

Microservice Architecture Reconstruction and Visualization Techniques: A Review

Tomas Cerny

*Department of Computer Science
Baylor University
Waco, Texas, United States
tomas_cerny@baylor.edu*

Amr S. Abdelfattah

*Department of Computer Science
Baylor University
Waco, Texas, United States
amr_elsayed1@baylor.edu*

Vincent Bushong

*Department of Computer Science
Baylor University
Waco, Texas, United States
vinbush@gmail.com*

Abdullah Al Maruf

*Department of Computer Science
Baylor University
Waco, Texas, United States
maruf_maruf1@baylor.edu*

Davide Taibi

*CloudSEA.AI Group
Tampere University
Tampere, FI-33720, Finland
davide.taibi@tuni.fi*

Abstract—Microservice system solutions are now mainstream. The older microservices-based systems are not more than 15 years old, and their architecture is by far different than the one originally designed because of several changes applied to the systems due to the implementation of new features and bug fixing. The evolution of these legacy systems is therefore subjected to degradation. One of the most important methods to identify degradation is being able to reconstruct the software architecture of a system based on the current system running in production. Different methods have been proposed in the past: methods based on the static analysis of the source code of the microservices and methods based on the analysis of the log traces collected at runtime. Both static and dynamic analysis-based methods have their pros and cons. In this work, we review the existing technologies for static and dynamic architectural reconstruction and related tools adopted to visualize the reconstructed architecture. The result of this work can be useful both to practitioners and researchers that can further develop these methods to provide better support for architectural degradation.

Index Terms—Microservices, Software Architecture Reconstruction, Visualization, System-centric view, Decentralization

I. INTRODUCTION

The software architecture provides the major perspective for the system's development and design. The software architecture serves as the blueprint for systems. The system architecture must often be reconstructed to determine if the system was built accurately. Various works described such a process. O'Brien's report [1] defines architecture reconstruction as "the process by which the architecture of an implemented system is obtained from the existing system." The results are then used to "evaluate the conformance of the as-built architecture to the as-documented architecture" to reconstruct "architecture descriptions for systems that are poorly documented or for which documentation is not

This material is based upon work supported by the National Science Foundation under Grant No. 1854049, grant from Red Hat Research, and Ulla Tuominen (Shapit).

available", and "to analyze and understand the architecture of existing systems to enable modification of the architecture to satisfy new requirements and to eliminate existing software deficiencies."

The reconstruction process derives a representation of software architecture from artifacts such as documentation or, more commonly, the source code or runtime traces. The result assists developers in better understanding the system in question, and it plays a key role in other tasks, such as architecture verification, conformance checking, and trade-off analysis [2]. Besides these, it is also relevant when facing issues related to software architecture degradation, the process where, due to changes in the codebase, a system's architecture drifts away from the originally intended architecture.

In the context of service-oriented architectures, particularly microservices, the reconstruction process has profound importance and the potential to derive the view of the overall decentralized system. Such a view then shows how the system works [2]. Obviously, there is a significant difference between assessing monolith systems and decentralized systems like cloud-native microservices. Specifically, the codebase is specific to each cloud-native microservice [3]. Each codebase is self-contained and possibly managed possibly by different teams. Moreover, each microservice can follow different conventions, use different versions of libraries, and even follow different platforms. This all makes the reconstruction process more challenging for microservices.

One categorization of reconstruction methods is based on how the analysis is performed. There are three broad groupings: *dynamic* or runtime analysis, where a tool constructs the view at runtime; *static* analysis, where the view is constructed from artifacts available before deployment [4]; and *manual* analysis, where a human examines the system and manually constructs a representation of it [5].

Manual analysis, while not involving an automated tool, is an important step in establishing a proposed method or validating the results.

In order to shed light on the existing approaches and technologies for the Software Architectural Reconstruction (SAR) of microservices, in this paper, we provide a review of the different methods available.

The result of this work could help practitioners and researchers to understand which technique is most suitable for reconstructing microservices-based systems and which tool they can adopt to visualize them. Moreover, the results can also be useful to tool providers that could address the gaps by providing better support for architectural reconstruction or microservices.

This paper is organized as follows. Section II details the dynamic analysis approach, followed by the static analysis approach in Section III. One major outcome of SAR is architectural visualization through views, which we detail and discuss in Section IV, along with current tools for this. We discuss the approaches in Section V and conclude this paper in Section VI.

II. DYNAMIC SAR

Dynamic analysis can operate on several different runtime data sources. It has been used for a myriad of end goals, ranging from analyzing runtime traces extracted from logs to find timing errors [6] to extracting service dependency graphs by extracting remote procedure calls from network logs in a microservice mesh [7] or by uncovering dependencies between monitored metrics for components of a distributed system [8].

Runtime analysis-based SAR has taken many forms. One technique is to use instrumentation to insert logging statements to report events. It is also possible to use a specialized framework or custom-made annotations in a program, with calls intercepted producing events to be analyzed at runtime. Such events can be used to describe dependencies in microservice systems by detecting where microservices call each other. This approach has been used to model microservice dependencies and find incomplete test coverage of calls across an entire microservice mesh [9]. This approach can also be used to detect discrepancies between required and provided service versions, as well as generate performance metrics from service error data [10].

Interceptors can be used to a similar end; Mayer and Weinreich use the Spring framework's interceptors to monitor runtime calls between services to generate an architectural view of a microservice system [11]. Similarly, calls could be intercepted and rerouted through a security gateway [12], but this can bring significant performance overhead and violation of the distribution with potential bottlenecks. In general, approaches dependent on code instrumentation [11] bring additional development difficulty and overhead.

Another runtime approach is to build on what we have mentioned above and utilize the underlying containeriza-

tion engine. Since microservices are often deployed using containers, container configuration files can be a valuable source of information about the application's architecture. Granatelli et al. [13] query the containerization framework to retrieve calls between microservices at runtime. The extracted calls are used along with deployment metadata collected from service descriptors to create an architectural model for a microservice system. This approach is limited in extracting further system concerns because not all information is available through the containerization engine, especially information relating to how the application represents and operates on data.

