

# Development Frameworks for Microservice-based Applications: Evaluation and Comparison

Hai Dinh-Tuan  
Technische Universität Berlin  
Ernst-Reuter-Platz 7  
10587 Berlin  
hai.dinhtuan@tu-berlin.de

Maria Mora-Martinez  
Technische Universität Berlin  
Ernst-Reuter-Platz 7  
10587 Berlin  
m.moramartinez@tu-berlin.de

Felix Beierle  
Technische Universität Berlin  
Ernst-Reuter-Platz 7  
10587 Berlin  
beierle@tu-berlin.de

Sandro Rodriguez Garzon  
Technische Universität Berlin  
Ernst-Reuter-Platz 7  
10587 Berlin  
sandro.rodriguezgarzon@tu-berlin.de

## ABSTRACT

The microservice architectural style has gained much attention from both academia and industry recently as a novel way to design, develop, and deploy cloud-native applications. This concept encourages the decomposition of a monolith into multiple independently deployable units. A typical microservices-based application is formed of two service types: *functional services*, which provide the core business logic, and *infrastructure services*, which provide essential functionalities for a microservices ecosystem. To improve developers' productivity, many software frameworks have been developed to provide those reusable infrastructure services, allowing programmers to focus on implementing microservices in arbitrary ways. In this work, we made use of four open source frameworks to develop a cloud-based application in order to compare and evaluate their usability and practicability. While all selected frameworks promote asynchronous microservice design in general, there are differences in the ways each implements services. This leads to interoperability issues, such as message topic naming convention. Additionally, a key finding is the long startup times of JVM-based services that might reduce application's resiliency and portability. Some other advantages come directly from the programming language, such as the ability of Go to generate native binary executables, which results in very small and compact Docker images (up to 78% smaller compared to other languages).

## CCS Concepts

• **Software and its engineering** → **Software design tradeoffs; Software design engineering; Development frameworks and environments.**

## Keywords

microservices; software engineering; framework-based software development.

## 1. INTRODUCTION

The microservice architectural style emerged from the *Service-oriented architecture* and the *Domain-Driven Design* architectural patterns with a strong focus on DevOps practices [15]. Not only has it been successfully adopted in multimillion-user companies such as Netflix or Amazon, but also received significant attention from the academia as a robust and scalable method to build cloud applications. In short, the microservice architecture advocates a new software design, in which the application is structured into services. These services are independently developed, updated, and deployed, thus significantly improving the flexibility of the application. Similarly, the overall scalability and fault tolerance of the application are enhanced. However, these advantages also come at a cost. Applications get more complex and latency increases because the communication channel is more susceptible to network delays.

To reduce the burdens created by distributed systems' complexities, several frameworks have been developed specifically for microservices. They provide boilerplate code and abstractions to allow the developers to focus on the business logic. By exploring implementation options for an example application, this paper proposes three criteria for evaluating and comparing frameworks. For *features*, we compare frameworks across three main categories, which are maturity, development supports, and built-in features. For *design patterns*, we categorized supported ones into six groups. Lastly, for *performance*, we relied on two main criteria, which are end-to-end latency performance and resource consumption. This paper therefore delivers a three-fold contribution:

1. A review of different approaches to implement microservices, which is exhibited in four microservice frameworks selected.
2. A comparison model for microservices frameworks, based on three main criteria: features, design patterns, and performance.
3. The implications of adopting a framework in developing microservices.

In the remainder of this paper, Section 2 gives a brief account of the evaluations and comparisons have been done so far for microservices. Then, we introduce the motivating use case and the architecture for a toll system's prototype in Section 3. Section 5

presents a feature-wise comparison while Section 6 shows practical performance results. We discuss the results in Section 7, draw the overall conclusion in Section 8, and list future research directions in Section 9.

## 2. RELATED WORK

One of the most well known definitions of software framework proposed by Johnson describes it as "a reusable design of all or part of a system that is represented by a set of abstract classes and the way their instances interact" or "the skeleton of an application that can be customized by an application developer" [12]. With a software framework, developers are able to reuse both design and code to focus on solving business specific needs [13]. Therefore, framework-based software development can be seen as an activity to modify and extend the framework [19]. In addition, a framework can improve programmer productivity and enhance the quality, performance, reliability and interoperability of software [7, 13, 17]. However, a software framework can also lead to prolonged development time, as developers have to learn and get familiar with the framework [23]. Also, adopting a framework might result in more complicated code structure and larger software footprints.


As frameworks are usually designed to solve a narrow set of problems by targeting a certain type of applications [8], there are more and more frameworks specialized to ease, speed up and improve development of applications with a microservice-based architectural design. There have been several works study the performance of microservices. [2, 25]. Others propose criteria to evaluate microservices such as [1, 24]. However, to the best of our knowledge, there is no research article that compares the performance and approaches by using different programming languages and frameworks to build the same application. This paper addresses this exact gap to compare and evaluate different framework-based development approaches for microservices.

