi

Abstract. Migrating from monolithic systems into microservice is a very complex task. Companies are commonly decomposing the monolithic system manually, analyzing dependencies of the monolith and then assessing different decomposition options. The goal of our work is two-folded: 1) we provide a microservice measurement framework to objectively evaluate and compare the quality of microservices-based systems; 2) we propose a decomposition system based on business process mining. The microservice measurement framework can be applied independently from the decomposition process adopted, but is also useful to continuously evaluate the architectural evolution of a system. Results show that the decomposition framework helps companies to easily identify the different decomposition options. The measurement framework can help to decrease the subjectivity of the decision between different decomposition options and to evaluate architectural erosion in existing systems.

Keywords: Microservices · Cloud-native · Microservice slicing · Microservice decomposition · Microservice migration

1 Introduction

Software evolves through its life-time, and often a large part the effort and costs is spent on software maintenance [14]. Furthermore, the incremental development practices in modern software development makes the nature of all the development processes to resemble a maintenance one [18]. A major obstacle to efficient maintenance is the tight coupling between the internal components of the software. In monolithic systems, most changes require modifications to several parts of the systems, and often size and complexity of the modification is hard to estimate in advance.

One approach to tackle the maintenance problem is to decompose the system into small and independent modules [20, 25]. Often, at the same time, companies want to utilize benefits of service-oriented architectures and even *microservices*, such as independent development, scaling and deployment [31].

Microservices are an adaptation service-oriented architecture but focuses on of relatively small and independently deployed services, with a single and clearly

Author Version. Please cite as:

Taibi D., Systä K. (2020) A Decomposition and Metric-Based Evaluation Framework for Microservices. CLOSER 2019. Communications in Computer and Information Science, vol 1218. Springer, Cham. https://doi.org/10.1007/978-3-030-49432-2_7

defined purpose [7]. The independent development and deployment bring several advantages. Different microservices can be developed in different programming languages, they can scale independently from other services, and each microservice can be deployed on the most suitable the hardware. Moreover, small services are easier to maintain and the split to independent responsibilities increases fault-tolerant since a failure of one service will not break the whole system. From the architectural perspective a well-designed microservice encapsulates its data and design choices. Thus, the internal logic of a microservice can be changed without affecting the external interface. This reduces the need for interaction between the teams [29,30].

However, decomposing a monolithic system into a set of independent microservices is a very difficult and complex tasks [31,35]. Decomposition of a system into separately maintained services is difficult as such, but microservice architecture adds further challenges related to performance. Calls inside a microservice are significantly lighter than calls between microservices. Still, the quality of the decomposition – the optimal slicing of the monolith to services – is critical for gaining the assumed benefits of using microservices. The software architects usually perform the decomposition manually but the practitioners have claimed that a tool to support identification different possible slicing solutions [25,31,32] would greatly help the task. Typically, the only helping tools for the software architects have been based on the static analysis of dependencies such as Structure 101¹. The actual discovery of slicing options has still been done by experienced software architects. In microservices, the dynamic behavior of the system is important too since it affects the performance and maintainability. Since static dependency analysis tools are not able to capture the dynamic behavior we decided to explore slicing based on runtime behavior instead of only considering static dependencies.

In our previous work, we proposed a microservice decomposition framework [36] based on process-mining to ease the identification of splitting candidates for decomposing a monolithic system into separate microservices. The framework is based on logs produced by process mining of the original monolithic system. The decomposition framework has been also validated in our previous study [36] in collaboration with an SME. The results of [36] shows that dynamic call history can be effectively used in the decomposition of microservices. This approach can also identify architectural issues in monolithic systems. The approach can be used by companies to proposed different slicing options to the software architects and to provide additional analysis of the software asset. This would reduce the risk of wrong slicing solutions.

In this paper, we extend the previous decomposition framework [36] proposing a new measurement framework to objectively compare two decomposition options. The measurement framework can be used independently from the decomposition strategy adopted.

The remainder of this paper is structured as follows. Section 2 presents the background on processes for migrating and splitting monolithic systems

¹ Structure101 Software Architecture Environment - <http://www.structure101.com>.

into microservices. Section 3 presents the measurement framework. Section 4 describes our proposed decomposition approach. Section 5 discusses the results. Section 6 presents related works while, finally Sect. 7 draws conclusions.

2 Background and Assumptions

Decomposing a system into independent subsystems is a task that has been performed for years in software engineering. Parnas [20] proposed the first approach for modularizing systems in 1972. After Parnas’s proposal, several works proposed different approaches [15]. Recently, the decomposition of systems took on another dimension thanks to cloud-native systems and especially microservices. In microservices, every module is developed as an independent and self-contained service.