Finally, the industry practice is to use existing monitoring, tracing, and metrics tools to capture data about the microservices [14]. Established tools recognize enterprise frameworks and utilize existing mechanisms, such as method call interception, instrumentation, or interaction with API-gateway. The additional step to deal with decentralization is that for cloud-native solutions, each logged statement includes correlation ID and origin location. This allows the analytics tools to connect statements related to the same distributed transaction to improve analytics and reasoning. Among industry tool and technology examples, consider OpenTelemetry, Kiali, Zipkin, or Jaeger ¹. These tools help determine system architecture with specific perspectives they provide from the dynamic information they collect, typically in the form of dependency graphs, directed acyclic graphs, or a topology view. On top of existing tools, others can build additional perspectives. For instance, Khan [15] constructed the tool MsViz using Jaeger, Grafana, and Prometheus to capture the architecture of microservice systems. He created a graph showing the architecture and overlaid performance data and other metrics onto the graph.

Dynamic analysis approaches benefit from the ability to access runtime data (such as performance metrics and real-time service calls), but they require the system to be deployed, running and used by users. Researchers point out many challenges with distributed tracing. For instance, Bento et al. [16] indicate problems with complexity, application specificity, the volume of information, and lack of tools to abstract, navigate, filter, and analyze trace data in an automated fashion relying on administrators to do it themselves.

Moreover, with dynamic analysis, we might identify the system endpoints; however, without a complete interaction trace over the system, we will fail to identify the complete system view, and this might be a critical aspect sometimes overlooked when compared to static analysis. For instance, if we intend to perform SAR or microservice correlations and consistency reasoning before production deployment, we need to integrate SAR into the software development lifecycle. Furthermore, unless robust testing and complete

¹<https://istio.io/latest/docs/tasks/observability/kiali/>;
<https://opentelemetry.io>; <https://zipkin.io>; <http://jaegertracing.io>

test coverage exist, we might not identify all possible communication paths properly and sufficiently. A broader view and precise detail can be uncovered with access to the codebase, but that involves static analysis.

III. STATIC SAR

Static analysis has on great advantage over dynamic analysis. It can be performed on a system before it is deployed, extracting information from existing artifacts that would otherwise have to be manually analyzed. The most common association with static analysis might be formal verification, where the program's codebases are analyzed to determine the system's correctness [17], [18]. However, static analysis has been used in automatically generating test cases for a program, for example, by identifying points for performance analysis instrumentation [19] or by extracting and analyzing an abstract syntax tree to identify all execution paths that need to be tested [20]. Developers can also use static analysis to better understand a program at a higher level. For example, UML models can automatically be generated by static analysis for legacy systems to better understand how to maintain or replace them [21], and it is integral in identifying code clones [22]–[24].

Static code analysis has also been applied to microservices [4]. It has been used to analyze monolithic systems and recommend splits for converting to microservices [25]. Static analysis cannot determine how instances will be instantiated and interact when deriving a precise service dependency graph [7]. However, we must consider that deployment descriptors are available in the codebase indicating remote services by aliases. Next, remote calls are evident in the source code. Then, when a system does not contain an indirection (i.e., enterprise service bus), which is the case of microservices. The service dependency graph can be constructed from the source code. Using Java source files and Docker/Spring configuration files have been demonstrated suitable [26] to reconstruct the architecture of microservices-based systems to identify cyclic dependencies between microservices. Analyzing code to visualize call-graphs between microservices was also proven feasible in [27].

Source code is not the only artifact available for static analysis. It is important to mention that many core artifacts like maven or docker files are typically included in the codebase and can be used for static analysis. Ibrahim et al. use a project's Dockerfiles to search for known security vulnerabilities of the container images being used, which they overlay on the system topology extracted from Docker Compose files to generate an attack graph showing how a security breach could be propagated through a microservice mesh [28]. This allows the creation of a centralized security concern for the system, but since it does not extend to source code, it cannot include security flaws in the programs deployed in the containers, only flaws with the images themselves. Another static source of information is

in the API definitions. Mayer and Weinreich used API definitions generated by Swagger as an input to their architecture generation system, but their system is also dependent on runtime data extracted from calls between services [11].

Another approach to pre-runtime SAR is to embed a source of information in the microservices as part of their development. For example, Salvadori et al. propose creating semantic microservices that expose information about their resources, allowing them to be automatically composed [29]. In this way, a centralized view of microservice communication is always available. However, this approach depends on using a fundamentally different approach to development, and it cannot be used to analyze existing codebases.

IV. ARCHITECTURAL VIEWS

Software architectures can be described by architectural views [30]. These views capture certain system qualities or aspects. For instance, they are elaborated by the 4+1 architectural view model [31] involving a logical view, process view, development view, physical view, and scenarios, which do not have a visual format. Furthermore, the 4+1 model can be generalized to the N+1 model [32].

The foundation for successful SAR is the ability to reconstruct effective architectural views of a system [33]. Existing SAR work related to microservices by Rademacher et al. [2] has considered four views as their outcome. In particular, it operated with domain, technology, service, and operation views.

- *Domain view* illustrates the domain concepts of in microservice system. More specifically, it describes the data entities of the system along with datasource connections of those entities.
- *Technology view* details technologies used for microservice implementation and operation.
- *Service view* copes with the service models that specify microservice interfaces and endpoints.
- *Operation view* then helps the operation to better understand service deployment in the infrastructure. It details containerization, service discovery, and monitoring.

Each of these views considered a specific perspective and related concerns within the system. However, each also relates to other views. As an example, consider the service view overlapping with the domain view to detail which data entities are involved in endpoints. The technology and domain view will then show where the data entities persist.

As suggested by Walker et al. [33], a key point for the construction of these views is that each view is an aggregation of smaller views, each illustrating a disparate microservice. Each microservice can be seen as operating within its bounded context [34]–[36] of its microservice concerns, but these can be aggregated into a fully centralized perspective of the system's architecture.

