

Sliceable Monolith: Monolith First, Microservices Later

Fabrizio Montesi
University of Southern Denmark
fmontesi@imada.sdu.dk

Marco Peressotti
University of Southern Denmark
peressotti@imada.sdu.dk

Valentino Picotti
University of Southern Denmark
picotti@imada.sdu.dk

Abstract—We propose **Sliceable Monolith**, a new methodology for developing microservice architectures and perform their integration testing by leveraging most of the simplicity of a monolith: a single codebase and a local execution environment that simulates distribution. Then, a tool compiles a codebase for each microservice and a cloud deployment configuration. The key enabler of our approach is the technology-agnostic service definition language offered by Jolie.

Index Terms—Microservices, Monolith, Jolie, Containers, Cloud Computing.

I. INTRODUCTION

Microservices represent the prominent software paradigm for building distributed applications that strive for scalability, maintainability, and tight development and deployment cycles [1]. Microservices enforce strong boundaries and interact by message passing, leading to modular and independently executable software components. However, they require dealing with multiple codebases (one per microservice), making prototyping and testing more challenging compared to a monolith—a standard application that consists of a single executable.

The complexity introduced by microservices can easily outweigh their benefits and, especially when it comes to greenfield project development, experts have mixed opinions on whether to start with microservices or with a monolith [2–4]. Thus we ask: Can we recover some of the simplicity of monoliths in the development of microservices? A positive answer would contribute to making the greenfield development of microservice systems more approachable, which is important because migrating monoliths to microservices is difficult [3].

In this article, we propose a new development methodology whereby an entire microservice architecture has a single codebase. Thus, our approach reduces drastically the complexity of reaching a working prototype to iterate on. We depict our methodology in Figure 1 and outline it in the following.

The main artifact in the codebase is a “sliceable monolith”: the definition of a microservice system that looks like a monolith, but where all components are enforced to be services with clear boundaries and data models (e.g., the structures of Data Transfer Objects). We achieve these features by using the Jolie programming language [5]. Jolie enforces linguistically some best practices for microservice development, e.g., interaction among components happens necessarily through formally-defined service interfaces. Thanks to the built-in facilities of

the Jolie interpreter, the application can then be tested locally straight away, enabling fast refinement cycles of the prototype.

The structure of a sliceable monolith make it possible to automatically extract the implementation of each microservice into its own codebase. We implement this procedure with an automatic *slicer* tool (called Jolie Slicer), emphasising the fact that the sliceable monolith is cut alongside the sharp boundaries of the microservices. Our slicer tool also produces the necessary configuration for the containerisation and distributed deployment of the microservice system on the cloud. At this point, developers are free to choose between iterating on the sliceable monolith codebase, or to start developing some (even all) of the microservices independently. The technology-agnostic nature of Jolie interfaces makes it possible to mix different languages for the implementation of each microservice (Jolie currently supports its own behavioural language, Java, and JavaScript, with a plug-in architecture for adding more [5]).

II. A USE CASE FROM SMART CITIES

In this section, we present our development methodology with a running example based on (the relevant parts of) the implementation of a use case from smart cities. We model a scenario in which a microservice architecture manages a set of private parking areas whose owners, in agreement with the power grid operators, have decided to share their charging stations in exchange for monetary incentives. The end-users of the application are owners of electrical vehicles looking for available charging stations near a given location. The described scenario is borrowed from [6, 7], where microservices interact according to the Domain Event pattern. The application follows the CQRS pattern (Command Query Responsibility Segregation): two services, one for querying data (*QuerySide*) and the other for updating data (*CommandSide*) interact indirectly through an event store service (*EventStore*).

A. Development and Local Testing

Service *QuerySide* allows clients to obtain information about parking spaces equipped with charging stations, performing queries based on geolocation. Service *CommandSide* offers an API for updating information about parking spaces. Service *EventStore* supports the coordination of the other two services by offering an API for event-driven communication.

