







A Toolchain for Checking Domain- and Model-driven Properties of Jolie Microservices

Saverio Giallorenzo¹, Fabrizio Montesi², Marco Peressotti²,
Florian Rademacher³, Sabine Sachweh⁴, and Philip Wizenty⁴

¹ Università di Bologna, Italy and INRIA, France saverio.giallorenzo@gmail.com

² University of Southern Denmark, Denmark [{fmontesi,peressotti}@imada.sdu.dk">{fmontesi,peressotti}@imada.sdu.dk](mailto)

³ codecentric AG, Germany rademacher@se-rwth.de

⁴ University of Applied Science and Arts Dortmund, Germany
[{sabine.sachweh,philip.wizenty}@fh-dortmund.de">{sabine.sachweh,philip.wizenty}@fh-dortmund.de](mailto)

Abstract We present Jolie Checker Toolchain (JCT), a plugin-based toolchain for model-code consistency and compliance analysis aimed at helping developers ensure that evolving implementations of Jolie microservices remain aligned with their intended architectural design—specified with domain- and model-driven engineering approaches like LEMMA and MDSL. One of JCT’s strengths lies in providing a uniform surface for plugin development, based on the analysis of abstract syntax trees, that leverages existing code generation tools, framing checks as correspondence relations between the expected and actual programs. We present two JCT plugins to respectively check the consistency of domain-driven design annotations on Jolie APIs and verify Jolie code conformity to a given LEMMA data model. We illustrate both plugins via a use case drawn from the Lakeside Mutual architecture.

1 Introduction

Background: microservices and their model-driven development Microservices are small, independent, and reusable distributed components composed via message passing [21, 9]. Thanks to their granularity, they can make systems more flexible and scalable [10]. To be effective, the design of microservice architectures needs to be carefully carried out, with challenges including designing appropriate APIs and defining (micro)services around business capabilities [21, 28].

To facilitate the design and implementation of microservice architectures, researchers have investigated several tools based on high-level linguistic abstractions. Some of these tools apply Model-Driven Engineering (MDE) [6]: they provide modelling languages for the specification of microservice domain models, following the principles of Domain-Driven Design (DDD) [11]. Examples of MDE tools for microservices include LEMMA [26], MDSL [32], MicroBuilder [29], and JHipster⁵. Other tools are general programming languages oriented towards the

⁵ <https://www.jhipster.tech/jdl>

development of services, like Jolie [20] and Ballerina [22]. Jolie, in particular, offers both a technology-agnostic API language and a language for implementing the behaviour of microservices based on a process-algebraic approach [20, 19].

Recently, Giallorenzo et al. [16] observed that the high-level constructs to represent services and APIs used in some metamodels for MDE (like LEMMA and MDSL) share strong similarities with linguistic abstractions found in Jolie. This observation paved the way for the development of code generators that, given the model of a microservice architecture, automatically produce skeletons of Jolie implementations with model-compliant APIs [16, 15, 32, 13]. The vision behind these code generators is to support developers in modelling microservice architectures by applying MDE, and then seamlessly switching to a service-oriented programming language to define concrete implementations by refining the (correct by construction) skeletons generated from the source models [15].

Problem statement: keeping service implementations consistent with architecture modelling Software evolves. Once developers translate models into code skeletons (ideally following the above-mentioned correct-by-construction approach), they have to refine the code to reach executable implementations. This refinement, the regular maintenance of the software, and interventions for its performance improvement can induce a drift between the models and their implementation.

Essentially, all these activities share a model-code compliance challenge. After any edit of the artefacts involved (models and code), the developers need to perform the work of comparing the code with its models, to make sure that the former still complies with the properties of the latter. This task suffers of at least two major disadvantages. First, developers lose momentum during the development process as they need to integrate the task of checking the consistency/correspondence of the properties. Second, manual checks do not guarantee the reliability of the process. This phenomenon is further exacerbated by large-scale codebases, where breaking consistency can be as easy as removing a field from the type of a data-transfer object. This calls for automated tool support.

This work We present Jolie Checker Toolchain (JCT), a plugin-based toolchain that automates consistency and conformance-violations checks between models and Jolie service implementations. We base our development on an innovative approach to code analysers for compliance. Since there exist already model-to-code generators, we want to avoid redeveloping part of the logic that maps models to service code to obtain a code analyser for consistency/conformance. Contrarily, we leverage the existing generators, avoiding logic duplication and the possible misalignments due to the parallel development of these tools. The code analysis in JCT builds on top of existing code generators, where a model-to-code tool generates code skeletons from a model and the plugin uses the skeletons as specifications for the service code under analysis. Technically, JCT plugins base their comparison on the abstract syntax trees (ASTs) of the model-generated code and the one under scrutiny and implement their checks as correspondence relations between the two pieces of code. To enable the general reuse model-to-code generators, we offer a uniform surface for checkers, which come as plugins.

In this way, JCT can host a plethora of checkers, each verifying the properties of a given model w.r.t. the given Jolie code.