A. Views in the Context of Separation of Duty

These views can be put in contrast to the cloud-native perspectives suggested by Carnell et al. [14] when building

microservices. Their book (chapter 3) highlights that "the foundation for successful microservice development starts with the perspectives of three critical roles. The *architect* that sees the big picture of decomposing an application into individual microservices and understanding their interactions. The *software developer* who codes and understands the language and development frameworks to deliver a microservice (its functions and use cases). The *DevOps engineer* that determines how the services are deployed and managed throughout production and non-production environments". All three roles are essential to ensure proper microservices development.

Clearly, the architect knows the big picture, while the developer knows one piece of the puzzle in broad detail; finally, the DevOps engineer does not need to know much about what is encapsulated in each microservice but cares about the deployment process, operations, and monitoring. Architectural views should benefit each of these three roles, while their goals and system knowledge are diametrically different. There are, however, overlaps. As of now, it is difficult for the architect to observe whether the developers did their job properly to assure consistency and dependencies across microservices. Proper architectural views would help the architect to ensure the specification was met by the implementation. The DevOps engineers lack a broader understanding of what are the specific needs of the microservice they were asked to deploy. Thus, access to more details (e.g., persistence) would help them to better optimize custom deployment rather than using a one-size-fits-all approach. Finally, developers sometimes need to know what other knowledge microservices encapsulate but might lack time due to the business pressure to assess the specific and resort to wheel reinvention leading to cloned knowledge, business rules, or functionality, introducing technical debt, that sooner or later leads to inconsistencies due to decentralized system evolution leading to architecture degradation.

Moreover, other perspectives can be involved, such as security audits. It is unrealistic to expect one individual to be an expert in all these roles with their different perspectives. Each of these roles needs to have access to distinct architectural views and to assess the system per microservice and holistically. Each microservice is developed by developers who are not architects, not DevOps engineers, and likely not security experts. It is unrealistic to expect one individual to be an expert in all these roles, and while there are probably some, generally, we can expect separate experts. Thus, security experts need to assess the system regarding security mechanisms, and privacy and determine weaknesses and vulnerabilities. Similarly, we can account for performance analysis, considering the system for an entirely different goal. Each of these roles needs to have access to distinct architectural views, assess the system per microservice, and also holistically.

Thus, in the result, we might again refer back to the N+1 model [32]. These views have been detailed in [37], such

as *logical view*, *process view*, *deployment/physical view*, *data view*, *security view*, *implementation view*, *development view*, and *use case view*. Not all views can be easily extracted, such as the use case view. We detail the other views considered in the N+1 model:

- *Logical view* illustrates the conceptual organization of the software in terms of the most important layers, subsystems, packages, frameworks, classes, and interfaces. Typically uses UML package, class, and interaction diagrams.
- *Process view* details processes and threads, their responsibilities, collaborations, and the allocation of logical elements to them. Typically uses UML interaction and activity diagrams or possibly Business Process Model Notation (BPMN).
- *Deployment/Physical view* shows the physical deployment of processes and components to processing nodes and the physical network configuration between nodes. This overlaps with the Operation view suggested earlier by [2]. UML deployment diagrams are the natural fit for this view.
- *Data view* gives an overview of the data flows, persistent data schema, the schema mapping from objects to persistent data (data source), the mechanism of mapping from objects to a database, and database stored procedures and triggers. This partially overlaps with the domain view by [2] but is more broad considering data flow.
- *Security view* is an overview of the security schemes and points within the architecture that security is applied, such as HTTP authentication, database authentication, and so forth. It can take advantage of other views starting with the deployment view through the logical view.
- *Implementation view* details the implementation, these days mostly the specific components involved in the code, packages, and other resources, including other used libraries, which brings an overlap to the technology view.
- *Development view* detailing the organization of each codebase to help developers orient themselves in the code.
- *Use Case view* giving a summary of the most architecturally significant use-cases and their non-functional requirements

B. Traditional View Modeling

When considering enterprise architecture, the main emphasis is on business processes. To model enterprise architecture, available languages include ArchiMate, UML, Business Motivation Model (BMM), and BPMN, among others [38]. The major focus of UML is objects or components. ArchiMate has derived several concepts from UML but focuses mostly on services. This makes the ArchiMate more suiting to large systems with less detail than when modeling with UML.

Frameworks for the enterprise architecture practice include The Open Group Architecture Framework (TOGAF). ArchiMate then uses Architectural Development to extend TOGAF. With TOGAF, architectural modeling considers four levels with different specializations: Business, Application, Data, and Technology, which aligns a similarity with Rademacher et al. [2]. However, the business architecture levels are not covered by Rademacher et al. since this kind of architecture consists of motivation, organization, and mapping of assets, which, while encoded in the system, rather drive the motivation for the implementation. In particular, these levels are:

- *Business architecture* considering governance, organization, the key business processes, actors, services, and qualities when adapting existing processes.
- *Applications architecture* detailing the individual system deployment including services, logical and physical components, the interactions between the systems, and their relationships to the business processes. (Sharing similarities with the service and technology views).
- *Data architecture* describing the structure of an organization's logical and physical data assets and the associated data management resources, including data entities. (Similar to the domain view).
- *Technology architecture* describes the hardware, software, and network infrastructure supporting the deployment of core systems. (Sharing similarities with the technology and operation views).

C. Hierarchical View Modeling

It is recognized that microservices produce complex systems that might need a different level of abstraction when browsing throughout the system. The hierarchical approach of one of the natural choices. The Context, Containers, Components, and Code or just C4 model is worth mentioning as a practical approach for modeling software architecture [39]. It is a hierarchical model consisting of four levels of abstraction. The high-level system context can let operators transfer to individual code elements. C4 is a natural fit for microservices. While it does not prescribe a method of analysis, its key feature is important for analysis tools to follow in that it allows varying levels of abstraction for the users to see. Such a hierarchical analysis is useful especially for microservices, as it allows views of the system as a whole and inspection of individual services.

Different tools have been proposed to visualize the service call graph.