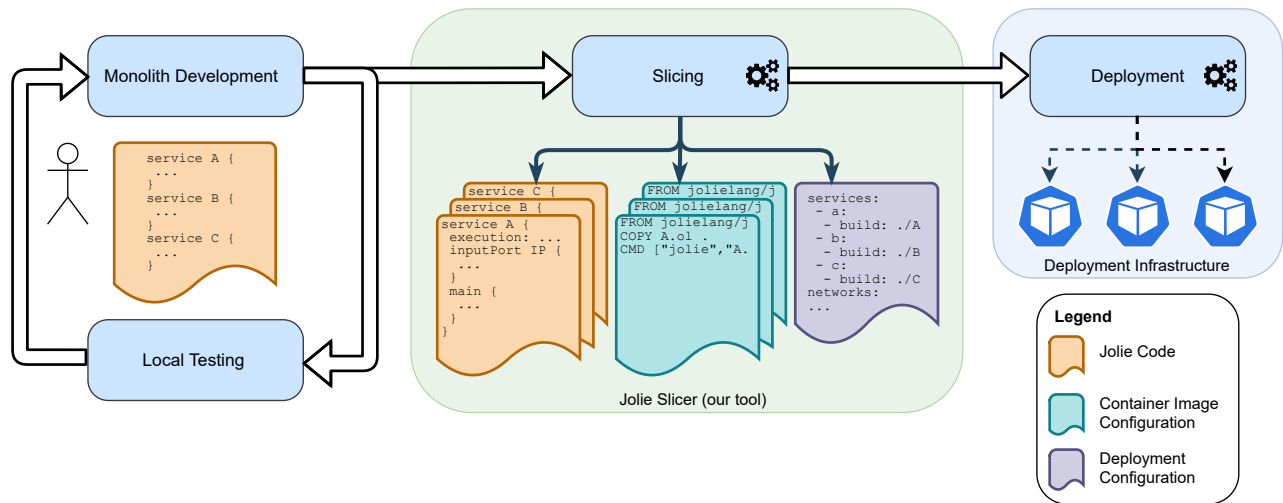


Figure 1. The development methodology of Sliceable Monolith.

Following the Sliceable Monolith approach, in this subsection we define the interfaces (data models and APIs) and implementations of these services in a single codebase.

We start by writing a Jolie program (smart-city.ol) that defines a service block for each one of our services:

```

service QuerySide( config ) { ... }
service CommandSide( config ) { ... }
service EventStore( config ) { ... }

```

Each service is parameterised on an externally-provided configuration, called `config`.

We show the most interesting parts of the definition of `CommandSide`, which is the most involved. First, we use Jolie data types to define the data model of the messages that the service exchanges. These include parking area identifiers and structures for providing information about each area: name, the time period in which it is available, the speed supported by the charging station in the area, and geolocation.

```

type PAID:long // Parking Area Identifier
type ParkingArea {
  id:PAID
  info:ParkingAreaInformation
}
type ParkingAreaInformation {
  name:string
  availability*:TimePeriod
  chargingSpeed:ChargingSpeed
  geolocation:Location
}

```

Using the data model, we build the API of `CommandSide` as a Jolie interface that comprises three RPCs for creating, updating, and deleting parking areas (RequestResponse is Jolie for RPC):

```

interface CommandSideInterface {
RequestResponse:
  createParkingArea( ParkingAreaInformation )( PAID ),
  updateParkingArea( ParkingArea )( string ),
  deleteParkingArea( PAID )( string )
}

```

In the definition of service `CommandSide`, we offer the API that we have just defined to clients (through a Jolie `inputPort`) and we declare a dependency towards the `EventStore` service (through a Jolie `outputPort`). The locations at which these communication ports should be deployed at are parameters that we get from the externally-provided configuration.

```

/* ... data types and API definitions ... */
service CommandSide( config:Configuration ) {
  execution: concurrent
  inputPort InputCommands {
    location: config.CommandSide.location
    protocol: http { format = "json" }
    interfaces: CommandSideInterface
  }
  outputPort EventStore {
    location: config.EventStore.location
    protocol: http { format = "json" }
    interfaces: EventStoreInterface
  }
  main { /* business logic implementation */ }
}

```

Note that data types and APIs appear outside of the `service` block that implements `CommandSide`. This allows us to share these same types across the implementations of all services, which aids in keeping the data models of interacting services consistent. For example, we can conveniently reuse the type `ParkingArea` in the data types of events that can be exchanged

through the EventStore service. This convenience comes at zero cost: our slicer tool (next subsection) uses a dependency analysis to produce optimised code for each microservice.