## 3. MOTIVATING USE CASE

One of the biggest issues in urban areas is the decline of air quality, which causes 4.2 million premature deaths every year, according to a recent report by the World Health Organization (WHO) [20]. In the same report, human mobility activities has been identified as the major sources of outdoor air pollution, including fossil fuels combustion from motor vehicles.

In many countries, policy-makers have introduced various regulations to address these problems, one of them is to implement a pollution toll. However, current toll schemes charge fixed fees based on static criteria such as vehicle type or toll area [26]. In this work, inspired by [10], we propose a novel fine-grained system to calculate toll fee proportionally to the pollution level of each area. This means the more polluted an area is, the higher fee per driven kilometre is applied for all vehicle in that area. One obvious advantage of this approach is that this system allows a dynamic toll rate based on real-time air pollution level. In addition, this is very scalable, as it only uses location data of the vehicle in short intervals to calculate the toll fee. The fairness of this system can be further improved by adding information about the vehicle's model and speed (estimated from the location data).

The hypothetical toll system adopts the event-driven microservice architecture, in which each microservice is responsible for a specific task and data is exchanged as JSONs through a message broker using the publish-subscribe asynchronous messaging pattern.



**Figure 1. Architecture overview of the proposed airpollution-aware toll system.**

Fig. 1 illustrates the architecture, in which the arrows among microservices are not direct communication but rather an indirect communication via a message broker. The processing pipeline is explained in detail below:

- The vehicles update their GPS coordinates information by sending JSON-based messages via the message broker.
- The *Map Matcher Service* subscribes to *Location update messages* from the vehicles. Then it uses the APIs provided by a local instance of OSRM (Open Source Routing Machine) [11] to map the reported coordinates to a route on streets. Then the driven route will be sent to the message broker in the *Route Messages*.
- The *Pollution Matcher Service* receives the *Route messages* and compares them with an air pollution database. In this PostGIS database, polluted areas are stored as polygons with corresponding pollution levels. The received *route* is matched against these polygons and in case the route crosses multiple areas with different pollution levels, it will be broken down into several segments. These segments together with corresponding pollution levels are sent by the *Pollution Matcher Service* to the message broker via *Segment Messages*.
- The *Toll Calculator Service* calculates the fare for each vehicle based on the *Segment Messages*. It calculates the travel distance of each segment and, based on the pollution level, returns the accumulative toll fee for each vehicle. Those values are sent to the message broker as *Toll Messages*.
- Upon receipt of the *Toll Messages*, the *Dashboard* displays the cumulative toll fee for each car and visualizes each vehicle's route on an interactive map.

## 4. FRAMEWORKS

There is a large availability of frameworks for different programming languages. For the implementation of the toll system, we selected four different microservices frameworks: Go Micro (Go), Moleculer (JavaScript with Node.js), Lagom (Scala), and Spring Boot/Spring Cloud (Java). They are chosen to reflect the variety of tools that can be used to develop microservices. With respect to programming languages, Java is consider an "old" language while Scala and Go both have been introduced in the last decade, and Node.js is a new run-time environment for JavaScript. In terms of framework design philosophy, although Spring Boot is widely used for microservices, it is not designed specifically for such applications like the other three. Each framework provides a different approach to build microservice-based applications. In this section, we briefly discuss the main differences.

*Go Micro* provides an abstraction for distributed systems development using Remote Procedure Call (RPC) and Event-

driven communication written in Go. This opinionated framework is provided as a set of pluggable building blocks that can be included or excluded in a project depending on the application's specific requirements. This design reduces the complexity in development. They cover a wide range of functionalities, from sync/async messaging, message encoding, service discovery, to dynamic configuration. To define a microservice with Go Micro, the framework encourages developers to firstly define the interfaces (default by ProtoBuf), and then implement them. Afterwards, the service can be initialized and registered with the built-in service registry.

*Lagom* is another opinionated microservice framework to build reactive microservices in Scala or Java. Developed and maintained by the same company behind Akka and Play Framework, Lagom is built to be easily integrated into these two technologies. Lagom supports and encourages using Event sourcing and Command Query Responsibility Segregation (CQRS) for data persistence. Lagom recommends developers to separate the service interface and service implementation, which are organized as sub-projects within the same build. While the former contains the interface descriptions, along with data models, the latter implements the service logic.

*Moleculer* is a microservice framework for JavaScript using Node.js runtime built on Chrome's V8 JavaScript engine. The main design philosophy of this framework is to produce low latency event-driven applications. Moleculer also provides pluggable modules for various functions such as gateways, database, serialization, etc. The framework consists of three main components: the transporter, the service and the service broker. While the service defines the general structure for each microservice, the transporter is an abstraction for the communication bus, and the service broker binds service with transporter.

*Spring Boot & Spring Cloud* are two projects from the Spring ecosystem. While Spring Boot is an opinionated way to build standalone Spring applications, Spring Cloud provides a set of tools to quickly build several common patterns in distributed systems such as service discovery, configuration management, etc. That means Spring Boot is not specifically designed for developing microservices but rather a full stack development environment. Spring Boot has to rely on Spring Cloud for microservices-specific capabilities.