Rahman and Taibi [27] proposed a MicroDepGraph², a visualization tool to present the connected services together with the MicroServices Dataset [27]. However, MicroDepGraph does not distinguish between type of services, adopting the same type of shapes to visualize services, databases and other components (e.g. message buses). An example of

the visualization proposed by MicroDepGraph is shown in Figure 1.

In our previous work [40], we developed Prophet³, a static analysis tool to parse Java-based applications to reconstruct the service call graph of a microservice-based system. Prophet recognizes component-based constructs behind Spring Boot and Enterprise Java. Prophet provides an intermediate graph representation of the system accessible through REST API, to enable system reasoning (Fig 2).

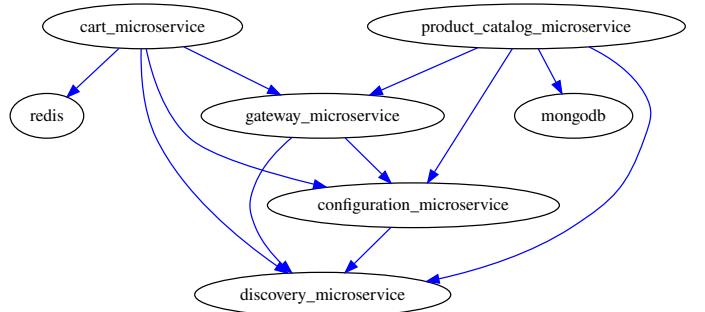


Fig. 1. MicroDepGraph (from [27])

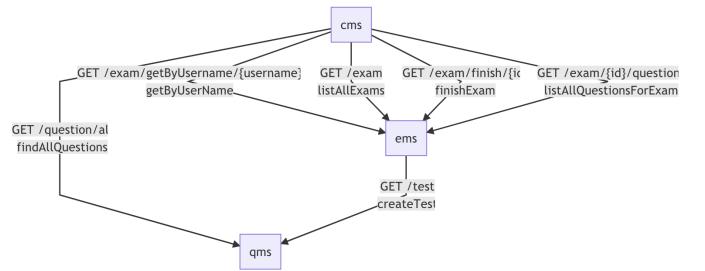


Fig. 2. Prophet service view (from [40])

D. Alternative Visualization for Software Architectures

In the area of more general visualization approaches for software architectures [41], [42] we can operate with various graph-based visualizations showing nodes and edges similar to ontologies, notation-based visualization such as UML or SysML or matrix-based approaches that act as a complementary representation of a graph. Another visualization area of great research interest is metaphor-based visualization. The visualization uses familiar physical world contexts (e.g., cities, islands, or landscapes).

To select the right visualization category, it has been noted [41] that these visualizations often serve a specific purpose. Among the main motivation to visualize architecture is architecture recovery, followed by architecture evolution, impact analysis, general analysis, synthesis, implementation, and reuse.

³Prophet: <https://github.com/cloudee/Prophet-utils>, <https://github.com/cloudee/Prophet-utils-app>, <https://github.com/cloudee/prophet>

²MicroDepGraph <https://github.com/cloudee/MicroDepGraph>

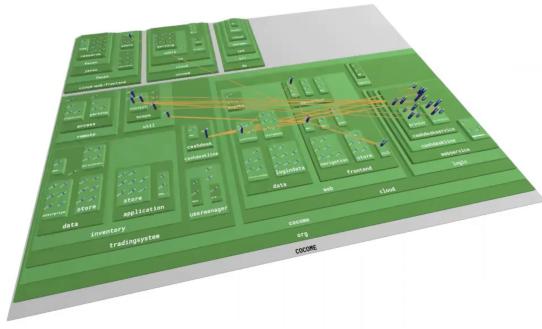


Fig. 3. The “software city” metaphor displays an application’s package structure as a series of structures building on each other. [43]



Fig. 4. Software components can be displayed using the “software island” metaphor, showing components and their inter-dependencies. [44]

The 3D visualization space has been assessed in the literature, but not that much in the context of microservices. In particular, for software architecture visualization, Virtual and Augmented Reality (VR/AR) are good candidates. These can be used as a visual metaphor to make the system more understandable. This example is a metaphor for a software city where software packages are represented as buildings and their dependencies as streets. Fittkau et al. implemented the software city metaphor in virtual reality [43] from Fig. 3. Steinbeck et al. [45] presented an even more advanced and scalable derivative called EvoStreets, which gives a better view of the software’s hierarchical makeup. Consider another example shown in Fig. 4. The approach shows individual software modules as islands in an ocean displayed in AR. This visualization method [44] is more closely applicable to a microservice architecture. Software packages and classes in each module are represented as regions and buildings on the module island, and, importantly, module imports and exports are displayed as ports that connect the different islands. While this approach has only been used on monolithic applications, the island metaphor is suitable for displaying the relationships between independent modules in a microservice architecture.

Large microservice-based systems are prime candidates for being visualized using VR/AR. One such approach is VR-EA tool [46]. Instead of doing dynamic or static analysis to extract a model, VR-EA uses modeling tools as inputs to generate a 3D VR view in the virtual reality of business processes and their relationships with enterprise resources. This approach can provide a comprehensive view of the enterprise system, as it can show a large group of interconnected components. However, it depends on a set of models that must be custom-created to capture the relationships



Fig. 5. Semantic information is a candidate for being displayed in large, three-dimensional graphs due to the natural connections between the elements. [49]

and complexities inside the large system, requiring manual creation of additional configuration and artifacts.

Virtual reality was used by Ma et al. to monitor a distributed set of servers, visualizing each server as a physical machine in the same VR room [47]. Although the monitoring capabilities were limited to system resource usage, the tool showed that physically disparate systems could be virtually collocated to provide a centralized view of a system.

More generally, large systems beyond software architecture have been explored in virtual reality. For example, Toumpalidis et al. used augmented reality to visualize data from IoT networks [48]. A user could see a summary of a device’s data overlaid on that device. This experiment showed that AR is useful for displaying and aggregating distributed data.