2.1 Microservices

Microservices are small and autonomous services deployed independently, with a single and clearly defined purpose [7, 19]. In microservices each service can be developed using different languages and frameworks. Each service is deployed to their dedicated environment whatever efficient for them.

The communication between the services can be based on either REST or message queue. So, whenever there is a change in business logic in any of the services, the others are not affected as long as the communication endpoint is not changed. As a result if any of the components of the system fails, the failure will not affect the other components or services, which is a big drawback of monolithic system [7].

As we can see in Fig. 1, components in monolithic systems are tightly coupled with each other so that failure of one component will affect the whole system. Also if there is any architectural changes in a monolithic system it will also affect other components. Due to these advantages, microservice architecture is way more effective and efficient than monolithic systems.

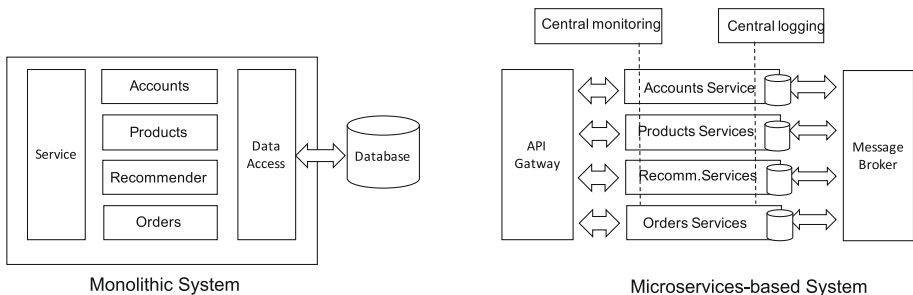


Fig. 1. Architectures of microservices and monolithic systems.

2.2 The Microservice Decomposition Process

Taibi et al. [31] conducted a survey among 21 practitioners who adopted microservices at least two years ago. The aim of the survey was to collect their motivation for, as well as the evolution, benefits, and issues of the adoption of microservices. Based on the results, they proposed a migration process framework composed of two processes for the redevelopment of the whole system from scratch and one process for creating new features with a microservice architecture on top of the existing system. They identified three different processes for migrating from a monolithic system to a microservices-based one. The goal of the first two processes is to support companies that need to migrate an existing monolithic system to microservices by re-implementing the system from scratch. The aim of the third approach is to implement new features only as microservices, to replace external services provided by third parties, or to develop features that need important changes and therefore can be considered as new features, thus gradually eliminating the existing system. All three of the identified processes are based on four common steps but differ in the details.

- *Analysis of the System Structure.* All processes start by analyzing dependencies mainly with the support of tools (Structure101, SchemaSpy², or others)
- *Definition of the New System Architecture.* Architectural guidelines or principles, and proposal of a decomposition solution into small microservices are defined. The decomposition is always done manually.
- *Prioritization of feature/service development.* In this step, all three processes identify and prioritize the next microservices to be implemented. Some processes prioritize microservices based on customer value; others according to components with more bugs; and yet others prioritize the development of new features as microservices, expecting that, in the long run, the new ecosystem of microservices will gradually replace each feature of the existing monolith.
- *Coding and Testing* are then carried out like any other software development project. Developers adopt the testing strategy they prefer. However, in some cases, testing of the different microservices is performed by doing unit testing at the microservices level and black-box testing at the integration level.

In this work, we focus mainly on the first two steps, supporting companies in the analysis of the system structure and in the identification of decomposition alternatives. The architectural guidelines should be defined by the company based on their internal policies.

² <http://schemaspy.sourceforge.net/>.

2.3 Architectural Goals

The microservices should be as cohesive and decoupled as possible [7]. The motivation of such architectural characteristics is to keep the maintenance as local as possible. In other words, the changes to the source code should be local to one microservice. Such *decoupled* architecture also supports independent development and deployment of the microservices. Sam Newman [19] describe loose coupling as follows: “a loosely coupled service knows as little as it needs to about the services with which it collaborates”.

Cohesion is related to decoupling and measures the degree to which the elements of a certain class belong together. Cohesion measures how weakly the functionalities of different modules are related to each other [6]. High cohesion often relates to low coupling [10, 13]. If the components of the software have high cohesion, the reasoning of the system is easier [13]. Thus, high cohesion supports efficient development and maintenance of the system.

In the design of microservice-based systems the developers target at high cohesion and low coupling by grouping the functionality and components according to the business processes. Then, changes to a functionality should lead changes to one microservice only [19].

Because cohesion and decoupling are key qualities of microservices, the dependency information is needed in the decomposition process. The commonly used dependency analysis tools, such as Structure 101, are based on static analysis of dependencies. They do not know which inter-component calls are really made and they do not recognize the full call paths. Our approach uses the dynamic dependency information that a process mining can provide. The mining provides recommendations, and analysis can then be used for reasoning. At this point, we do not aim at completely automated decomposition. In the next Subsection we report the underlying assumptions of our approach and details of the decomposition process.