We represent the general structure of JCT in Figure 1. In JCT, developers can compose the different checker plugins in successive stages, where each stage re-

ceives the AST of the Jolie program under verification, performs its checks, and outputs its results to the user. The reason behind this structure is, besides accommodating the development of plugins by developers, the integration of JCT within existing Continuous Integration/Deployment (CI/CD) DevOps pipelines.

As shown in Figure 1, we present two illustrative plugins for JCT. The first, introduced in Section 2, is a checker for consistency of Domain-Driven Design patterns defined as annotations of Jolie code. The second, described in Section 3, leverages the state-of-the-art LEMMA2Jolie code generator (which generates Jolie service code from LEMMA domain models [13]) to implement the LEMMA4Jolie plugin for checking the compliance between LEMMA models and their Jolie implementation. Since we build both plugins from existing work, in Section 2 and Section 3 we briefly summarise the related background work.

In Section 4, we discuss how developers can use JCT and our plugins to check DDD and model-code properties of Jolie programs. This discussion also provides an initial validation of our approach by presenting implementations of a scenario based on a reference MSA from the literature: Lakeside Mutual [32]. Finally, we compare our work with the related one in Section 5, and we draw closing remarks and future steps in Section 6. The tool and use case presented in this work are available at <https://anonymous.4open.science/r/vcxbdxn/> together with additional examples.

2 DDD Concepts and Patterns in Jolie

2.1 Background: Jolie IDL

Jolie provides an Interface Description Language (IDL) for specifying the operations offered and consumed by each microservice and the types of the semi-structured data accepted and returned by those operations. Jolie IDL is designed to be technology-agnostic: its types model Data Transfer Objects (DTOs) built on native types generally available in most architectures [7].

Figure 2 contains the grammar for Jolie IDL taken from [20] and updated to the latest major release at the time of writing (Jolie 1.11). An **interface** is a

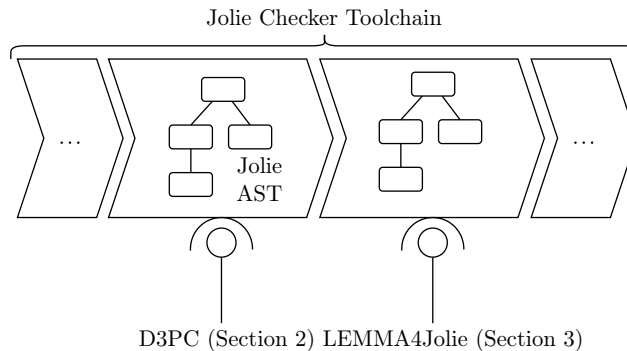


Figure 1. Schema of the Jolie Checker Toolchain (JCT).

```

I ::= MD interface id {RequestResponse MD id(TP1)(TP2)}
TP ::= id | B
TD ::= MD type id : T
T ::= B [{MD id C : T}] | undefined
C ::= [min, max] | * | ?
B ::= int(R) | string(R) | void | ...
R ::= range(min, max) | length(min, max) | enum(...) | ...
MD ::= ε | /** metadata */ | /// metadata

```

Figure 2. Jolie IDL, simplified syntax (types and interfaces).

collection of named operations (**RequestResponse**), where the sender delivers its message of type TP_1 and waits for the receiver to reply with a response of type TP_2 . Jolie also supports **oneWay** operations, where the sender delivers its message to the receiver, without waiting for the latter to process it (fire-and-forget); we omit this class of operations for conciseness. Operations have types describing the shape of the data structures they can exchange, which can either define custom, named types (id) or basic ones (B) like **integers**, **strings**, etc. Jolie **type** definitions (TD) have a tree-shaped structure to characterise semi-structured data typical of web technologies such as XML, JSON. At their root, we find a basic type (B)—which may include a refinement (R) to express constraints that further restrict the possible inhabitants of the type [12]. The possible branches (or fields) of a **type** are a set of nodes, where each node associates a name (id) with an array with a range length (C) and a type T .

We can decorate interfaces, types, and their members (operations and fields) with syntactic metadata (MD) using line ($/// [..]$) and multi-line ($/** [...] */$) documentation comments. These elements are not part of the API but are available to the Jolie interpreter and exposed to plugins by JCT.

2.2 D3PC: Realising DDD Patterns and Their Contracts

Recently, Giallorenzo et al. [16] established a connection between the metamod-els of LEMMA and Jolie, paving the way for the principled mapping of DDD concepts and patterns to Jolie developed in consecutive work [15, 14]. To elicit and document instances of DDD concepts in Jolie APIs, Giallorenzo et al. [15] introduced *DDD annotations* i.e., syntactic metadata comprised by expressions of the form $@Concept$ and $@Concept(params)$ for parametric concepts. In *op. cit.*, DDD annotations document the intent of using DDD patterns leaving the onus of correctly applying these patterns to programmers. Our Domain-driven Design Pattern Checker (D3PC) JCT plugin fills this gap. In the remainder, we summarise the patterns found in *op. cit.* and present the pattern-induced verifications performed by D3PC (implementation and a use case are in Section 4).