Three-dimensional visualizations can also be employed to visualize complex information relationships. Halpin et al. use virtual reality to display the relationships between patent registrants [49] as shown in Fig. 5, and Royston et al. use a similar idea to display connections from social media sites [50]. Neither of these approaches uses a specific visual metaphor, instead opting to display their contents as simple graphs in three-dimensional space. These examples have a common structure with the microservice architectures, as microservices architecture can be viewed as a network of services communicating with each other based on semantic relationships.

Moreover, Moreno-Lumbreras et al. [51] have suggested using VR to visualize development metrics and analytics in three-dimensional space. Acquiring a broad range of aspects and views for this kind of visualization may be regarded as similar to architectural reconstruction. They analyzed and compared the comprehension of metrics in code reviews when aided by VR or 2D visualization. However, no results are available yet.

In our previous work [40], we developed and we empirically evaluated Microvision, an AR tool that uses the intermediate representation of the system built by the static analysis tool Prophet, which is capable of multi-codebases

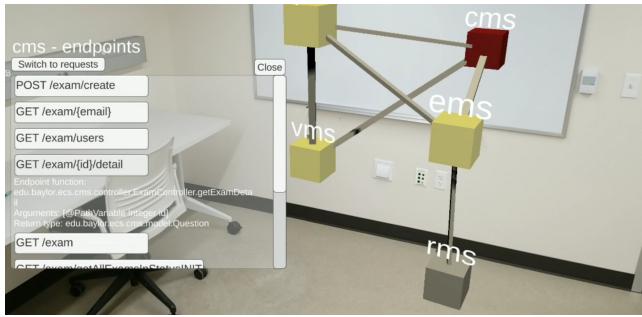


Fig. 6. Microvision Augmented Reality Visualization. The context menu shows a selected services API endpoints, in this case the “cms” service highlighted in red.

analysis for microservices to visualize the service endpoints of microservices, confirming the usefulness of AR for system understanding. An example of Microvision visualization is reported in Figure 6.

E. Microservice visualization in the industry

Microservice development comes from practitioners, and research tends to come later, so publications about microservices are still limited in a lot of areas. Thus, grey literature may hold valuable insights that academic literature simply cannot provide yet [52].

Amazon provides a solution called X-Ray console⁴. The provided approach is a map visual representation that consists of service nodes that serve requests, upstream client nodes that represent the origins of the requests, and downstream service nodes that represent web services and resources used by an application while processing a request (as depicted in Fig. 7). The X-Ray console provides embedded views that enable the user to view service maps and traces of applications' requests.

Netflix provides an interactive visualization technique for their system⁵. Fig. 8 shows the service graph representation of the system. It illustrates service dependencies in the whole system and enables the user to reconstruct the services communication graph to analyze different topologies. However, such a topology view is not particularly useful in debugging where a specific service is experiencing an issue.

Jaeger tracing⁶ is a common tool using dynamic analysis that provides Jaeger UI to render service dependencies. Fig. 9 shows a visual Directed Acyclic Graph (DAG) from Jaeger UI along with frequencies of calls. It can render a view to observe the system architecture.

When using the istio service mesh, a visualization is provided by the Kiali mesh visualization tool⁷. Kiali produces graphs representing traffic flowing through the service mesh for a period of time. There are several graph types provided, such as application, versioned application, workload, or service. Application type aggregates all versions of an app into a single graph node. The Versioned application graph

type shown at Fig. 10 shows a node for each version of an app, but all versions of a particular app are grouped together. The workload graph type shows a node for each workload in the service mesh. Finally, the service graph type shows a high-level aggregation of service traffic in the mesh.

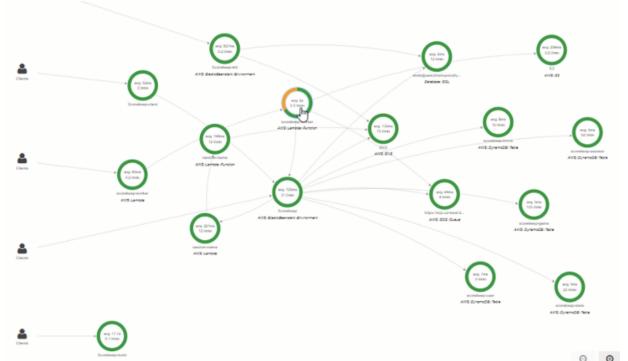


Fig. 7. Amazon X-Ray Console for Microservices Visualization

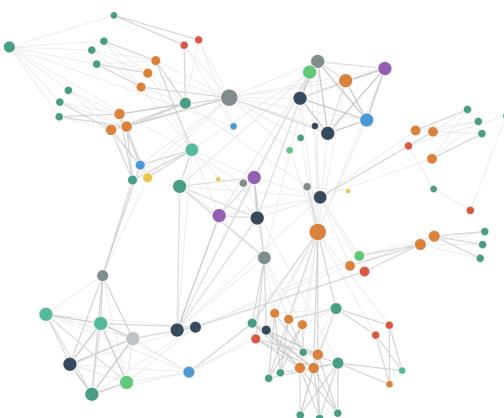


Fig. 8. Netflix Interactive Microservices Visualization

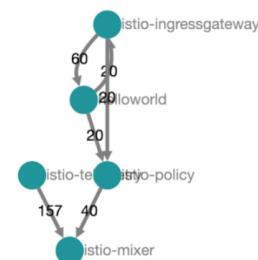


Fig. 9. Jaeger's directed acyclic graph with call frequencies

V. DISCUSSION

Different approaches have been proposed for SAR.