2.4 Decomposition Framework Assumptions

The core assumption of our approach is existence of an extended log trace that has been collected at runtime. This means that the whole chain of operations after any external trigger can be traced from the log files. Examples of such external events include any user operations (e.g., clicking on a button) and calls from other applications (e.g., APIs or command line). The log file must include information about all methods and classes involved in serving of the request. The complete execution path must be completely traceable from the entry point to the access to the database (if any) and to the results returned to the client. The log must also include the start and end events. A hypothetical example of the data reported in the log file is shown in Table 1. A trace in Table 1 is identified by a session ID. That ID distinguishes the trace from other sessions.

Table 1. Example of log traces (timestamps are shortened for reasons of space).

Start time	End time	Sess. ID	Class	Method
00:00	00:36	S1	Form.jsp	btnClick()
01:00	01:39	S1	A.java	a()
01:40	01:45	S1	A.java	b()
01:45	01:55	S1	B.java	b()
01:56	02:05	S1	B.java	c()
02:05	02:13	S1	DB.java	query()
02:14	02:21	S1	DB	TABLE A
02:22	03:28	S1	DB	TABLE B
02:29	02:36	S1	B.java	c()
02:36	02:45	S1	B.java	b()
02:46	02:55	S1	A.java	b()
02:56	03:03	S1	A.java	c()
03:04	03:16	S1	Results.jsp	render()

There are several ways to collect the traces. One possible method is to instrument the source code, but using Aspect-Oriented Programming (AOP) can be done too, like in the work done by Suonsyrjä [26]. For some runtime systems it is also possible to instrument the executable file with tools like Elastic APM³. For Java programs our current recommendation is to use Elastic APM since the instrumentation with it requires a minimal effort. Depending on the language and on the technology adopted, other tools such as Dynatrace⁴ or Datadog⁵ could be also used.

3 The Microservice Measurement Framework

In this Section, we propose our microservice measurement framework. The framework has the goal of supporting companies to compare different microservices-based solutions, but also to understand the high-level architectural quality of their current system.

The measurement framework is based on availability of log files from real execution and is composed of four measures: coupling (CBM), number of classes per microservice (CLA), number of duplicated classes (DUP), and frequency of external calls (FEC).

This measurement framework is used on the decomposition framework presented in Sect. 4.

³ The Elastic APM Libraries. <https://www.elastic.co/solutions/apm>.

⁴ Dynatrace <https://www.dynatrace.com>.

⁵ Datadog <https://www.datadoghq.com>.

3.1 Coupling (CBM)

As reported in Sect. 2.3, in successful decompositions the coupling between microservices should be minimized and cohesion should be maximized. A comprehensive calculation of these measures would require knowledge of information beyond the log traces – for instance knowledge about access to local variables. Thus, we need to rely on approximation. One way to approximate coupling is to estimate it as an inverse to cohesion. Coupling can be considered as inversely proportional to cohesion and therefore, a system with low coupling will have a high likelihood of having high cohesion [10].

In our framework we adopt the metric “Coupling Between Microservice” (CBM) [36], a coupling measure inspired by the well-known Coupling Between Object (CBO) metric proposed by Chidamber and Kemerer [3]. CBO counts the number of classes coupled with a given class. Classes can be coupled through several mechanisms, including method calls, field accesses, inheritance, arguments, return types, and exceptions.

In [36] we calculate the relative CBM for each microservice as follows:

$$\text{CBM}_{MS_j} = \frac{\text{Number of external Links}}{\text{Number of Classes in the Microservice}}$$

In this formula “Number Of External Links” represents the number of call paths to external services. So, external services that are called several times, even by different classes of the microservice, are only counted once. The external services could be other microservices or services external to the whole system. The call frequency of external calls should also be taken into account, but we have a separate measure presented in Subsect. 3.4 for that.

3.2 Number of Classes per Microservice (CLA)

This measure is an abstraction of the size of the microservice, and allows the developer to discover services that are either too big or too small compared with other microservices. In general, smaller microservices are easier to maintain and thus large microservices should be avoided.

In some cases optimizing for CLA measure leads to compromises to other measurements. For instance, larger number smaller microservices may lead to stronger coupling (CBM) and higher frequency of external calls (FEC).

3.3 Number of Duplicated Classes (DUP)

The execution traces often have common sub-paths, i.e., some classes and methods are common for several execution traces. If traces should be implemented in different microservices, one way to increase independence is to duplicate part of the code to several microservices. For example, method *j* of class *E* (Fig. 4) is used by two execution traces. In that example the decomposition option 1

has one duplicated class, while option 2 requires no classes to be duplicated. Duplicated classes increases size of the system and complicates the maintenance process.