Aggregate and Part In DDD, aggregates prescribe object graphs requiring their parts to maintain a consistent state [11]. Thus, aggregates are always loaded and stored in a consistent state and within a single transaction. Aggregates and their parts are realised as Jolie types annotated with *@aggregate* and their branches annotated with *@part*. Figure 3 shows an example of an aggregate representing a parking space booking. A DDD aggregate must consist of at least one part and at least one part must be an entity or a value object, which we discuss below. D3PC verifies that types representing aggregates follow these requirements.

```

/// @ctx(BookingManagement)
/** @aggregate
    @entity */
type ParkingSpaceBooking {
    ///@identifier
    bookingID: long
    /// @part
    timeSlot: TimeSlot
    priceInEuro: double
}
/** @valueObject */
type TimeSlot { . . . }

```

Jolie

Figure 3. An example of DDD annotations [15, Sec. 2.4].

Entity and Identifier Instances of DDD entities are distinguishable by a domain-specific identity [11], e.g., a unique ID. In Jolie, DDD entities are realised as types annotated with *@entity* and their identifier as a branch annotated with *@identifier*. D3PC verifies that types representing entities follow these requirements. Continuing our example, we extend our *PSB* aggregate with an identifier (a booking ID) and annotate them accordingly.

Value Object and Domain Event In DDD, value objects consist of data and logic that do not depend on a domain-specific identity, differently from entities. Therefore, value objects serve as data transfer objects (DTOs) to exchange data between microservices and, in asynchronous scenarios, domain events. These concepts are realised in Jolie by annotating types with *@valueObject* and *@event* respectively. D3PC verifies that types representing value objects and events follow these requirements. The type *TimeSlot*, shown in Figure 3, is marked as a value object meaning that it can serve as a DTO and cross boundaries between microservices (see bounded contexts).

Bounded Context In DDD for microservice architectures, the concept of bounded context makes the boundaries for domain concepts explicit. Contexts are realised as modules which are constituted by Jolie files marked with the annotation *@ctx(ContextName)* parametric in the context name. To comply with this concept, Jolie codebases must avoid imports of types and interfaces across modules marked as distinct bounded contexts if they realise DDD concepts other than value object and domain event. Likewise, operations of interfaces meant to be exposed outside a context (even if they are not themselves marked with any DDD annotation) must not expose return types that realise DDD concepts other than value object and domain event. D3PC verifies all that the elements of a module representing a context are used accordingly.

<pre> context UserManagement { structure User <entity> { int id <identifier>, string firstName, string lastName, Username username } </pre> <p style="text-align: right;">LEMMA</p>	<pre> structure Username<valueObject> { immutable string username } enum Status { ENABLED, DISABLED, INACTIVE } </pre> <p style="text-align: right;">LEMMA</p>
---	--

Figure 4. Illustrative LEMMA domain model: the UserManagement bounded context.

3 LEMMA domain models and Jolie

The second plugin we present builds on work on LEMMA and Jolie [13] for developing MSA-based software systems. In the following, we introduce in Section 3.1 LEMMA’s domain data modelling language for the DDD-supported design of MSAs by Domain Experts and Service Developers [21]. In Section 3.2, we elaborate on using LEMMA’s domain data models to create Jolie source code artefacts by applying methods and techniques from Model-driven Engineering (MDE). Then, in Section 3.3, we introduce the methodology followed to implement the correspondence relation that underpins the comparison between the LEMMA domain data model and the Jolie service under verification.

3.1 Background: LEMMA Domain Data Modeling Language

The LEMMA ecosystem [25] for Model-Driven Engineering (MDE) offers a range of modeling languages that capture different concerns in Microservices Architecture (MSA) engineering from the perspective of architecture stakeholders. Developers can integrate these MSA models via an import mechanism—supporting referencing between elements of diverse models—to enhance the information content of captured viewpoints in a microservice architecture, also facilitating reuse.

Domain Data Modelling Language (DDML). The Domain Data Modelling Language (DDML) is a language that enables domain experts and microservice developers to construct models addressing concerns from the Domain Viewpoint. Domain experts and developers can use DDML to create domain models that capture relevant concepts from the application domain. These concepts can be enhanced with patterns from DDD [11]—a popular methodology in microservice development. DDML also implements LEMMA’s type system, making domain concepts usable for typing parameters of modelled microservice operations. In this way, typing relationships help to identify the portion of the application domain on which a microservice operates and for which it is responsible.

We report an example in Figure 4 of DDML model with integrated concepts from DDD of a UserManagement context relating to the bounded context [11].

The UserManagement bounded context contains two data structures that capture domain information. The User data structure clusters information data

```

[[context id {CT}]] = /** @ctx(id) */ [[CT]]
((structure id [{STRF}] {FLD OPS})) = [/** @STRF */] interface id_interface {((OPS))_id}
[[structure id [{STRF}] {FLD OPS}]] = type [[structure id [{STRF}] {FLD}]]
[[OPS]]_id ((structure id [{STRF}] {OPS}))_id
[[procedure id [{OPSF}] (FLD)]_id_s = type id_type: void {self?: id_s [[FLD]]} ...
[[structure id [{STRF}] {FLD}]] = [/** @STRF */] id: void {[[FLD]]}
[[S id [{FLDF}]] = [/** @FLDF */] id: [[S]]
...
[[unspecified]] = undefined

```

Figure 5. Extract of the Jolie encoding for LEMMA’s domain modeling concepts [16].

about the user of the software system, e.g., the primitive type attributes firstName and lastName. The User data structure contains the complex attribute UserName, which is specified as a type to make the domain concept explicit.