## 5. FEATURE COMPARISON

In this section, we provide an extensive evaluation of those frameworks with respect to 2 main criteria: *features* and *design patterns*. For the former, we further sub-categorize it into (1) Maturity, (2) Supports for microservice development, and (3) Built-in features. All of these details will be summarized in Table 1 and Table 2. Note that, a "N/A" does not necessary mean that the framework lacks a certain feature, but rather that the feature requires external libraries, which has not yet been officially integrated into the framework.

### 5.1 Main features and design philosophy

All four microservice frameworks focus on providing infrastructure services for microservice-based applications. In fact, the separation between infrastructure services and functional services has been well discussed in the literature [5]. While infrastructure services provide essential capability such as service registry or service discovery, functional services are the ones that provide actual business logic.

As observed in the four implementations, there are two main approaches from the different frameworks: they can either implement their own services or provide integration with existing solutions. In addition, as open source projects, besides the official code, they can provide community-maintained plugins and modules such as in Go Micro and Moleculer.

### 5.2 Maturity

To quantify the maturity of a open sourced software framework, we take a look at the release cycle and metrics from open source websites for code hosting such as GitHub as used in previous research [6, 18].

Given that Spring Boot has been initially releases in 2014, while Lagom's first release was in 2016 and both Go Micro and Moleculer in 2017, Spring Boot has become one of the most mature frameworks. That also explains how the amount of releases from Spring outperforms the others twice, as seen in Table 1. However, as Spring Boot is also a popular framework to build generic Java applications, this might not fully reflect the maturity of Spring in developing microservices specifically.

### 5.3 Supports for microservice development

Given the difference in maturity among the selected frameworks, one can expect different levels of support and learning curves. Indeed, the documentation of Spring Boot is very extensive compared to the others. With a huge community and long existence on the market, it is easy to find support from developer communities such as StackOverflow. Besides, Spring and Lagom are backed by Pivotal and Lightbend respectively, so they can offer commercial support for large organizations, while Moleculer and Go Micro must rely on community-based supports. Build tools and dependency management solution plays an important roles to significantly improve the productivity. Boilerplate codes for new microservices can be quickly generated following certain templates, however each framework supports different level of customization. For example, Lagom and Moleculer only allow to set up project name and service name, while the other two allow specifying dependency packages. In addition, features such a hot reloading or interactive development console can further shorten development cycles, as they help developers to immediately see the effects of new code at runtime.

### 5.4 Design patterns

To achieve robustness in microservices, there are certain recommendations to follow, such as minimizing service dependencies or separating data storage [21]. These non-functional properties materialize in different design patterns. To understand how each selected framework design philosophy differs from each other, we briefly discuss them in this section and provide a summary in Table 2.

To provide API access for external users, all four frameworks support the API Gateway pattern. In this design pattern accesses to individual microservices are not directly provided to the user. Instead, a gateway is used as a single entry point for the application, exposing APIs, routing messages to responsible microservices, or even translating messaging between different protocols, such as in Go Micro.

Another design pattern advocated by all frameworks is the *Database per service* design pattern. Microservice design helps

Table 1. Features of microservices frameworks.

		Spring	Lagom	Moleculer	Go Micro	
Maturity	GitHub stats	Initial release date	July 10th 2014	March 3rd 2016	February 16th 2017	December 6th 2017
		Releases	146	76	86	70
		Commits	25357	2397	3222	2700
		Watch	3348	165	123	452
		Star	45626	2370	3322	11748
	Contributors	656	126	65	77	
Development supports	Documentation	Documentation	Very Extensive	Extensive	Less Extensive	Limited
		Sample code	Yes	Yes	Yes	Yes
	Supported languages	Supported languages	Java, Kotlin, Groovy, Beanshell	Java, Scala	Javascript/NodeJs	Go
		Community supports	Stack OverFlow tags	162688	296	38
	Chat platform		Gitter	Gitter	Gitter	Slack
	Forum		Yes	Yes	Yes	Yes
	GitHub		Yes	Yes	Yes	Yes
	Commercial supports	Training course	Yes	Yes	No	No
		Certification	Yes	Yes	No	No
	Development environment	Build tool & Dependency Management	Maven (default), Gradle, Ant	sbt	Default to npm	Default to Go Build
		Hot reload	Yes, but not by default	Yes, supported by default with sbt	Yes, enabled by default	N/A
		Template generator	Spring Initializr (web-based)	Lagom Tech Hub Project Starter (web-based)	moleculer-cli command	micro new command
		Interactive development console	N/A	sbt development console, only for development purposes	REPL console: list nodes, services, actions, events, environment variables, etc.	Interactive Micro CLI, to manage the services
Built-in features	Essential services	API Gateway	Yes	Yes	Yes	Yes
		Service discovery	Yes	Yes	Yes	Yes
		Service registry	Yes	Yes	Yes	Yes
		Load balancing	Spring Cloud Load Balancer	Relies on underlying Akka	server-side load balancing by default	proxy package
		Serialization	JSON by Jackson	JSON (default), ProtoBuf, Apache Avro, Kryo	JSON (default), Avro, MsgPack, Notepack, ProtoBuf, Thrift, and custom modules	JSON or ProtoBuf
	Other features	Integration with Slack	N/A	N/A	N/A	bot package
		Interactive CLI	Spring Boot CLI	sbt	moleculer REPL	Micro CLI package
		Web dashboard	N/A	N/A	N/A	web package
		Caching	JCache, EhCache 2.x, Hazelcast, Infinispan, Couchbase, Redis, Caffeine, etc.	N/A	built-in cache, heap memory, LRU cache module, Redis-based, or custom cacher	N/A