3.4 Frequency of External Calls (FEC)

Calls between microservices are computationally substantially heavier than calls within a microservice. Thus, reducing of the frequency of external calls optimizes the performance and delays. Since our approach is based on log-file analysis, we have the frequency information available.

We use the call frequency presented in Table 3 as an input to a measure relative Frequency of External Calls (FEC):

$$FEC_{MS_j} = \frac{\text{Number of Call Instances}}{\text{Number of Classes in the Microservice}}$$

As an example consider execution paths and decomposition options in Fig. 4. For the sake of the example we assume that

- Path A.a() → A.b() → B.c() → B.d() is called 200 times
- Path C.e() → C.f() → D.g() → D.h() is called 200 times.
- Path C.e() → C.f() → F.j() → D.g() → D.h() is called 50 times
- Path E.i() → E.j() → F.k() → F.l() is called 100 times.

With the input data we can calculate the number of internal calls, external calls, and FEC per each microservice. In Table 2 we also show the total number of internal in microservices (internal C), total number of calls between microservices (external C) and relative computational load (load). In this example we assume that an external call is 1000 times heavier than an internal call.

Table 2. Example analysis of call frequencies between microservices.

MS split	internal c	external c	load	FEC_{MS1}	FEC_{MS2}	FEC_{MS3}
0: A+B, C+D, E+F	1150	100	101500	0	25	25
1: A+B, C+D+E.j, E+F	1650	0	1650	0	0	0
2: A+B, C+D+E+F	1650	0	1650	0	0	

4 The Decomposition Framework

In this Section, we describe a decomposition framework that uses the data from the execution path analysis to discover the optimal split to micro services. A top-level description of the framework is given Fig. 2.

When the log files are available, the decomposition process defined in (Fig. 2) can be started. The process consists of six steps that are outlined in the following subsections.



Fig. 2. The decomposition process (from [36]).

4.1 Step 1: Execution Path Analysis

As our approach aims to optimize the system for the often used call sequences, the first step of our approach is to identify the most frequently executed call paths from the log files. One way to do that is use of a process-mining tool. In our case, we used DISCO⁶ to graphically represent the business processes found in the log files. Other similar tools can be used instead. The outcome of step 1 is a graphical representation of the processes and call graphs. One example of such graphical diagram is presented in Fig. 3. The diagram shows the call paths between classes, methods and data bases with arrows. This figure provides the user with the following information:

- The actually executed call paths in the system. Possible but never executed paths are not shown in this figure.
- Inter-class dependencies in the system. The dependencies are visualized with arrows between methods and classes. The external dependencies to libraries or web-services are also visualized.
- The usage frequency of each path. Process mining tools may present the frequency with thickness of the arrows or in a separate table as in Table 3.
- Branches and circular dependencies. If the system has circular dependencies or branches in the call path, those can easily be found from the visualization.

The call paths, shown with chains of arrows in Fig. 3, form candidates for business processes that are later used in the decomposition to microservices. For example, the path documented in Table 1 is visualized in a business process shown in Fig. 3.

Table 3. Frequency analysis of each execution path (from [36]).

Path	Freq.
A.a(); A.b(), B.b(), C.c(), DB.query, Table A, Table B, ...	1000
A.b(); A.c(), B.a(), C.c(), DB.query, Table A, Table B, ...	150

4.2 Step 2: Frequency Analysis of the Execution Paths

In our approach, the call frequency is a major contributor for the produced recommendations. Thus the frequency should be studied and analyzed. For visual

⁶ <https://fluxicon.com/disco/>.




Fig. 3. Simplified process example (from [36]).

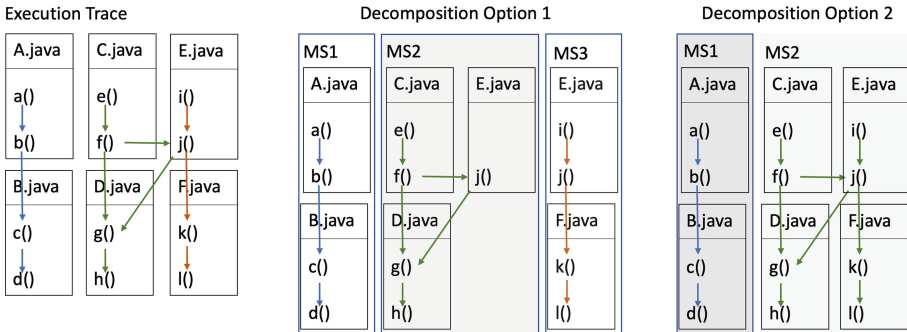


Fig. 4. Microservice decomposition example.

inspection, the process analysis tools can provide help. For instance, in the DISCO tool, the thickness of the arrows reflect the call frequency. In addition to visual inspection, we use concrete numeric data for all execution paths with their usage frequency. So, the output of this step is a table similar to Table 3.