The domain model contains DDD concepts, e.g., **entity** assigned to User relates to the eponymous DDD concept, whose **identifier** attribute is id.

3.2 Background: From Domain Models to Jolie APIs

Our LEMMA4Jolie plugin uses in one of its steps LEMMA2Jolie [15, 13], a *Model-to-Text Transformation* that generates Jolie source code from LEMMA domain models. Figure 5 contains an excerpt of this transformation.

The encoding specifies, e.g., that a LEMMA context is transformed into the Jolie annotations of `@ctx(...)` or that the complex type of data structure becomes a correspondingly-structured Jolie **type**. Figure 6 shows the Jolie source code generated by the transformation of the LEMMA domain model from Figure 4.

<pre> /// @ctx(User) /// @entity type User { ///@identifier id: int firstName: string lastName: string username: Username } </pre> <p style="text-align: right;">Jolie</p>	<pre> /// @valueObject type Username { username: string } type Status { literal : string(enum(["ENABLED", "DISABLED", "INACTIVE"])) } </pre> <p style="text-align: right;">Jolie</p>
--	--

Figure 6. Example Jolie code generated from the model in Figure 4 by LEMMA2Jolie.

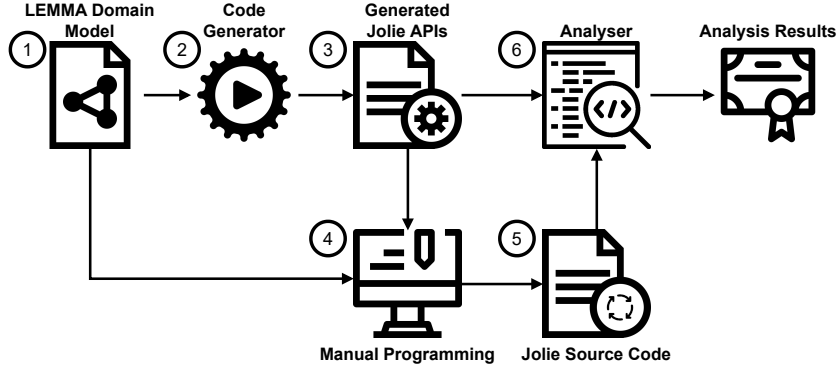


Figure 7. Schema of the LEMMA4Jolie compliance routine.

Briefly, in Figure 6, we find the Jolie types `User` and `Username` from the related LEMMA data structures, including the DDD concepts applied on them as Jolie comments, starting with a `@`-symbol to indicate DDD-specific annotations.

3.3 LEMMA4Jolie: Compliance of Jolie APIs with Domain Models

We draw in Figure 7 the steps that characterise our LEMMA4Jolie JCT plugin, which includes code generation, manual programming, and the overall development process of the software systems.

Step ① consists of modelling the software systems domain concepts using LEMMA’s domain data modelling language. The domain model, constructed by domain experts and service developers, is part of the software systems specification and, therefore, serves as the Architectural Design [30] of the system.

In step ②, we focus on the software systems’ initial development. This step uses a code generator (LEMMA2Jolie) to generate the skeleton implementation of the system as Jolie source code, referenced in ③. Note that the approach does not necessarily requires the developer to start from the generated code, which they can also write from scratch (as shown in ④). Hence, the code generated from ② is only mandatory for the LEMMA4Jolie plugin as the reference specification to compare against the provided Jolie implementation.

Step ④ indicates the possible manual modifications of the Jolie code which we are going to check. This step ranges on all the ways developers can manually intervene on the Jolie code, including from-scratch implementations of the model to modifications of the model-generated code—e.g., to realise architectural designs from other models or improve performance, fix bugs, and integrate new features. All these modifications result in step ⑤, which references the Jolie source code that we want to check for correspondence with the LEMMA model.

Step ⑥ implements the correspondence relation at the heart of the compliance checks between the LEMMA2Jolie-generated Jolie specification and the Jolie source code. Concisely, the relation checks that each Jolie type and interface

found in the LEMMA-generated source has a corresponding type and interface in the code under verification, including the related annotations for DDD patterns. Any structural deviation between the former and the latter results in a violation reported to the developer—the report indicates the line (if any) where the code presents a violation and its explanation (e.g., a missing expected annotation).

