Model-Driven Generation of Microservice Interfaces: From LEMMA Domain Models to Jolie APIs

Saverio Giallorenzo, Fabrizio Montesi, Marco Peressotti, Florian Rademacher

To cite this version:

Saverio Giallorenzo, Fabrizio Montesi, Marco Peressotti, Florian Rademacher. Model-Driven Generation of Microservice Interfaces: From LEMMA Domain Models to Jolie APIs. 24th IFIP WG 6.1 International Conference, 17th International Federated Conference on Distributed Computing Techniques (DisCoTec 2022), Jun 2022, Lucca, Italy. pp.223-240, 10.1007/978-3-031-08143-9_13 . hal-03915132

HAL Id: hal-03915132

https://inria.hal.science/hal-03915132

Submitted on 29 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Model-Driven Generation of Microservice Interfaces: From LEMMA Domain Models to Jolie APIs

Saverio Giallorenzo[1], Fabrizio Montesi[2], Marco Peressotti[2], and Florian Rademacher[3]

[1] Università di Bologna, Italy and INRIA, France saverio.giallorenzo@gmail.com
[2] University of Southern Denmark {fmontesi,peressotti}@imada.sdu.dk
[3] University of Applied Sciences and Arts Dortmund
florian.rademacher@fh-dortmund.de

Abstract We formally define and implement a translation from domain models in the LEMMA modelling framework to microservice APIs in the Jolie programming language. Our tool enables a software development process whereby microservice architectures can first be designed with the leading method of Domain-Driven Design, and then corresponding data types and service interfaces (APIs) in Jolie are automatically generated. Developers can extend and use these APIs as guides in order to produce compliant implementations. Our tool thus contributes to enhancing productivity and improving the design adherence of microservices.

1 Introduction

Microservice Architecture (MSA) is one of the current leading patterns in distributed software architectures [17]. While widely adopted, MSA comes with specific challenges regarding architecture design, development, and operation [5,24]. To cope with this complexity, researchers in software engineering and programming languages started proposing linguistic approaches to MSA: language frameworks that ease the design and development of MSAs with high-level constructs that make microservice concerns in the two different stages syntactically manifest.

Regarding development, Ballerina and Jolie are examples of programming languages [18,16] with new linguistic abstractions for effectively programming the configuration and coordination of microservices. Regarding design, Model-Driven Engineering (MDE) [3] has gained relevance as a method for the specification of service architectures [1], crystallised in MDE-for-MSA modelling languages such as MicroBuilder, MDSL, LEMMA, and JHipster [27,12,21,11]. Jolie's abstractions have been found to offer a productivity boost in industry [10]. LEMMA provides linguistic support for the application of concepts from Domain-Driven Design [6,21], and has been validated in real-world use cases [25,22].

Recently, it has been observed that the metamodels of LEMMA's modelling languages and the Jolie programming language have enough contact points to

consider their integration [9]. In the long term, such an integration could bring (quoting from [9])

> "*an ecosystem that coherently combines MDE and programming abstractions to offer a tower of abstractions [15] that supports a step-by-step refinement process from the abstract specification of a microservice architecture to its implementation*".

The aim is to provide a toolchain that enables people to apply MDE to the design of microservices in LEMMA, and then seamlessly switch to a programming language with dedicated support for microservices like Jolie in order to develop an implementation of the design. To this end, three important parts of the metamodels of LEMMA and Jolie need to be covered and integrated [9]:

1. *Application Programming Interfaces* (API), describing what functionalities (and their data types) a microservice offers to its clients;
2. *Access Points*, capturing where and how clients can interact with the API;
3. *Behaviours*, defining the internal business logic of a microservice.

Since the API is the layer the other two build upon, in this paper we focus on concretising the relationship between LEMMA and Jolie API layers. To this end, we contribute a formal encoding between LEMMA's Domain Data Modelling Language (DDML) and Jolie types and interfaces. This encoding enables systematic translation of LEMMA domain models, which, following DDD principles, capture domain-specific types including operation signatures, to Jolie APIs. As a second contribution, we present LEMMA2Jolie—a code generator that allows automatic translation of LEMMA domain models to Jolie APIs based on the introduced encoding. Specifically, LEMMA2Jolie not only shows the encoding's feasibility and practicability, but also constitutes a crucial contribution towards improving the adoption of DDD in microservice design, which in practice is often perceived complex given the lack of formal guidelines on how to map DDD domain models to microservice code [2]. We have evaluated LEMMA2Jolie in the context of a nontrivial microservice architecture that had previously been used to validate LEMMA [22], which covers all the aspects of the formal encoding. The generated Jolie code is as expected, in the sense that it is faithful to the formal encoding and the model defined in LEMMA. We use snippets of this code to exemplify our method throughout the paper.