The Frequency Analysis step helps the architect to select the potential decomposition options. The numbers are used to calculate the measures presented in Sect. 3 and used in step 6 (see Subject. 4.5).

4.3 Step 3: Removal of Circular Dependencies

In this step, we first find circular dependencies by analyzing the execution paths reported in the table generated in the first Step (e.g. Table 3). This can be done with a simple algorithm to discover cycles in the execution paths. In the case of circular dependencies, software architects should discuss with the development team how to break these cycles. One example of the patterns that can be applied to break the cycles is Inversion of Control [17]. However, every cyclic dependency could need a different breaking solution that must be analyzed carefully. The result is a refined version of the execution path table (see Table 3 as example).

Although this step is not part of the decomposition, it is important to consider cyclic dependencies to avoid possible deadlocks, and to better architect the system.

4.4 Step 4: Identification of Decomposition Options

In this step the execution paths from Step 3 are used to identify different decomposition opportunities. Visual inspection of the call graphs like those shown in Fig. 4 is used. We have relied on manual expert-based decomposition, but an algorithm to create different decompositions could be developed.

The execution paths may merge to common sub-paths or split into several branches. This leads to alternative decomposition solutions. This is demonstrated in Fig. 4. If the six source files each providing implementation of a class are assigned to three different microservices as A.java+B.java, C.java+D.java and E.java+F.java, the calls from C.f() to E.j() and E.j() to D.g() are inter-service calls. These calls are heavier operations than local calls and expand the external interfaces of the microservices. If use of external calls is not feasible, there are two other alternatives that can be proposed. The first option is to use three microservices so that class E (or at least function j() of it) is duplicated in two microservices. The other option is to decompose into two microservices as shown in the rightmost part of Fig. 4. Obviously, there is also the alternative to allow external calls and have three microservices with no duplications.

All these options have their advances and disadvantages, and the team should discuss the alternatives from multiple view points. The consideration could include both functionality of the software – if the paths belong logically together, and development process – what are the consequences of duplication to the development team(s). In addition, the call frequency has to be taken into account. For example in above split and merge case, the team has to consider both the development effort and run-time cost of making the two call external. The metrics discussed Sect. 3 help in analysing the run-time costs.

4.5 Step 5: Metric-Based Ranking of the Decomposition Options

In this step, we apply the measures identified in the Measurement Framework (Sect. 3), to help software architects to assess the quality of the decomposition options.

Sometimes optimization of the measures contradict with each other. Currently, we propose use of judgment of the team, but in the future approaches like Pareto [5] optimization could be used.

4.6 Step 6: Selection of the Decomposition Solution

In final step, the decomposition alternatives identified in Step 4 and the measures collected in Step 5 are used by the software architects to decide which solution to take.

Our approach does not automatically rank the solutions to any order. The software architects should consider the provided recommendations and measurements before selecting the solution. The team should discuss the relevance of the differences in their case.

5 Discussion

In this work we proposed a microservice measurement framework and we applied in our previously proposed decomposition process [36].

The measurement framework is based both on static and dynamic measures collected at runtime. The main benefit of analyzing runtime information is the availability of the data on the usage of each component, together with the dynamic analysis of dependencies. The dynamic measures allow to better understand the behavior of the system, and to analyze the dynamic coupling. Moreover, thanks to the dynamic measures collected, such as the frequency of usage of each method, the software architects can better understand which features is used more, we prioritized the development and the slicing of the monolithic system differently.

Companies could benefit from our lessons learned, by applying our proposed frameworks to decompose their monolithic system, but also to evaluation and monitoring the runtime behaviors or existing microservices to continuously understand possible issues. Moreover, the microservice measurement framework will allow software architects to clearly evaluate different decomposition options, with the usage of repeatable and objective measures.

Despite this approach being very beneficial in our case company, the results could have a different impact on other companies. Researchers can benefit from this approach and extend it further. New optimization metrics could be defined, and in theory, it would be possible to propose an automated decomposition approach that would identify the slices by maximizing the metrics identified. Genetic algorithms could be a possible solution for this idea.

6 Related Work

Fritzscht et al. [8] present a classification for the refactoring approaches. Our approach should be categorized as *Workload-Data aided* in this classification since we use operational data, i.e., dynamic data, in decomposition and analysis.

Bogner et al. have conducted a literature review of maintenance metrics of microservices [1]. The report summarizes several metrics that specialize in service-based systems instead of metrics designed for object-oriented systems. Although that research assume use of static info only, these metrics should be interesting for us in future research.