**Table 2. Design patterns of microservice frameworks.**

		Spring	Lagom	Moleculer	Go Micro
Integration	API Gateway	Spring Cloud Gateway, built on top of Spring MVC, Netty	Default embedded gateway based on Akka HTTP	moleculer-web for request routing	Go API libraries for protocol translation and request routing
	Database				
Database	Database per service	Supported	Cassandra, MySQL, PostgreSQL, Oracle, Couchbase, Microsoft SQL, H2	MongoDB, MySQL, PostgreSQL, SQLite, MSSQL. Default to CRUD actions	Only provides a Store interface for additional storage in the future
	Shared database	Supported	Discouraged	Discouraged	N/A
	CQRS/Event Sourcing	Supported	Default	N/A	N/A
Observability	Log Aggregation	Logback(default), Java Util Logging, Log4J2	SLF4J based on Logback (default), Log4J2	Console output, Pino, Bunyan, Winston, Log4js, Datadog	Logrus, Zap, Zerolog
	Performance Metrics	Prometheus, Datadog, Netflix Atlas	Provided by Lightbend platform	Prometheus, StatsD, Datadog	Prometheus
	Distributed Tracing	Supported with Spring Cloud Sleuth	Provided by Lightbend, which supports OpenTracing	Built-in tracing module works with Zipkin, Jaeger, Datadog	Amazon AWS X-Ray, UUID plugin to trace messages manually
	Health Check	Supported with Spring Actuator	N/A	Built-in	Sidecar implementation via RPC
Cross-cutting concerns	External configuration	Spring Cloud Netflix for client and server-side, Netflix Archaius	N/A	External configuration files can be loaded via moleculer-runner	Supported via Go Config
	Service Registry	Netflix Eureka. Also Kubernetes, Consul, Apache Zookeeper, Alibaba Nacos	Provided by default	Built-in service registry	Built-in, can be accessed via a Web Dashboard
	Client-side service discovery	Client-side load balancing supported with Ribbon	Provided in the ServiceLocator	N/A	mdns, etcd
	Server-side service discovery	N/A	Supported	Default	mdns, etcd
	Deployment platform	Docker, Kubernetes, Cloud Foundry, Azure, Heroku, Amazon AWS, OpenShift, Boxfuse	Docker, Kubernetes, OpenShift, MesosphereDC/OS	Docker/Docker Compose, Kubernetes under development	Local, Docker, Kubernetes
Communication	Synchronous messaging	Supported	Supported	Supported	Supported
	Asynchronous messaging	RabbitMQ, Apache Kafka, Amazon Kinesis, Google PubSub	Message Broker API for Kafka	Preferred: NATS, Kafka, Redis, MQTT, AMQP	PubSub is preferred. Default embedded NATS server
	Remote Procedure Invocation	RMI, Spring's HTTP Invoker, Hessian, Burlap	Akka gRPC library HTTP/JSON	RPC is implemented by default	Supports gRPC
Fault Tolerance	Circuit Breaker	Supported with Netflix Hystrix	Used for all service calls by default	Built-in, thresholdbased	Plugins supported
	Retry	Spring-retry, Retryable annotation	N/A	Exponential back-off retry	Resend the requests to another node
	Timeout	For HTTP components	Circuit breaker: call and reset timeouts	Set timeout values for requests	N/A

decoupling services as they can be developed, deployed, and updated without breaking other services. However, data can create hidden dependencies among services. Therefore, each microservice should manage its own data and no service can modify data belonging to another service. In contrast, shared databases among multiple services is not encouraged by any framework. A more advanced pattern, such as Command-Query Responsibility Segregation with Event Sourcing, is even part of the framework philosophy like in the case of Lagom.

In an environment with multiple services being independently deployed, it is crucial for the operators to have an overview of the application and insights into underlying operations [3]. Therefore, all four frameworks provide extensive integration with different logging tools. However, when it comes to metrics collecting, while Moleculer has a built-in module collecting and publishing metrics to various outputs, Spring Boot and Go Micro provide a metrics collector as a pluggable module; Lagom has to rely on Lightbend cloud platform for those capabilities. This is because Lagom has been designed to simplify microservices application development and deployment on Lightbend’s current offer in cloud platform.