The remainder of the paper is organised as follows. Section 2 introduces and exemplifies the encoding between LEMMA's DDML and Jolie APIs. Section 3 describes the architecture and implementation of LEMMA2Jolie. Section 4 presents future work and concludes the paper.

2 Encoding LEMMA Domain Modelling Concepts in Jolie

This section describes and exemplifies domain modelling with LEMMA (cf. Section 2.1), and the development of types and interfaces with Jolie (cf. Section 2.2). Next, it reports a formal encoding from LEMMA domain models to Jolie APIs and illustrates its application (cf. Sections 2.3 and 2.4).

$$
\begin{array}{lll}
CTX & ::= & \textbf{context}\ id\ \textbf{\{}\overline{CT}\textbf{\}} \\
CT & ::= & STR \mid COL \mid ENM \\
STR & ::= & \textbf{structure}\ id\ [\langle\overline{STRF}\rangle]\ \textbf{\{}\overline{FLD}\ \overline{OPS}\textbf{\}} \\
STRF & ::= & \textbf{aggregate} \mid \textbf{domainEvent} \mid \textbf{entity} \mid \textbf{factory} \\
 & & \mid \text{service} \mid \text{repository} \mid \textbf{specification} \mid \textbf{valueObject} \\
FLD & ::= & id\ id\ [\langle\overline{FLDF}\rangle] \mid S\ id\ [\langle\overline{FLDF}\rangle] \\
FLDF & ::= & \textbf{identifier} \mid \textbf{part} \\
OPS & ::= & \textbf{procedure}\ id\ [\langle\overline{OPSF}\rangle]\ \textbf{(}\overline{FLD}\textbf{)} \mid \textbf{function}\ (id \mid S)\ id\ [\langle\overline{OPSF}\rangle]\ \textbf{(}\overline{FLD}\textbf{)} \\
OPSF & ::= & \text{closure} \mid \textbf{identifier} \mid \text{sideEffectFree} \mid \textbf{validator} \\
COL & ::= & \textbf{collection}\ id\ \textbf{\{}(S \mid id)\textbf{\}} \\
ENM & ::= & \textbf{enum}\ id\ \textbf{\{}\overline{id}\textbf{\}} \\
S & ::= & \textbf{int} \mid \textbf{string} \mid \textbf{unspecified} \mid \ldots
\end{array}
$$

Figure 1. Simplified grammar of LEMMA's DDML. Greyed out features are out of the scope of this paper and subject to future work.

2.1 LEMMA Domain Modelling Concepts

LEMMA's DDML supports domain experts and service developers in the construction of models that capture domain-specific types of microservices. Figure 1 shows the core rules of the DDML grammar[4].

The DDML follows DDD to capture domain concepts. DDD's Bounded Context pattern [6] is crucial in MSA design as it makes the boundaries of coherent domain concepts explicit, thereby defining their scope and applicability [17]. A LEMMA domain model defines named bounded **context**s (rule CTX in Figure 1). A **context** may specify domain concepts in the form of complex types (CT), which are either structures (STR), collections (COL), or enumerations (ENM).

A **structure** gathers a set of data fields (FLD). The type of a data field is either a complex type from the same bounded context (id) or a built-in primitive type, e.g., **int** or **string** (S). The **unspecified** keyword enables continuous domain exploration according to DDD [6]. That is, it supports the construction of underspecified models and their subsequent refinement as one gains new domain knowledge [20]. Next to fields, **structure**s can comprise operation signatures (OPS) to reify domain-specific behaviour. An operation is either a **procedure** without a return type, or a **function** with a complex or primitive return type.

LEMMA's DDML supports the assignment of DDD patterns, called *features*, to structured domain concepts and their components. For instance, the **entity** feature (rule $STRF$ in Figure 1) expresses that a structure comprises a notion of domain-specific identity. The **identifier** feature then marks the data fields ($FLDF$) or operations ($OPSF$) of an **entity** which determine its identity. For

[4] The complete grammar can be found at `https://github.com/SeelabFhdo/lemma/blob/main/de.fhdo.lemma.data.datadsl/src/de/fhdo/lemma/data/DataDsl.xtext`.

compactness, we defer the detailed presentation of the considered DDD features to Section 2.4, when discussing their relationship with our encoding to Jolie.