Externally-provided configurations for Jolie services can be given as JSON files. By using a JSON file that provides locations at the local host, we can test the entire architecture locally, both with unit and integration tests. While the latter is typically problematic in general, in our case it is simply a matter of writing a few lines of Jolie code that can be run locally. For example, the following test checks that deleting a parking area triggers the right event notification from EventStore.

```
subscribe@EventStore( {
  location = testLocation
  topics[0] = "PA_DELETED"
} )( res )
deleteParkingArea@CommandSide( 123L )()
notify( event )
if( event.type != "PA_DELETED" || event.id != 123L )
  throw( AssertionFailed )
```

Jolie

B. Slicing and Deployment

To switch to a distributed architecture, we write another configuration file (`deploy.json`), where locations are abstract DNS names (resolved by Docker, in our concrete deployment). We then run our Jolie Slicer tool with the following command.

```
jolie-slicer --config deploy.json smart-city.ol
```

This produces an output directory that contains a Docker Compose file for cloud deployment and a subfolder for each service. Each subfolder contains the code of its respective service alongside a `Dockerfile` that instructs Docker on how the service should be containerised, like the following.

```
FROM jolielang/jolie
COPY CommandSide.ol .
COPY ../deploy.json .
CMD ["jolie", "-p", "deploy.json", "CommandSide.ol"]
```

Dockerfile

The generated `docker-compose.yml` file is compatible with Docker Swarm. It can be used as is, but the programmer can also refine it, e.g., with the desired load balancing configuration:

```
services:
  commandside:
    build: ./commandside
    deploy:
      replicas: 1
    ...
```

Docker Compose

From a Docker Swarm manager node, the following command will deploy the entire architecture as a composition of independent microservices, as expected.

```
docker stack deploy -c docker-compose.yml smartcity
```

III. CONCLUSION

When developing a new service system, developers have to choose whether to start with a monolith or jump directly to microservices. The first choice leads to quick prototyping, but poses the risk of discarding the monolith entirely when migrating to microservices [2]. The second choice gives immediately a more flexible architecture, but at the cost of severely slowing down early development [3, 8].

Sliceable Monolith is a new middle ground between these approaches, which retains some of the simplicity of developing a monolith and automates the migration to a full-fledged distributed system of microservices.

We have omitted the business logic implementations of our services in the example; they are not surprising, and apply typical best practices (especially statelessness, where good scalability is desired). Our methodology abstracts from the technologies used for such implementations. Our tool already supports different technologies through Jolie, which has a plug-in architecture for implementing the business logics of services (the “main” block) in other languages, like Java and JavaScript [5]. Our tool currently supports Docker Swarm for cloud deployment, but in principle our methodology applies also to other technologies, like Kubernetes, Pulumi, or Terraform.

Our approach gives a first positive answer to the question in the introduction. It joins multiparty languages, an emerging class of distributed programming paradigms [9]. Interesting future work includes the systematic study of its impact.

Acknowledgments: We thank Florian Rademacher for useful discussions about the use case. Work partially supported by Independent Research Fund Denmark, grant no. 0135-00219.

REFERENCES

- [1] N. Dragoni, S. Giallorenzo, A. L. Lafuente, M. Mazzara, F. Montesi, R. Mustafin, and L. Safina, “Microservices: Yesterday, today, and tomorrow,” in *Present and Ulterior Software Engineering*. Springer, 2017, pp. 195–216.
- [2] M. Fowler. (2015) Monolith first. [Online]. <https://martinfowler.com/bliki/MonolithFirst.html>
- [3] S. Tilkov. (2015) Don’t start with a monolith. [Online]. <https://martinfowler.com/articles/dont-start-monolith.html>
- [4] S. Newman. (2015) Microservices for greenfield? [Online]. <https://samnewman.io/blog/2015/04/07/microservices-for-greenfield/>
- [5] F. Montesi, C. Guidi, and G. Zavattaro, *Service-Oriented Programming with Jolie*. Springer, 2014, pp. 81–107.
- [6] F. Rademacher, “A language ecosystem for modeling microservice architecture,” Ph.D. dissertation, University of Kassel, 2021, forthcoming.
- [7] (2020) Puls: Parken und laden in der stadt. [Online]. <https://parken-und-laden.de/>
- [8] S. Newman, *Building Microservices: Designing Fine-Grained Systems*, 1st ed. O’Reilly, 2015.
- [9] S. Giallorenzo, F. Montesi, M. Peressotti, D. Richter, G. Salvaneschi, and P. Weisenburger, “Multiparty languages: The choreographic and multitier cases,” in *ECOOP 2021*, ser. LIPIcs, vol. 194, 2021, pp. 22:1–22:27.