Similarly, a technique used for monitoring and analyzing distributed systems called *Distributed Tracing* is also supported by all frameworks. In this technique, as a request travels through the system, a unique identifier is added, allowing details about the path, time spent in each services to be centrally captured and visualized. Although no framework provides its own implementation of a monitoring service supporting Distributed Tracing, Spring Boot, Moleculer, and Go Micro are fully integrated with external solutions such as Prometheus. In contrast, Lagom microservices have to rely again on Lightbend’s platform.

With the recent growth of container-based cloud platforms, Docker and Docker-based orchestration platforms like Kubernetes, Mesosphere DC/OS, OpenShift, etc., have become the de facto standard. All of the frameworks have supports for Docker deployment, except Moleculer’s support for Kubernetes still missing. Both Spring Boot and Lagom can build executables that can be deployed on Heroku, Azure, etc.

Regarding the communication patterns among microservices, all of the frameworks encourage using asynchronous messaging as the main way for exchanging information. Asynchronous messaging can decouple services, minimizing dependencies, and improve the fault tolerance of the application as a whole [22]. In fact, being able to operate asynchronously is one of the design criteria to get the most out of the microservice architectural style [9]. While most of the widely used protocols such as AMQP, MQTT, are fully supported by all, each framework implements the communication channel differently: Spring, Lagom, and Go Micro provides different integration modules for message brokers, while Moleculer differentiates two communication types. *Local services* use an inmemory local bus for exchanging messages. *Remote services* relies on an abstraction layer called *Transporter*, which can deliver events, requests, and responses. By default, NATS is used as the broker, however, other implementations cannot be added directly, because Moleculer automatically adds prefixes to messages topics.

In a microservices environment with a large number of distributed services, it is utmost crucial to prevent cascading failures. All four frameworks support the *Circuit Breaker* pattern. In this design, in case of failure, requests are rejected immediately, thus avoiding errors.

## 6. PERFORMANCE EVALUATION

### 6.1 Test setup

We implemented the pipeline as described in section 3 with each framework. For the evaluation test, the following versions of the tools and frameworks have been used: Spring Boot 2.2.4 (Java 1.8), Go Micro 1.18.0, (Go 1.13), Moleculer 0.13.0 (Node.js 13.1.4), Lagom 1.6.0 (Scala 2.12.8), NATS 2.1.2. To ensure the consistency of input data across tests, we built a *Traffic Simulator* to simulate vehicles. This component is based on the Simulation of Urban MObility (SUMO) [4], and provides a traffic simulation in a European city using data from the *OpenStreetMap*. From the *Dashboard*, the simulation can be configured by the number of simulated vehicles and update frequency. The results presented in this paper is from simulations with 10 vehicles, and they update their location every 5 seconds.

The machine where the tests were carried out has an Intel Core i5 processor and 16GB of RAM. All microservices are containerized using Docker and deployed using Docker Compose, as this is a popular choice to achieve platform-independent deployment strategy [25]. We have considered latency and resource consumption as the key criteria to measure the performance of the different frameworks. Another resource metric is the footprint of the application, given by the size of the Docker images.

The selected usecase represents some typical microservices types, which should be taken into account when application being evaluated. The *Map Matcher* portrays a service which consumes external API, the *Pollution Matcher* constantly exchanges data with an external database, while the *Toll Calculator* performs all the calculation internally.

### 6.2 Slow start of JVM-based implementations

We noticed long latencies in the first messages processed in Spring Boot and Lagom implementations compared to subsequent requests as illustrated in Fig. 2.




Figure 2. Long latency for first messages in Lagom.

The reason for such a strange behavior, which is not found in other implementations, lies in the Java Virtual Machine (JVM) warm up phase [16]. In this test, to reduce the effects of the warm up phase, we only evaluate the framework performance after five message iterations, i.e. skipping the first 50 messages.

### 6.3 End-to-end latency performance

Fig. 3 summarizes the end-to-end latency of each framework. All four implementations show similar behavior in regard to latency,

although to different extents. There are still some noticeable patterns shared by all implementations. For example, the slowest microservice is always the *Pollution Matcher* that has to access the PostGIS database, it accounts from 73.27% in Lagom to 90.59% of the total time in Go Micro implementation. By contrast the *Toll Calculator* has a minimal contribution to the total end-to-end time since it can perform all the processes internally. From the results, it is clear that database access is very slow compared to other I/O bound tasks. This illustrates difficulties in managing service's state, in which they should be externalized into a database to create totally stateless services. In many designs, the microservices are forced to follow the 12-factor app principle, i.e., stateless, but generally it is either very difficult or not optimized to build applications solely from stateless microservices.