4 Implementation and Use Case

To illustrate the usage of JCT and the D3PC and LEMMA4Jolie plugins and provide a preliminary validation of their design, we apply them to an illustrative use case based on the textbook MSA Lakeside Mutual [32]. Specifically, we modelled in LEMMA the use case to elicit DDD Patterns, we generated the Jolie APIs via LEMMA2Jolie, introducing to the latter illustrative modifications to exemplify the usage of the plugins. JCT is a terminal tool designed for both interactive and automatic reporting within a CI/CD infrastructure. Within a CI/CD, JCT shall stop the process when it detects a violation, while the interaction with the users mirrors that of a compiler, where warnings and errors carry the reason behind the reported violation to help the users identify the offending lines (if any). We exemplify both compliance/conformance reporting and the fixing of illustrative violations with the Jolie code in Figure 8, stored within a `example.ol` file. Using D3PC, we obtain the following output.

```

///ctx(User)
/** @entity
    @aggregate */
type Customer {
  customerRef: int
  firstName: string
  lastName: string
  ///@identifier
  username: Username
}
...
Jolie

```

Figure 8. `example.ol`

DDD	Line	Description	Status
Username	16	Value Object has no entity annotation	CORRECT
Customer	7	Entity has an identifier	CORRECT
Customer	7	Aggregate is an entity	CORRECT
Customer	7	Aggregate misses a part	VIOLATION

From the report, the code complies with the property of *@entity*, i.e., Customer has an *@identifier* subnode (username). Similarly, *@aggregate* is correctly attributed to Customer—i.e., an *@entity* can also be an *@aggregate*—but we have a VIOLATION because we need to associate an aggregate with one or more *@parts*. Since Username (omitted for brevity) is a complex type aggregated by Customer, we mark username *@part* to fix the violation. Running JCT with D3PC again confirms we fixed the problem.

```

type Customer {
  [...]
  /** @identifier
    @part */
  username: Username
}
Jolie

```

DDD	Line	Description	Status
Username	17	Value Object has no entity annotation	CORRECT
Customer	7	Entity has an identifier	CORRECT
Customer	7	Aggregate is an entity	CORRECT
Customer	7	Aggregate misses a part	CORRECT

The LEMMA4Jolie plugin needs a source Jolie file generated by LEMMA2Jolie from the LEMMA data model intended for the considered Jolie file. Specifically, the LEMMA model for our example is the following on the left (a slight extension of the one from Figure 3), of which we report the Jolie code generated on the right (lightly reformatted for presentation).

<pre>context User { structure Customer <entity,aggregate> { int customerRef (identifier), string firstName, string lastName, Username username <part> } ... }</pre>	<pre>///<i>@ctx</i>(User) /** <i>@entity</i> <i>@aggregate</i> */ type Customer { ///<i>@identifier</i> customerRef: int firstName: string lastName: string ///<i>@part</i> username: Username } ...</pre>
LEMMA	Jolie

Using LEMMA4Jolie, we obtain the following output.

Name	Line	Description	Status	Type
Customer	7	Matching annotations	CORRECT	ANNOTATION
customerRef	9	Type present	CORRECT	TYPE
customerRef	9	Annotation variation, Identifier	VIOLATION	ANNOTATION
firstName	10	Type present	CORRECT	TYPE
lastName	12	Type present	CORRECT	TYPE
username	13	Type present	CORRECT	TYPE
username	13	Annotation variation, Identifier	VIOLATION	ANNOTATION
username	13	Matching annotation	CORRECT	ANNOTATION

By comparing the codes of `example.ol` (updated) and the related LEMMA model, the two programs differ by the assignment of the *@identifier*. Indeed, in `example.ol`, we assign the *@identifier* annotation to the `username` leaf. This assignment was fine for the D3PC checker, which just verified that the `Customer` *@entity* had an *@identifier* subcomponent. Contrarily, that association violates the LEMMA model, which specifies that `customerRef` is the identifier of the `Customer` entity. LEMMA4Jolie reports the above error at two locations. From the top of the output, we first find that `customerRef` is missing the *@identifier* annotation, which is incorrectly attributed to the `username`—reported as an unexpected annotation by the tool as the second violation.

To make it valid for the LEMMA model, we need to fix the annotations of the `customerRef` and `username` nodes as shown in the Jolie file reported above.

5 Related Work

Looking at related work from the perspective of model-driven engineering, we see our proposal close to architecture-conformance analysis [18, 23], which aims to check whether a software’s architecture is consistent with its intended organization of structural elements (such as packages, subsystems, and layers) to reduce software erosion. Works the closest to ours, on microservices, include measures to assess architecture conformance to microservice patterns [31]

Architecture conformance checking is an approach to address erosion in software development. Caraccio et al. introduced a unified approach for conformance checking [5]. The approach is based on the domain-specific language `Dicto` to specify architecture rules to support the software developer in formalising functional and non-functional aspects of the software, e.g., that the architecture design of the software systems conforms to its implementation. `Probo`, a tool coordination framework, uses the specified rules to check the conformance of the software system. Our work differs from the approach described in [5] in various factors. The main difference is that our approach does not require additional artefacts in the software development process, such as the architecture rules defined via `Dicto`. `JCT` plugins (e.g., `D3PC` and `LEMMA4Jolie`) work with artefacts that already exist in the software development process. Both approaches enable the integration of additional tools to analyse the architecture conformance. However, the approach described by Caracciolo et al. [5] addresses architecture conformance generally, whereas our approach is tailored for microservices.