One case of refactoring a legacy system to a service based system has been reported by Khadka et al. [12]. Their case has substantial similarities to our approach and case. They also stress the importance and difficulty of finding the right set of services. They also analyze the call paths to find hotspots in

the code. However, they do not present a systematic and repeatable process for decomposition. Actually, only a limited set of research works propose systematic approaches for developers in decomposing their systems into an optimal set of microservices.

Chris Richardson [24] put scalability into focus and propose a decomposition approach based on the “scalability cube” where applications can be scaled based on the X, Y or Z axis. The X-axis and Z-axis scaling consists of running multiple copies of an application behind a load balancer. The Y-axis axis scaling is the real microservice decomposition approach, that propose to split the application into multiple, different services. Each service is responsible for one or more closely related functions. The decomposition is then based on two approaches: decomposing based on verbs used into the description of the service or decomposing by noun creating services responsible for all operations related to a particular entity such as customer management. Richardson also recommend to use combination of verb-based and noun-based decomposition when needed.

Richardson [23] also mention this approach in his two decomposition strategies:

- “Decompose by business capability and define services corresponding to business capabilities”;
- “Decompose by domain-driven design sub-domain”;

In an older version of this page [23] (2017), Richardson proposed other two patterns:

- “Decompose by verb or use ‘cases’ and define services that are responsible for particular actions”;
- “Decompose by nouns or resources by defining a service that is responsible for all operations on entities/resources of a given type”.

Kecskemeti et al. [11] proposed a decomposition approach based on container optimization. The goal is to increase the elasticity of large-scale applications and the possibility to obtain more flexible compositions with other services.

Another decomposition possibility is to start from a monolithic system and progressively move towards a microservices-based architecture [39] or delivering separate microservices by splitting a development team into smaller ones responsible for a limited group of microservices.

Vresk et al. [38] defined an IoT concept and platform based on the orchestration of different IoT system components, such as devices, data sources, data processors, and storage. They recommend an approach similar to the one proposed by Richardson’s Scalability Cube [24] combining verb-based and noun-based decomposition approaches. The proposed approach hides the complexity stemming from the variation of end-device properties thanks to the application of a uniform approach for modeling both physical and logical IoT devices and services. Moreover, it can foster interoperability and extensibility using diverse communication protocols into proxy microservice components.

Gysel et al. [9] proposed a clustering algorithm approach based on 16 coupling criteria derived from literature analysis and industry experience. This approach is an extensible tool framework for service decomposition as a combination of a criteria-driven methods. It integrates graph clustering algorithms and features priority scoring and nine types of analysis and design specifications. Moreover, this approach introduces the concept of coupling criteria cards using 16 different instances grouped into four categories: Cohesiveness, Compatibility, Constraints, and Communications. The approach was evaluated by integrating two existing graph clustering algorithms, combining actions research and case study investigations, and load tests. The results showed potential benefits to the practitioners, also confirmed by user feedback.

Chen et al. [2] proposed a data-driven microservices-oriented decomposition approach based on data flow diagrams from business logic. Their approach could deliver more rational, objective, and easy-to-understand results thanks to objective operations and data extracted from real-world business logic. Similarly, we adopt process mining to analyze the business processes of a monolithic system.

Alwis et al. [4] proposed a heuristic to slice a monolithic system into microservices based on object subtypes (i.e., the lowest granularity of software based on structural properties) and functional splitting based on common execution fragments across software (i.e., the lowest granularity of software based on behavioral properties). This approach is the closer to our work. However, they analyzed the system by means of static analysis without capturing the dynamic behavior of the system and they did not propose measures to evaluate the quality of the slicing solution proposed.

Taibi et al. [28,33,34], proposed a set of patterns and anti-patterns that should be carefully considered during the microservice decomposition [28,33] recommending to avoid a set of harmful practices such as cyclic dependencies and hard-coded endpoints but also to consider critical anti-patterns and code smells [27] that can be generated into the monolithic system.

7 Conclusions

The decomposition of monolithic systems into microservices is a very complex and error-prone task, commonly performed manually by the software architect.

In this work, we first proposed a new microservice measurement framework based on 4 measures: coupling, number of classes per microservices, number of duplicated classes and frequency of external calls. The goal of our framework is to support software architects to compare different microservice decompositions, by means of a set of objective and repeatable measures.

We instantiated our measurement framework in the context of the previously proposed process-mining decomposition approach [36].

Our goal is not the automated slicing of monolithic systems but to provide extra support to software architect, to help them in identifying different slicing options reducing the subjectivity and to measure and compare the different solutions objectively.

The microservice measurement framework can be adopted independently from the decomposition process used. As example, software architect might manually identify two decomposition options for a monolithic system. The measurement framework will support them in the comparison of their decomposition options.

We recommend companies to apply periodically our measurement framework also in case of existing microservices-based systems. The historical analysis of the evolution of the system might provide useful information on the quality of the system and also be a trigger for future refactorings.