As for SAR-based on dynamic analysis, the industry seems to be a step forward compared to research. Industry proposed different visualization tools, mainly based on

⁴<https://aws.amazon.com/xray/>

⁵<http://simianviz.surge.sh/netflix>

⁶<https://www.jaegertracing.io>

⁷<https://istio.io/latest/docs/tasks/observability/kiali/>

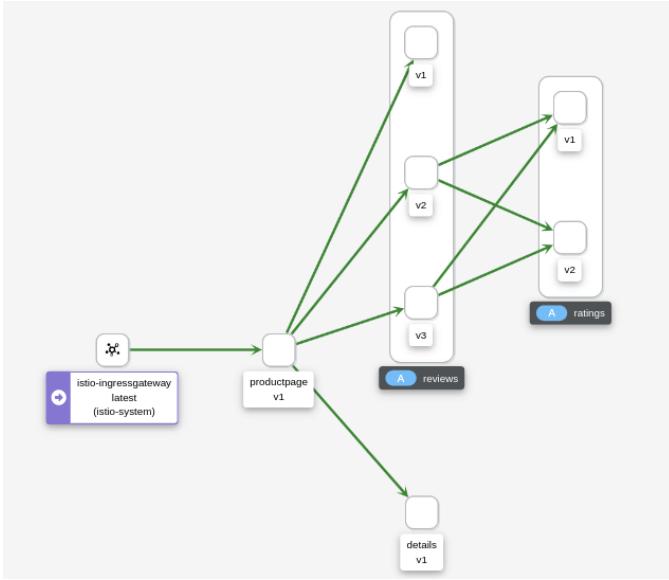


Fig. 10. Kiali traffic flowing through the service mesh

the service call-graph. However, no visualization provides insight into the quality of the architecture or its possible degradation.

Dynamic analysis might be very useful for making business decisions on the system and the development priority. As an example, companies might prioritize the maintenance of service more used by their customers; it would be possible to compare the maintenance effort with the actual usage of a service, but also to understand which service can be removed from the system. However, existing tools do not currently provide support for these decisions. However, An important lack of dynamic analysis tools is the support for architectural patterns [53], anti-patterns [54], and software metrics (e.g. Coupling [55]). For example, the service call-graph enables the detection of different patterns [56], anti-patterns and to calculate metrics such as coupling and cohesion. Therefore, we recommend tool providers introduce such features to provide better support to companies using their tools.

SAR based on static analysis received more attention from researchers, with little consideration from the practitioner's point of view. However, we would like to stress the importance of architectural reconstruction with static analysis. The main reason is that issues detected with static analysis could be immediately notified to the developers before the system goes into production. As an example, it could be possible to set quality gates in pull requests in case the developers introduce some architectural anti-pattern.

VI. CONCLUSIONS

This research was motivated by recurrent microservice system challenges regarding missing system-centric views. In addition, the microservices decentralization nature leads

to misaligned documentation and, consequently, architectural degradation, inefficiency, and broadening efforts.

Current approaches to determine the system-centric perspective prioritize dynamic analysis, targeting technology-neutrality. While it helps DevOps with their task, it does not necessarily fit developers as it requires dynamic system interaction. While user simulation tests can be developed, these easily degrade as the system evolves.

This work investigates the existing methods for static and dynamic architectural reconstruction and the tools adopted to visualize it.

Results show that static analysis is still not properly developed and might deserve special attention from the practitioner's point of view. Dynamic analysis, instead, is now widely used by several tools to visualize the service call-graph. However, there is no support for architectural degradation or investigation of quality issues in microservices.

Future work includes the development of extensions for dynamic analysis tools to detect architectural smells and to calculate software metrics such as coupling and cohesion. Moreover, we are planning to survey developers to understand their needs in terms of architectural degradation and further develop static and dynamic analysis tools based on developers' feedback.

ACKNOWLEDGMENTS

This material is based upon work supported by a National Science Foundation under Grant No. 1854049, a grant from Red Hat Research (<https://research.redhat.com>), and support from the Shapit Project (Ulla Tuominen Foundation - Finland).

REFERENCES

- [1] L. O'Brien, C. Stoermer, and C. Verhoef, "Software architecture reconstruction: Practice needs and current approaches," Carnegie Mellon University, Tech. Rep., 01 2002.
- [2] F. Rademacher, S. Sachweh, and A. Zündorf, "A modeling method for systematic architecture reconstruction of microservice-based software systems," in *Enterprise, Business-Process and Information Systems Modeling*, S. Nurcan, I. Reinhartz-Berger, P. Soffer, and J. Zdravkovic, Eds. Cham: Springer International Publishing, 2020, pp. 311–326.
- [3] A. Wiggins, "The twelve-factor app," 2017, (Accessed on 10/02/2021). [Online]. Available: <https://12factor.net/>
- [4] T. Cerny and D. Taibi, "Static analysis tools in the era of cloud-native systems," in *4th International Conference on Microservices*, 2022.
- [5] A. Al Maruf, A. Bakhtin, T. Cerny, and D. Taibi, "Using microservice telemetry data for system dynamic analysis," in *2022 IEEE Symposium on Service-Oriented System Engineering (SOSE)*, 2022.
- [6] M. Cinque, D. Cotroneo, R. D. Corte, and A. Pecchia, "A framework for on-line timing error detection in software systems," *Future Generation Computer Systems*, vol. 90, pp. 521 – 538, 2019. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167739X18309609>
- [7] S. Esparrachiari, T. Reilly, and A. Rentz, "Tracking and controlling microservice dependencies," *Queue*, vol. 16, no. 4, pp. 10:44–10:65, Aug. 2018. [Online]. Available: <http://doi.acm.org/10.1145/3277539.3277541>

[8] J. Thalheim, A. Rodrigues, I. E. Akkus, P. Bhatotia, R. Chen, B. Viswanath, L. Jiao, and C. Fetzer, "Sieve: Actionable insights from monitored metrics in distributed systems," in *Proceedings of the 18th ACM/IFIP/USENIX Middleware Conference*, ser. Middleware '17. New York, NY, USA: ACM, 2017, pp. 14–27. [Online]. Available: <http://doi.acm.org/10.1145/3135974.3135977>

[9] S. Ma, C. Fan, Y. Chuang, W. Lee, S. Lee, and N. Hsueh, "Using service dependency graph to analyze and test microservices," in *2018 IEEE 42nd Annual Computer Software and Applications Conference (COMPSAC)*, vol. 02, 2018, pp. 81–86.