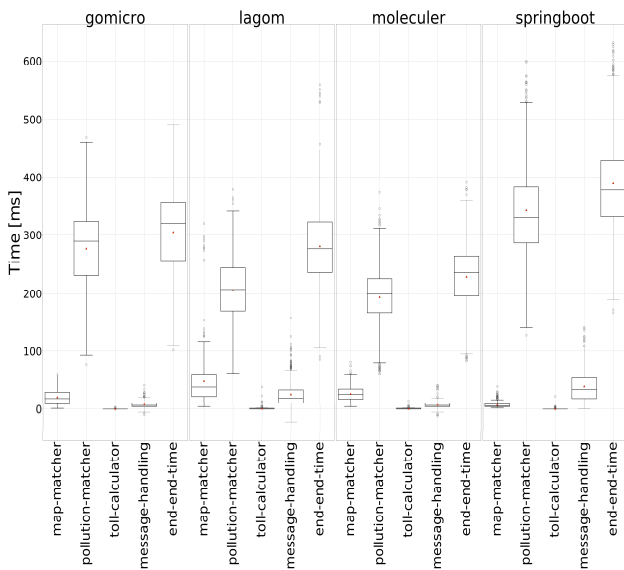


Figure 3. Comparison of latency performance.

Communication overhead is more dispersed for Lagom and Spring, while Go Micro and Moleculer have a more consistent delay. Even with slower communication, Lagom performs in general better than Go Micro.

## 6.4 Resource consumption

### 6.4.1 Docker images size

We evaluated the Docker image size produced by each framework, as smaller images can be deployed and created quickly, leading to faster update cycles [14]. In this setup, we used Alpine-based images: Spring uses *openjdk:8-jre-alpine*, Lagom uses *openjdk:8-jre-alpine*, Moleculer uses *mhart/alpine-node:10*, and Go Micro uses *golang* base image as build image and *scratch* as deployment container.

The results are summarized in Fig. 4. As Go applications can be built as binary files, they can use *scratch* - a minimal empty Docker image. Therefore services developed with this language are very compact compared to others, for example, the *Pollution Matcher* written in Go Micro is 78% smaller the implementation with Lagom (33.6 MB vs. 156 MB). Moleculer also has the ability to create smaller Dockerized services compared to Lagom or Spring Boot, as these two frameworks require a JVM to run the source code.

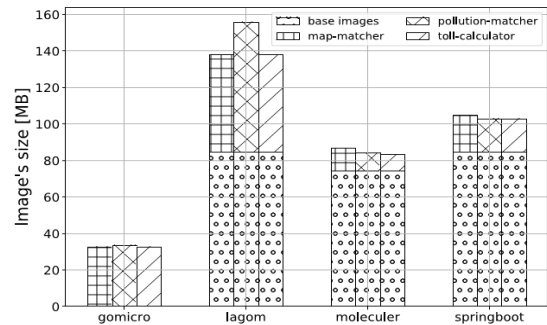


Figure 4. Docker images size comparison.

### 6.4.2 CPU and memory usage

As visible in Fig. 5, Lagom and Spring yield much higher resource consumption than the other two frameworks. It is interesting that while Lagom has a higher CPU usage on the Map Matcher service than Spring, and lower than for the Pollution Matcher, in the case of memory usage the results are reversed. Looking again at Fig. 3, the latency for the Map Matcher is the lowest with Spring, but this framework also shows the highest latency in the Pollution Matcher. Lagom obtains a slightly lower total end-to-end latency than Go Micro, but with a much higher resource consumption, so it is important to look at both aspects when deciding for a framework.

The lower CPU and memory consumption and small image size of the Go Micro services make it the most efficient framework in resource usage out of the four.

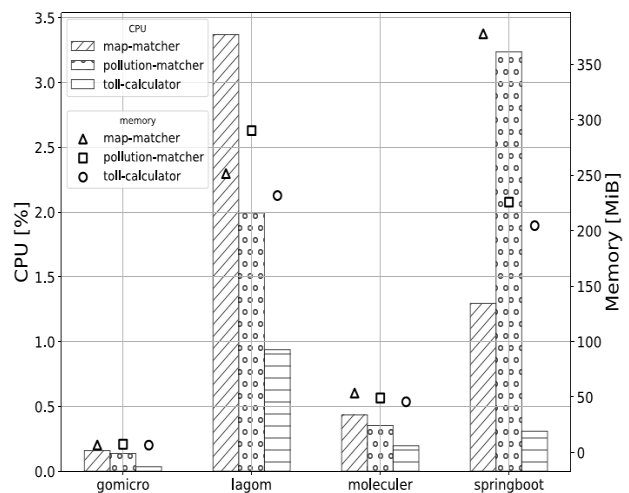


Figure 5. CPU and memory consumption of microservices implemented by different frameworks.

## 7. DISCUSSION

Three out of four frameworks studied in this work (Lagom, Moleculer, and Go Micro) encourage developers to structure the communication protocol before developing actual microservices. In contrast, Spring Boot doesn't force developers to follow a certain code style to define application's API. One of the possible reasons is that Spring Boot is a general-purpose framework, not designed specifically for microservices. This fact emphasizes the importance of communication pattern in distributed systems in general and in microservices in particular.