Buckley et al. presented `JITTAC` [4], a tool that leverages real-time reflection modelling to inform the software developers about the architectural consequences of their programming actions and, therefore, try to raise awareness towards promoting consistency between the design and implementation of the software system. `JITTAC` relies on an architectural model that consists of the different components of the software systems and their relationship with each other. The components of the architecture model are mapped to the various source code artefacts to monitor their evolution. The difference between `JCT` and `JITTAC` lies in the latter considering the extensibility, scope, and integration into the development process. Both tools aim to improve the architecture compliance of the software systems. However, `JITTAC` is directly integrated into the Eclipse IDE and, therefore, forces the IDE usage with a fixed number of analyses. In contrast, `JCT` does not have those constraints for developing the software system due to the loose integration via terminal or into the CI/CD pipeline. Additionally, `JCT` focuses specifically on the needs of MSA developers. `JITTAC` supports only a central model specifying the architecture, which is a broader approach and lacks detailed MSA-specific concepts.

Another tool, `InMap` [27], focuses on automated interactive Code-to-Architecture Mappings recommendations. The tool uses reflection modelling to recom-

mend a mapping between source code artefacts to their corresponding model specifications. The mappings are then used to ensure the conformance between the software systems architecture and implementation. Differently from JCT, InMap does not analyse the compliance between design and implementation but rather specifies a possible mapping between the two.

6 Discussion and Conclusion

We presented JCT, a plugin-based toolchain for checking the respect of MDE and DDD properties of Jolie programs. The main idea behind JCT is leveraging the AST of Jolie programs to abstract away syntactic details and concentrate on the shape of programs. We conjecture that this abstraction step is particularly suitable for efficiently integrating different model-driven engineering code-generation technologies within a single toolchain. Indeed, besides providing facilities to obtain ASTs designed for the analysis of Jolie programs, JCT supports a plugin development methodology where constraints induced by models manifest as correspondence checks between Jolie code generated from models and the code developed by programmers. We present two illustrative JCT plugins namely, D3PC and LEMMA4Jolie, which both work on Jolie APIs. D3PC performs a consistency check among DDD annotations of Jolie code while LEMMA4Jolie verifies the conformance of a Jolie program w.r.t. a given LEMMA data model.

We conducted a preliminary validation of our toolchain with a use case consisting of the textbook MSA Lakeside Mutual [32] drawing positive results that support the design of JCT. However, the preliminary nature of this use case limits our conclusions and calls for more systematic investigations. To address the scope and limitations of our use case, we are planning to pursue two main directions: 1) retrace the steps in Section 4 with a larger sample of use cases and reference MSAs and 2) conduct qualitative analyses with Jolie programmers not involved in the development of JCT. We leave these as future work since the size of such studies evades the scope of this work.

We cast our work within the larger effort to provide Jolie programmers with tools to help them check properties and test their applications [13, 17] in CI/CD pipelines. Immediate future steps to evolve JCT regard the implementation of new plugins for tools that already support from-model Jolie code generation, like MDSL[32], but also integrate checkers that analyse code looking for “smells” of microservices, e.g., on security [24, 30].

While, in this paper, we provide a concrete implementation of the model-code consistency toolchain methodology targetting the Jolie language, we argue that the approach we propose is generally applicable to any language whose ecosystem of tools includes model-to-code generators [1]. Indeed, we conjecture that one can rely on the approach we introduced with JCT not only on paradigmatic programming languages for microservices like Jolie but also with libraries used to inject service-oriented concerns into programming languages based on other programming paradigms, e.g., with Java and the Spring framework⁶. Hence,

⁶ <https://spring.io/microservices>.

since LEMMA can generate Java/Spring microservice skeletons, one could build a tool similar to JCT for Java/Spring microservices. A concrete step in this direction is investigating the usage of the Java Checker Framework [8] for implementing a LEMMA4Jolie-like plugin for Java/Spring-based microservices that leverages a LEMMA-to-Java/Spring code generator to compare Java/Spring programs against their expected Java-based specifications. With more support from LEMMA generators, one can envision a generic plugin for performing the same ASTs comparison performed by LEMMA4Jolie but for arbitrary languages: instead of relying on deep integration with a generator for a specific language (as in the case of LEMMA4Jolie), such JCT plugin would work in tandem with LEMMA and use generic markers left by its generators to identify the ‘generated gap’. Widening our scope, we conjecture that our approach is not limited to one language per architecture and one could use a multi-target model framework, like LEMMA, to check in the same pipeline a mixed codebase of microservices, e.g., written in Jolie and Java/Spring—this extension would require adding annotations to indicate which microservices are written in which language. One could further broaden the coverage of the models and target languages. For example, one could integrate the tool by Bucchiarone et al. [3] to target also container configuration files (viz. Docker files) or the one by Belguidoum et al. [2] to cover IoT deployments.