[10] S. Ma, I. Liu, C. Chen, J. Lin, and N. Hsueh, "Version-based microservice analysis, monitoring, and visualization," in *2019 26th Asia-Pacific Software Engineering Conference (APSEC)*, 2019, pp. 165–172.

[11] B. Mayer and R. Weinreich, "An approach to extract the architecture of microservice-based software systems," in *2018 IEEE Symposium on Service-Oriented System Engineering (SOSE)*, 2018, pp. 21–30.

[12] K. A. Torkura, M. I. Sukmana, and C. Meinel, "Integrating continuous security assessments in microservices and cloud native applications," in *Proceedings of The 10th International Conference on Utility and Cloud Computing*, ser. UCC '17. New York, NY, USA: Association for Computing Machinery, 2017, pp. 171–180. [Online]. Available: <https://doi.org/10.1145/3147213.3147229>

[13] G. Granchelli, M. Cardarelli, P. D. Francesco, I. Malavolta, L. Iovino, and A. D. Salle, "Towards recovering the software architecture of microservice-based systems," in *2017 IEEE International Conference on Software Architecture Workshops (ICSAW)*, 2017, pp. 46–53.

[14] J. Carnell and I. H. Sánchez, *Spring microservices in action*. 2nd ed. Shelter Island, NY, USA: Manning Publications Co., 2021. [Online]. Available: <https://www.manning.com/books/spring-microservices-in-action-second-edition>

[15] F. Khan, "Microservices metrics visualization," 2020, tamper University. [Online]. Available: <https://urn.fi/URN:NBN:fi:tuni-202011157967>

[16] A. Bento, J. Correia, R. Filipe, F. Araujo, and J. Cardoso, "Automated analysis of distributed tracing: Challenges and research directions," *Journal of Grid Computing*, vol. 19, no. 1, p. 9, 2021. [Online]. Available: <https://doi.org/10.1007/s10723-021-09551-5>

[17] A. Chlipala, "The bedrock structured programming system: Combining generative metaprogramming and hoare logic in an extensible program verifier," *SIGPLAN Not.*, vol. 48, no. 9, pp. 391–402, Sep. 2013. [Online]. Available: <http://doi.acm.org/10.1145/2544174.2500592>

[18] E. Albert, M. Gómez-Zamalloa, L. Hubert, and G. Puebla, "Verification of java bytecode using analysis and transformation of logic programs," in *Practical Aspects of Declarative Languages*, M. Hanus, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 124–139.

[19] H. Cho, "Using metaprogramming to implement a testing framework," in *Proceedings of the 47th Annual Southeast Regional Conference*, ser. ACM-SE 47. New York, NY, USA: ACM, 2009, pp. 55:1–55:2. [Online]. Available: <http://doi.acm.org/10.1145/1566445.1566519>

[20] P. Tonella, "Evolutionary testing of classes," *SIGSOFT Softw. Eng. Notes*, vol. 29, no. 4, pp. 119–128, Jul. 2004. [Online]. Available: <http://doi.acm.org/10.1145/1013886.1007528>

[21] P. E. Papotti, A. F. do Prado, and W. L. de Souza, "Reducing time and effort in legacy systems reengineering to mdd using metaprogramming," in *Proceedings of the 2012 ACM Research in Applied Computation Symposium*, ser. RACS '12. New York, NY, USA: Association for Computing Machinery, 2012, p. 348–355. [Online]. Available: <https://doi.org/10.1145/2401603.2401681>

[22] I. Keivanloo, C. K. Roy, and J. Rilling, "Sebyte: Scalable clone and similarity search for bytecode," *Science of Computer Programming*, vol. 95, pp. 426 – 444, 2014, special Issue on Software Clones (IWSC'12). [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167642313002773>

[23] —, "Java bytecode clone detection via relaxation on code fingerprint and semantic web reasoning," in *Proceedings of the 6th International Workshop on Software Clones*, ser. IWSC '12. Piscataway, NJ, USA: IEEE Press, 2012, pp. 36–42. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2664398.2664404>

[24] D. Rattan, R. Bhatia, and M. Singh, "Software clone detection: A systematic review," *Information and Software Technology*, vol. 55, no. 7, pp. 1165 – 1199, 2013. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0950584913000323>

[25] S. Eski and F. Buzluca, "An automatic extraction approach: Transition to microservices architecture from monolithic application," in *Proceedings of the 19th International Conference on Agile Software Development: Companion*, ser. XP '18. New York, NY, USA: Association for Computing Machinery, 2018. [Online]. Available: <https://doi.org/ezproxy.baylor.edu/10.1145/3234152.3234195>

[26] I. Pigazzini, F. A. Fontana, V. Lenarduzzi, and D. Taibi, "Towards microservice smells detection," in *Proceedings of the 3rd International Conference on Technical Debt*, ser. TechDebt '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 92–97. [Online]. Available: <https://doi.org/10.1145/3387906.3388625>

[27] M. Rahman and D. Taibi, "A curated dataset of microservices-based systems," in *Joint Proceedings of the Summer School on Software Maintenance and Evolution*. CEUR-WS, September 2019.

[28] A. Ibrahim, S. Bozhinoski, and A. Pretschner, "Attack graph generation for microservice architecture," in *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing*, ser. SAC '19. New York, NY, USA: Association for Computing Machinery, 2019, pp. 1235–1242. [Online]. Available: <https://doi.org/ezproxy.baylor.edu/10.1145/3297280.3297401>

[29] I. Salvadori, A. Huf, R. d. S. Mello, and F. Siqueira, "Publishing linked data through semantic microservices composition," in *Proceedings of the 18th International Conference on Information Integration and Web-Based Applications and Services*, ser. iiWAS '16. New York, NY, USA: Association for Computing Machinery, 2016, pp. 443–452. [Online]. Available: <https://doi.org/ezproxy.baylor.edu/10.1145/3011141.3011155>

[30] L. Bass, P. Clements, and R. Kazman, *Software Architecture in Practice*, 3rd ed. Addison-Wesley Professional, 2012.