Future works include the validation of the framework, both in the context of manual decompositions and when using process-mining based approaches. Moreover, we want to evaluate the application of our approach in development of a tool to facilitate the identification of the process, the automatic calculation of the metrics, and identification of other measures for evaluating the quality of the decomposition. We already started to develop a tool to automatically identify dependencies between microservices [21] and we published a dataset containing the analysis of 20 projects [22].

We are also planning to further empirically validate this approach with other companies and to include dynamic measures for evaluating the quality of the system at runtime [16,37]. In the future, we are also planning to adopt mining software repositories techniques to identify the areas that changed simultaneously in the past, to help developers to understand pieces of code connected to each other.

Another possible future work is to include identification of partial migration, i.e., migration of a limited set of processes from a monolithic system. Finally, we are also considering to extend this work by proposing not only different decomposition options but also a set of patterns for connecting microservices based on existing common microservices patterns [19,32] and anti-patterns [28,33,34].

References

1. 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, IWSM Mensura 2017, pp. 107–115. ACM, New York (2017). <https://doi.org/10.1145/3143434.3143443>
2. Chen, R., Li, S., Li, Z.: From monolith to microservices: a dataflow-driven approach. In: 2017 24th Asia-Pacific Software Engineering Conference (APSEC), pp. 466–475, December 2017
3. Chidamber, S.R., Kemerer, C.F.: A metrics suite for object oriented design. *IEEE Trans. Softw. Eng.* **20**(6), 476–493 (1994)
4. De Alwis, A.A.C., Barros, A., Polyvyanyy, A., Fidge, C.: Function-splitting heuristics for discovery of microservices in enterprise systems. In: Pahl, C., Vukovic, M., Yin, J., Yu, Q. (eds.) ICSSOC 2018. LNCS, vol. 11236, pp. 37–53. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03596-9_3

5. Deb, K.: Evolutionary algorithms for multi-criterion optimization in engineering design. In: EUROGEN 1999, pp. 135–161 (1999)
6. Fenton, N., Bieman, J.: *Software Metrics: A Rigorous and Practical Approach*, 3rd edn. CRC Press, Boca Raton (2014)
7. Fowler, M., Lewis, J.: *Microservices* (2014). <http://martinfowler.com/articles/microservices.html>
8. Fritzsche, J., Bogner, J., Zimmermann, A., Wagner, S.: From monolith to microservices: a classification of refactoring approaches. In: Bruel, J.-M., Mazzara, M., Meyer, B. (eds.) DEVOPS 2018. LNCS, vol. 11350, pp. 128–141. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-06019-0_10
9. Gysel, M., Kölbener, L., Giersche, W., Zimmermann, O.: Service cutter: a systematic approach to service decomposition. In: Aiello, M., Johnsen, E.B., Dustdar, S., Georgievski, I. (eds.) ESOC 2016. LNCS, vol. 9846, pp. 185–200. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-44482-6_12
10. Jabangwe, R., Börstler, J., Šmite, D., Wohlin, C.: Empirical evidence on the link between object-oriented measures and external quality attributes: a systematic literature review. *Empir. Softw. Eng.* **20**(3), 640–693 (2014). <https://doi.org/10.1007/s10664-013-9291-7>
11. Kecskemeti, G., Marosi, A.C., Kertesz, A.: The entice approach to decompose monolithic services into microservices. In: 2016 International Conference on High Performance Computing Simulation (HPCS), pp. 591–596, July 2016
12. Khadka, R., Saeidi, A., Jansen, S., Hage, J., Haas, G.P.: Migrating a large scale legacy application to SOA: challenges and lessons learned. In: 2013 20th Working Conference on Reverse Engineering (WCRE), pp. 425–432, October 2013. <https://doi.org/10.1109/WCRE.2013.6671318>
13. Kramer, S., Kaindl, H.: Coupling and cohesion metrics for knowledge-based systems using frames and rules. *ACM Trans. Softw. Eng. Methodol.* **13**(3), 332–358 (2004)
14. Lehman, M.M.: On understanding laws, evolution, and conservation in the large-program life cycle. *J. Syst. Softw.* **1**, 213–221 (1984). [https://doi.org/10.1016/0164-1212\(79\)90022-0](https://doi.org/10.1016/0164-1212(79)90022-0)
15. Lenarduzzi, V., Sillitti, A., Taibi, D.: Analyzing forty years of software maintenance models. In: 39th International Conference on Software Engineering Companion, ICSE-C 2017, IEEE Press (2017). <https://doi.org/10.1109/ICSE-C.2017.122>
16. Lenarduzzi, V., Stan, C., Taibi, D., Venters, G.: A dynamical quality model to continuously monitor software maintenance. In: 11th European Conference on Information Systems Management (ECISM) (2017)
17. Martin, R.C.: *Agile Software Development: Principles, Patterns, and Practices*. Prentice Hall PTR, Upper Saddle River (2003)
18. Mikkonen, T., Systä, K.: Maximizing product value: continuous maintenance. In: Jedlitschka, A., Kuvaja, P., Kuhrmann, M., Männistö, T., Münch, J., Raatikainen, M. (eds.) PROFES 2014. LNCS, vol. 8892, pp. 298–301. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-13835-0_26
19. Newman, S.: *Building Microservices*, 1st edn. O'Reilly Media Inc., Newton (2015)
20. Parnas, D.L.: On the criteria to be used in decomposing systems into modules. *Commun. ACM* **15**(12), 1053–1058 (1972)
21. Rahman, M., Taibi, D.: Microservice dependency graph (microdegraph) (2019). <https://github.com/clowee/MicroDepGraph>
22. Rahman, M., Taibi, D.: The microservice dependency graph dataset. In: Proceedings of the Joint Summer School on Software Evolution, Software Quality and Bug Prediction with Machine Learning. CEUR-WS, September 2019