Acknowledgements Partially supported by the research project FREEDA (CUP: I53D23003550006) funded by the framework PRIN 2022 (MUR, Italy), the French ANR project SmartCloud ANR-23-CE25-0012, and Villum Fonden (grant no. 29518). Co-funded by the European Union (ERC, CHORDS, 101124225). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

References

1. Balasubramanian, K., Gokhale, A.S., Karsai, G., Sztipanovits, J., and Neema, S.: Developing Applications Using Model-Driven Design Environments. *Computer* 39(2), 33–40 (2006). DOI: [10.1109/MC.2006.54](https://doi.org/10.1109/MC.2006.54)
2. Belguidoum, M., Gourari, A., and Sehili, I.: MDMSD4IoT a Model Driven Microservice Development for IoT Systems. In: Fournier-Viger, P., Yousef, A.H., and Bellatreche, L. (eds.) *Model and Data Engineering: 11th International Conference, MEDI 2022, Cairo, Egypt, November 21-24, 2022, Proceedings*. LNCS, vol. 13761, pp. 176–189. Springer, Heidelberg (2022). DOI: [10.1007/978-3-031-21595-7_13](https://doi.org/10.1007/978-3-031-21595-7_13)
3. Bucchiarone, A., Soysal, K., and Guidi, C.: A Model-Driven Approach Towards Automatic Migration to Microservices. In: Bruel, J., Mazzara, M., and Meyer, B. (eds.) *Software Engineering Aspects of Continuous Development and New Paradigms of Software Production and Deployment - Second International Workshop, DEVOPS 2019, Château de Villebrumier, France, May 6-8, 2019, Revised Selected Papers*. LNCS, vol. 12055, pp. 15–36. Springer, Heidelberg (2019). DOI: [10.1007/978-3-030-39306-9_2](https://doi.org/10.1007/978-3-030-39306-9_2)

4. Buckley, J., Mooney, S., Rosik, J., and Ali, N.: JITTAC: a just-in-time tool for architectural consistency. In: 2013 35th International Conference on Software Engineering (ICSE), pp. 1291–1294 (2013)
5. Caracciolo, A., Lungu, M.F., and Nierstrasz, O.: A unified approach to architecture conformance checking. In: 2015 12th Working IEEE/IFIP Conference on Software Architecture, pp. 41–50 (2015)
6. Combemale, B., France, R.B., Jézéquel, J.-M., Rumpe, B., Steel, J., and Vojtisek, D.: *Engineering Modeling Languages: Turning Domain Knowledge into Tools*. CRC Press (2017)
7. Daigneau, R.: *Service Design Patterns*. Addison-Wesley (2012)
8. Dietl, W., Dietzel, S., Ernst, M.D., Muslu, K., and Schiller, T.W.: Building and using pluggable type-checkers. In: Taylor, R.N., Gall, H.C., and Medvidovic, N. (eds.) *Proceedings of the 33rd International Conference on Software Engineering, ICSE 2011, Waikiki, Honolulu, HI, USA, May 21-28, 2011*, pp. 681–690. ACM (2011). DOI: 10.1145/1985793.1985889
9. Dragoni, N., Giallorenzo, S., Lafuente, A.L., Mazzara, M., Montesi, F., Mustafin, R., and Safina, L.: *Microservices: Yesterday, Today, and Tomorrow*. In: *Present and Ulterior Software Engineering*. Ed. by M. Mazzara and B. Meyer, pp. 195–216. Springer (2017)
10. Dragoni, N., Lanese, I., Larsen, S.T., Mazzara, M., Mustafin, R., and Safina, L.: *Microservices: How To Make Your Application Scale*. In: Petrenko, A.K., and Voronkov, A. (eds.) *Perspectives of System Informatics*, pp. 95–104. Springer, Cham (2018)
11. Evans, E.: *Domain-Driven Design*. Addison-Wesley (2004)
12. Freeman, T., and Pfenning, F.: Refinement types for ML. In: *Proc. of the 1991 Conf. on Programming Language Design and Implementation*, pp. 268–277 (1991)
13. Giallorenzo, S., Montesi, F., Peressotti, M., and Rademacher, F.: LEMMA2Jolie: A tool to generate microservice APIs from domain models. *Sci. Comput. Program.* 228, 102956 (2023). DOI: 10.1016/J.SCICO.2023.102956
14. Giallorenzo, S., Montesi, F., Peressotti, M., and Rademacher, F.: Model-Driven Code Generation for Microservices: Service Models. In: Dorai, G., Gabbriellini, M., Manzonetto, G., Osmani, A., Prandini, M., Zavattaro, G., and Zimmermann, O. (eds.) *Joint Post-proceedings of the Third and Fourth International Conference on Microservices (Microservices 2020/2022)*. Open Access Series in Informatics (OASICS), 6:1–6:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2023). DOI: 10.4230/OASICS.Microservices.2020-2022.6
15. Giallorenzo, S., Montesi, F., Peressotti, M., and Rademacher, F.: Model-Driven Generation of Microservice Interfaces: From LEMMA Domain Models to Jolie APIs. In: Beek, M.H. ter, and Sirjani, M. (eds.) *Coordination Models and Languages - 24th IFIP WG 6.1 International Conference, COORDINATION 2022, Held as Part of the 17th International Federated Conference on Distributed Computing Techniques, DisCoTec 2022, Lucca, Italy, June 13-17, 2022, Proceedings*. LNCS, vol. 13271, pp. 223–240. Springer, Heidelberg (2022). DOI: 10.1007/978-3-031-08143-9_13
16. Giallorenzo, S., Montesi, F., Peressotti, M., Rademacher, F., and Sachweh, S.: Jolie and LEMMA: Model-Driven Engineering and Programming Languages Meet on Microservices. In: *Coordination Models and Languages*, pp. 276–284. Springer (2021)
17. Giallorenzo, S., Montesi, F., Peressotti, M., Rademacher, F., and Unwerawattana, N.: JoT: A Jolie Framework for Testing Microservices. In: Jongmans, S., and Lopes,