[31] P. Kruchten, "Architectural blueprints: The 4+1 view model of software architecture," *CorR*, vol. abs/2006.04975, 1995. [Online]. Available: <https://arxiv.org/abs/2006.04975>

[32] —, "The 4+1 view model of architecture," *IEEE Softw.*, vol. 12, no. 6, p. 42–50, nov 1995. [Online]. Available: <https://doi.org/10.1109/52.469759>

[33] A. Walker, I. Laird, and T. Cerny, "On automatic software architecture reconstruction of microservice applications," *Information Science and Applications: Proceedings of ICISA 2020*, vol. 739, p. 223, 2021.

[34] M. Fowler, "Bounded context," January 2014, (Accessed on 10/02/2021). [Online]. Available: <https://martinfowler.com/bliki/BoundedContext.html>

[35] E. Evans and E. J. Evans, *Domain-driven design: tackling complexity in the heart of software*. Addison-Wesley Professional, 2004.

[36] V. Vernon, *Implementing domain-driven design*. Addison-Wesley, 2013.

[37] C.-a. Sun, "A multi-view architectural model and its description and construction," in *2010 International Conference on Computational Intelligence and Software Engineering*, 2010, pp. 1–5.

[38] Z. Zhou, Q. Zhi, S. Morisaki, and S. Yamamoto, "A systematic literature review on enterprise architecture visualization methodologies," *IEEE Access*, vol. 8, pp. 96 404–96 427, 2020.

[39] A. Vázquez-Ingelmo, A. García-Holgado, and F. J. García-Péñalvo, "C4 model in a software engineering subject to ease the comprehension of uml and the software," in *2020 IEEE Global Engineering Education Conference (EDUCON)*, 2020, pp. 919–924.

[40] T. Cerny, A. Abdelfattah, V. Bushong, A. A. Maruf, and D. Taibi, "Microvision: Static analysis-based approach to visualizing microservices in augmented reality," in *2022 IEEE Symposium on Service-Oriented System Engineering (SOSE)*, 2022.

[41] M. Shahin, P. Liang, and M. A. Babar, "A systematic review of software architecture visualization techniques," *J. Syst. Softw.*, vol. 94, pp. 161–185, 2014.

[42] R. Wettel and M. Lanza, "Visually localizing design problems with disharmony maps," in *Proceedings of the 4th ACM Symposium on Software Visualization*, ser. SoftVis '08. New York, NY, USA: Association for Computing Machinery, 2008, p. 155–164. [Online]. Available: <https://doi.org/10.1145/1409720.1409745>

[43] F. Fittkau, A. Krause, and W. Hasselbring, "Exploring software cities in virtual reality," in *2015 IEEE 3rd Working Conference on Software Visualization (VISSOFT)*, 2015, pp. 130–134.

[44] A. Schreiber, L. Nafeie, A. Baranowski, P. Seipel, and M. Misiak, "Visualization of software architectures in virtual reality and augmented reality," in *2019 IEEE Aerospace Conference*, 2019, pp. 1–12.

[45] M. Steinbeck, R. Koschke, and M. O. Rüdel, "How evostreets are observed in three-dimensional and virtual reality environments,"

in *2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2020, pp. 332–343.

[46] R. Oberhauser and C. Pogolski, “VR-EA: Virtual Reality Visualization of Enterprise Architecture Models with ArchiMate and BPMN,” in *Business Modeling and Software Design*, B. Shishkov, Ed. Cham: Springer International Publishing, 2019, vol. 356, pp. 170–187, series Title: Lecture Notes in Business Information Processing.

[47] Z. Ma and Y. Bai, “A distributed system monitoring tool with virtual reality,” in *Proceedings of the 2nd International Conference on Computer Science and Application Engineering*, ser. CSAE ’18. New York, NY, USA: Association for Computing Machinery, 2018.

[48] I. Toumpalidis, K. Cheliotis, F. Roumpani, and A. Hudson Smith, “Vr binoculars: An immersive visualization framework for IoT data streams,” in *Proceedings of the IEEE Living in the Internet of Things: Cybersecurity of the IoT*, 2017.

[49] H. Halpin, D. J. Zielinski, R. Brady, and G. Kelly, “Exploring semantic social networks using virtual reality,” in *The Semantic Web - ISWC 2008*, A. Sheth, S. Staab, M. Dean, M. Paolucci, D. Maynard, T. Finin, and K. Thirunarayanan, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 599–614.

[50] S. Royston, C. DeFanti, and K. Perlin, “A collaborative untethered virtual reality environment for interactive social network visualization,” *CoRR*, vol. abs/1604.08239, 2016. [Online]. Available: <http://arxiv.org/abs/1604.08239>

[51] D. Moreno-Lumbrales, G. Robles, D. Izquierdo-Cortázar, and J. M. Gonzalez-Barahona, “To VR or not to VR: Is virtual reality suitable to understand software development metrics?” 2021. [Online]. Available: <https://arxiv.org/abs/2109.13768>

[52] J. Bogner, J. Fritzs, S. Wagner, and A. Zimmermann, “Industry practices and challenges for the evolvability assurance of microservices,” *Empirical Software Engineering*, vol. 26, no. 5, p. 104, 2021. [Online]. Available: <https://doi.org/10.1007/s10664-021-09999-9>

[53] D. Taibi, V. Lenarduzzi, and C. Pahl, “Architectural patterns for microservices: A systematic mapping study,” in *Proceedings of the 8th International Conference on Cloud Computing and Services Science - Volume 1: CLOSER*, INSTICC. SciTePress, 2018, pp. 221–232.

[54] ———, “Microservices anti-patterns: A taxonomy,” in *Microservices: Science and Engineering*. Springer International Publishing, 2020, pp. 111–128.

[55] S. Panichella, M. R. Imranur, and D. Taibi, “Structural coupling for microservices,” in *11th International Conference on Cloud Computing and Services Science*, 04 2021.

[56] A. Bakhtin, A. Al Maruf, T. Cerny, and D. Taibi, “Survey on tools and techniques detecting microservice api patterns,” in *IEEE International Conference on Services Computing (SCC)*, 2022.