In theory, service decoupling is described as one of the most important features of microservices. However, in reality, there are several challenges to realize total independence among microservices. For example, shared libraries among microservices improve productivity of developers but also create dependencies. In our implementation, JSON is selected as the message format exchanged among services because it does not create specific requirements for microservices in serialization/deserialization. If Protocol Buffer is used instead, it forces every microservice to use ProtoBuf. However, given that ProtoBuf implementations are available for many programming languages, developers still have many options to develop their services. In another extreme end, using some platform-specific protocols such as *Spring HTTP Invoker* forces both server and client to be written in a JVM language.

For a cloud native architecture such as microservices, interoperability is becoming more significant than ever. Adopting a microservice framework using standards or popular tools can be much more beneficial for an organization in the long run. In this regard, as one of the most popular software frameworks, Spring Boot/Spring Cloud provides developers a clear advantage to build microservices as this framework has an extensive level of support and integration available. Both Go Micro and Moleculer provide less opinionated approach to build microservices. Lagom relies heavily on Lightbend platform and Akka, so it might require additional efforts to deploy on other platforms.

The development of microservices in particular, or cloud native applications in general, relies on many different levels of abstraction, which can degrade the performance significantly. Software architects and developers should also keep that in mind when selecting a technology stack for an application. For example, our results show that Moleculer can achieve a good end-to-end latency performance, but when deploying using Docker, Go Micro can produce much smaller images, which could be beneficial on limited resources. Similarly, JVM-based languages can lead to bigger Docker images, and higher resource consumption. Together with the JVM warm-up issue, it might not be the best choice for latency-sensitive application.

## 8. CONCLUSION

In this work, we presented an evaluation of microservice frameworks through the development of an example application, which includes both internal and external communication, data query from external database, and uses Docker for deployment. We defined three main comparison categories: For *features*, we examined three categories, which are maturity, development supports, and built-in features; for *design patterns*, we categorized them into six groups as presented in Table 2; for *performance*, we relied on two main criteria, which are end-to-end latency performance and resource consumption. The criteria mentioned in our evaluation can be extended to use as a comparison framework for selecting a development stack for a microservice-based application.

From the results presented in this work, it is clear that there is no one-size-fits-all solution when choosing a microservice framework. It depends a lot on application's specific functional and nonfunctional requirements. All of the frameworks mentioned in this paper as well as most of other options available can provide abstraction and boilerplate code to help developers deal with the complexities commonly found in distributed systems such as microservices. As shown in the evaluation section, interoperability is one of the most important criteria to maximize the flexibility in

tool selection. For example, the fact that Lagom is designed to heavily rely on Akka and Lightbend platform, or Moleculer uses a custom naming convention for message queues, might cause integration issues with other microservices. Regarding design patterns, all frameworks follow asynchronous, event-driven architecture with supports for other capabilities such as distributed tracing or fault tolerance. When it comes to performance comparison, our results show how using JVM-based languages can result in heavier Docker images and higher resource consumption.

The design of our test is, however, still limited to provide a comprehensive comparison among selected frameworks. Although we have tried our best to have a similar implementation in different frameworks/languages, there are still some specific features that cannot be easily translated into other implementations. In addition, the interaction among microservices of our design is only a simple chain, which cannot convey more complicated interaction schemes.

## 9. FUTURE WORK

As mentioned in the conclusions, an evaluation scenario with more complicated interaction patterns instead of chained microservices can produce more insights into the application performance. Other possibilities are to evaluate on the effects of data persistence strategies on microservices performance and scaling of the system, with multiple instances of each microservice where also load balancing could be tested.

## 10. ACKNOWLEDGEMENT

This work has received funding by the Federal Ministry for Economic Affairs and Energy (BMWi), Germany under grant no. 01MA17008, project Industrial Communication for Factories (IC4F).

## 11. REFERENCES

- [1] Carlos M Aderaldo, Nabor C Mendonça, Claus Pahl, and Pooyan Jamshidi. 2017. Benchmark requirements for microservices architecture research. In *2017 IEEE/ACM 1st International Workshop on Establishing the Community-Wide Infrastructure for Architecture-Based Software Engineering (ECASE)*. IEEE, 8–13.
- [2] Marcelo Amaral, Jorda Polo, David Carrera, Iqbal Mohamed, Merve Unuvar, and Malgorzata Steinder. 2015. Performance evaluation of microservices architectures using containers. In *2015 IEEE 14th International Symposium on Network Computing and Applications*. IEEE, 27–34.
- [3] Kapil Bakshi. 2017. Microservices-based software architecture and approaches. In *2017 IEEE Aerospace Conference*. IEEE, 1–8.
- [4] Michael Behrisch, Laura Bieker, Jakob Erdmann, and Daniel Krajzewicz. 2011. SUMO—simulation of urban mobility: an overview. In *Proceedings of SIMUL 2011, The Third International Conference on Advances in System Simulation*. ThinkMind.
- [5] Hai Dinh-Tuan, Felix Beierle, and Sandro Rodriguez Garzon. 2019. MAIA: A Microservices-based Architecture for Industrial Data Analytics. In *2019 IEEE International Conference on Industrial Cyber Physical Systems (ICPS)*. IEEE, 23–30.
- [6] Mikhail G Dozmorov. 2018. GitHub statistics as a measure of the impact of opensource bioinformatics software.