- A. (eds.) Coordination Models and Languages - 25th IFIP WG 6.1 International Conference, COORDINATION 2023, Held as Part of the 18th International Federated Conference on Distributed Computing Techniques, DisCoTec 2023, Lisbon, Portugal, June 19-23, 2023, Proceedings. LNCS, vol. 13908, pp. 172–191. Springer, Heidelberg (2023). DOI: [10.1007/978-3-031-35361-1_10](https://doi.org/10.1007/978-3-031-35361-1_10)
18. Knodel, J., and Popescu, D.: A comparison of static architecture compliance checking approaches. In: 2007 Working IEEE/IFIP conference on software architecture (WICSA'07), pp. 12–12 (2007)
 19. Montesi, F., and Carbone, M.: Programming Services with Correlation Sets. In: Kappel, G., Maamar, Z., and Nezhad, H.R.M. (eds.) Service-Oriented Computing - 9th International Conference, ICSOC 2011, Paphos, Cyprus, December 5-8, 2011 Proceedings. LNCS, vol. 7084, pp. 125–141. Springer, Heidelberg (2011). DOI: [10.1007/978-3-642-25535-9_9](https://doi.org/10.1007/978-3-642-25535-9_9)
 20. Montesi, F., Guidi, C., and Zavattaro, G.: Service-Oriented Programming with Jolie. In: Web Services Foundations. Ed. by A. Bouguettaya, Q.Z. Sheng and F. Daniel, pp. 81–107. Springer (2014). DOI: [10.1007/978-1-4614-7518-7_4](https://doi.org/10.1007/978-1-4614-7518-7_4)
 21. Newman, S.: Building Microservices: Designing Fine-Grained Systems. O'Reilly (2015)
 22. Oram, A.: Ballerina: A Language for Network-Distributed Applications. O'Reilly (2019)
 23. Passos, L., Terra, R., Valente, M.T., Diniz, R., and Mendonça, N.: Static architecture-conformance checking: An illustrative overview. *IEEE software* 27(5), 82–89 (2009)
 24. Ponce, F., Soldani, J., Astudillo, H., and Brogi, A.: Smells and refactorings for microservices security: A multivocal literature review. *J. Syst. Softw.* 192, 111393 (2022). DOI: [10.1016/J.JSS.2022.111393](https://doi.org/10.1016/J.JSS.2022.111393)
 25. Rademacher, F.: A Language Ecosystem for Modeling Microservice Architecture. PhD thesis, Universität Kassel (2022).
 26. Rademacher, F., Sorgalla, J., Wizenty, P., Sachweh, S., and Zündorf, A.: Graphical and Textual Model-Driven Microservice Development. In: *Microservices: Science and Engineering*, pp. 147–179. Springer (2020)
 27. Sinkala, Z.T., and Herold, S.: InMap: automated interactive code-to-architecture mapping recommendations. In: 2021 IEEE 18th International Conference on Software Architecture (ICSA), pp. 173–183 (2021)
 28. Soldani, J., Tamburri, D.A., and Heuvel, W.-J.V.D.: The pains and gains of microservices: A Systematic grey literature review. *Journal of Systems and Software* 146, 215–232 (2018)
 29. Terzić, B., Dimitrieski, V., Kordić, S., Milosavljević, G., and Luković, I.: Development and evaluation of MicroBuilder: a Model-Driven tool for the specification of REST Microservice Software Architectures. *Enterprise Information Systems* 12(8-9), 1034–1057 (2018)
 30. Wizenty, P., Ponce, F., Rademacher, F., Soldani, J., Astudillo, H., Brogi, A., and Sachweh, S.: Towards Resolving Security Smells in Microservices, Model-Driven. In: Fill, H., Mayo, F.J.D., Sinderen, M. van, and Maciaszek, L.A. (eds.) *Proceedings of the 18th International Conference on Software Technologies, ICSoft 2023, Rome, Italy, July 10-12, 2023*, pp. 15–26. SCITEPRESS (2023). DOI: [10.5220/0012049800003538](https://doi.org/10.5220/0012049800003538)
 31. Zdun, U., Navarro, E., and Leymann, F.: Ensuring and assessing architecture conformance to microservice decomposition patterns. In: *Service-Oriented Computing: 15th International Conference, ICSOC 2017, Malaga, Spain, November 13–16, 2017*, Proceedings, pp. 411–429 (2017)

32. Zimmermann, O., Stocker, M., Lübke, D., Zdun, U., and Pautasso, C.: Patterns for API Design: Simplifying Integration with Loosely Coupled Message Exchanges. Addison-Wesley (2023)