23. Richardson, C.: Pattern: Microservice Architecture (2017). <http://microservices.io/patterns/microservices.html>
24. Richardson, C.: The Scale Cube (2017). <https://microservices.io/articles/scalecube.html>
25. Soldani, J., Tamburri, D.A., Heuvel, W.J.V.D.: The pains and gains of microservices: a systematic grey literature review. *J. Syst. Softw.* **146**, 215–232 (2018)
26. Suonsyrjä, S., Mikkonen, T.: Designing an unobtrusive analytics framework for monitoring Java applications. In: Kobyliński, A., Czarnacka-Chrobot, B., Świerczek, J. (eds.) *IWSM/Mensura -2015. LNBIP*, vol. 230, pp. 160–175. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24285-9_11
27. Taibi, D., Janes, A., Lenarduzzi, V.: How developers perceive smells in source code: a replicated study. *Inf. Softw. Technol.* **92**, 223–235 (2017)
28. Taibi, D., Lenarduzzi, V.: On the definition of microservice bad smells. *IEEE Softw.* **35**(3), 56–62 (2018)
29. Taibi, D., Lenarduzzi, V., Ahmad, M.O., Liukkunen, K.: Comparing communication effort within the scrum, scrum with Kanban, XP, and Banana development processes. In: *Proceedings of the 21st International Conference on Evaluation and Assessment in Software Engineering, EASE 2017*, pp. 258–263 (2017)
30. Taibi, D., Lenarduzzi, V., Janes, A., Liukkunen, K., Ahmad, M.O.: Comparing requirements decomposition within the scrum, scrum with Kanban, XP, and Banana development processes. In: *Agile Processes in Software Engineering and Extreme Programming*, pp. 68–83 (2017)
31. Taibi, D., Lenarduzzi, V., Pahl, C.: Processes, motivations, and issues for migrating to microservices architectures: an empirical investigation. *IEEE Cloud Comput.* **4**(5), 22–32 (2017)
32. Taibi, D., Lenarduzzi, V., Pahl, C.: Architectural patterns for microservices: a systematic mapping study. In: *8th International Conference on Cloud Computing and Services Science, CLOSER 2018* (2018)
33. Taibi, D., Lenarduzzi, V., Pahl, C.: Microservices anti-patterns: a taxonomy. *IWSM/Mensura -2015. LNBIP*, vol. 1073, pp. 111–128. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-31646-4_5
34. Taibi, D., Lenarduzzi, V., Pahl, C.: Continuous architecting with microservices and DevOps: a systematic mapping study. In: Muñoz, V.M., Ferguson, D., Helfert, M., Pahl, C. (eds.) *CLOSER 2018. CCIS*, vol. 1073, pp. 126–151. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-29193-8_7
35. Taibi, D., Lenarduzzi, V., Pahl, C., Janes, A.: Microservices in agile software development: a workshop-based study into issues, advantages, and disadvantages. In: *XP Workshops*, pp. 23:1–23:5. ACM (2017)
36. Taibi, D., Systa, K.: From monolithic systems to microservices: a decomposition framework based on process mining. In: *9th International Conference on Cloud Computing and Services Science, CLOSER 2019, Heraklion, Greece, 05 2019* (2019)
37. Tosi, D., Lavazza, L., Morasca, S., Taibi, D.: On the definition of dynamic software measures. In: *ESEM*, pp. 39–48. ACM (2012)
38. Vresk, T., Cavrak, I.: Architecture of an interoperable IoT platform based on microservices. In: *MIPRO*, pp. 1196–1201. IEEE (2016)
39. Zimmermann, O.: Microservices tenets. *Comput. Sci. Res. Dev.* **32**(3), 301–310 (2016). <https://doi.org/10.1007/s00450-016-0337-0>