- Frontiers in bioengineering and biotechnology* 6 (2018), 198.
- [7] Mohamed Fayad and Douglas C Schmidt. 1997. Object-oriented application frameworks. *Commun. ACM* 40, 10 (1997), 32–38.
- [8] Mohamed E Fayad, David S Hamu, and Davide Brugali. 2000. Enterprise frameworks characteristics, criteria, and challenges. *Commun. ACM* 43, 10 (2000), 39–46.
- [9] Martin Garriga. 2017. Towards a taxonomy of microservices architectures. In *International Conference on Software Engineering and Formal Methods*. Springer, 203–218.
- [10] Sandro Rodriguez Garzon and Axel Küpper. 2019. Pay-per-Pollution: Towards an Air Pollution-aware Toll System for Smart Cities. In *2019 IEEE International Conference on Smart Internet of Things (SmartIoT)*. IEEE, 361–366.
- [11] Stephan Huber and Christoph Rust. 2016. Calculate travel time and distance with OpenStreetMap data using the Open Source Routing Machine (OSRM). *The Stata Journal* 16, 2 (2016), 416–423.
- [12] Ralph E Johnson. 1997. Components, frameworks, patterns. In *Proceedings of the 1997 symposium on Software reusability*. 10–17.
- [13] Ralph E Johnson and Brian Foote. 1988. Designing reusable classes. *Journal of object-oriented programming* 1, 2 (1988), 22–35.
- [14] Jaewook Kim, Tae Joon Jun, Daeyoun Kang, Dohyeun Kim, and Daeyoung Kim. 2018. GPU Enabled Serverless Computing Framework. In *2018 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)*. IEEE, 533–540.
- [15] James Lewis and Martin Fowler. 2014. Microservices. *MartinFowler.com* (2014).
- [16] David Lion, Adrian Chiu, Hailong Sun, Xin Zhuang, Nikola Greveski, and Ding Yuan. 2016. Don’t get caught in the cold, warm-up your {JVM}: Understand and eliminate {JVM} warm-up overhead in data-parallel systems. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*. 383–400.
- [17] Michael Mattsson. 1999. Effort distribution in a six year industrial application framework project. In *Proceedings IEEE International Conference on Software Maintenance-1999 (ICSM’99). ‘Software Maintenance for Business Change’ (Cat. No. 99CB36360)*. IEEE, 326–333.
- [18] Nora McDonald and Sean Goggins. 2013. Performance and participation in open source software on github. In *CHI’13 Extended Abstracts on Human Factors in Computing Systems*. 139–144.
- [19] Maurizio Morisio, Daniele Romano, and Ioannis Stamelos. 2002. Quality, productivity, and learning in framework-based development: an exploratory case study. *IEEE Transactions on Software Engineering* 28, 9 (2002), 876–888.
- [20] World Health Organization. 2018. Ambient (outdoor) air pollution. [https://www.who.int/en/news-room/fact-sheets/detail/ambient-\(outdoor\)-airquality-and-health](https://www.who.int/en/news-room/fact-sheets/detail/ambient-(outdoor)-airquality-and-health)
- [21] Claus Pahl and Pooyan Jamshidi. 2016. Microservices: A Systematic Mapping Study.. In *CLOSER* (1). 137–146.
- [22] Daniel Richter, Marcus Konrad, Katharina Utecht, and Andreas Polze. 2017. Highly-available applications on unreliable infrastructure: Microservice architectures in practice. In *2017 IEEE International Conference on Software Quality, Reliability and Security Companion (QRS-C)*. IEEE, 130–137.
- [23] Savitha Srinivasan. 1999. Design patterns in object-oriented frameworks. *Computer* 32, 2 (1999), 24–32.
- [24] J Akshitha Sriraman and Thomas F Wenisch. 2018.  $\mu$  suite: a benchmark suite for microservices. In *2018 IEEE International Symposium on Workload Characterization (IISWC)*. IEEE, 1–12.
- [25] Takanori Ueda, Takuya Nakaïke, and Moriyoshi Ohara. 2016. Workload characterization for microservices. In *2016 IEEE international symposium on workload characterization (IISWC)*. IEEE, 1–10.
- [26] Ying Wang, Su Song, SHIYONG QIU, Lu Lu, Yilin Ma, Xiaoyi Li, and Ying Hu. 2017. Study on international practices for low emission zone and congestion charging. *Beijing, World Research Institute* (2017).