The DDML also enables the modelling of **collection**s (rule COL in Figure 1), which represent sequences of primitives (S) or complex (id) values, as well as **enum**erations (ENM), which gather sets of predefined literals.

The following listing shows an example of a LEMMA domain model constructed with the grammar of the DDML [22].

```
context BookingManagement {
 structure ParkingSpaceBooking⟨entity⟩ {
  long bookingID⟨identifier⟩,
  double priceInEuro,
  function double priceInDollars
 }
}
                                                LEMMA
```

The domain model defines the bounded **context** *BookingManagement* and its **structure**d domain concept *ParkingSpaceBooking*. It is a DDD **entity** whose *bookingID* field holds the **identifier** of an entity instance. The entity also clusters the field *priceInEuro* to store the price of a parking space booking, and the **function** signature *priceInDollars* for currency conversion of a booking's price.

2.2 Jolie Types and Interfaces

Jolie interfaces and types define the functionalities of a microservice and the data types associated with those functionalities i.e., the API of a microservice. Figure 2 shows a simplified variant of the grammar of Jolie APIs, taken from [16] and updated to Jolie 1.10 (the latest major release at the time of writing).

$$
\begin{array}{lll}
I & ::= & \textbf{interface}\ id\ \{\overline{\textbf{RequestResponse}\ id(TP_1)(TP_2)}\} \\
TP & ::= & id \mid B \\
TD & ::= & \textbf{type}\ id :\ T \\
T & ::= & B\ [\{\overline{id\ C :\ T}\}] \mid \textbf{undefined} \\
C & ::= & [min, max] \mid * \mid \textbf{?} \\
B & ::= & \textbf{int}[(R)] \mid \textbf{string}[(R)] \mid \textbf{void} \mid \ldots \\
R & ::= & \textbf{range}([min, max]) \mid \textbf{length}([min, max]) \mid \textbf{enum}(...) \mid \ldots
\end{array}
$$

Figure 2. Simplified syntax of Jolie APIs (types and interfaces)

An **interface** is a 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—although Jolie also supports **oneWay**s, where the sender delivers its message to the receiver, without waiting for the latter to

process it (fire-and-forget), we omit them here because they are not used in the encoding (cf. Section 2.3). 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) (**int**egers, **string**s, etc.).

Jolie **type** definitions (TD) have a tree-shaped structure. At their root, we find a basic type (B)—which can include a refinement (R) to express constraints that further restrict the possible inhabitants of the type [8]. The possible branches 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.

Jolie data types and interfaces are technology agnostic: they model Data Transfer Objects (DTOs) built on native types generally available in most architectures [4].

Based on the grammar in Figure 2, the following listing shows the Jolie equivalent of the example LEMMA domain model from Section 2.1.

```
///@beginCtx(BookingManagement)
///@entity
type ParkingSpaceBooking {
 ///@identifier
 bookingID: long
 priceInEuro: double
}
interface ParkingSpaceBooking_interface {
 RequestResponse:
   priceInDollars(ParkingSpaceBooking)(double)
}
///@endCtx
```
Jolie

Structured LEMMA domain concepts like *ParkingSpaceBooking* and their data fields, e.g., *bookingID*, are directly translatable to corresponding Jolie **type**s.

To map LEMMA DDD information to Jolie, we use Jolie documentation comments (*///*) together with an @-sign. It is followed by (i) the string *beginCtx* and the parenthesised name of a modelled bounded context, e.g., *BookingManagement*; (ii) the DDD feature name, e.g., *entity*; or (iii) the string *endCtx* to conclude a bounded context. This approach enables to preserve semantic DDD information for which Jolie currently does not support native language constructs. The comments serve as documentation to the programmer who will implement the API. In the future, we plan on leveraging these special comments also in automatic tools (see Sections 2.4 and 4).

LEMMA operation signatures are expressible as **RequestResponse** operations within a Jolie **interface** for the LEMMA domain concept that defines the signatures. For example, we mapped the domain concept *ParkingSpaceBooking* and its operation signature *priceInDollars* to the Jolie interface *ParkingSpaceBooking_interface* with the operation *priceInDollars*.

2.3 Encoding LEMMA Domain Models as Jolie APIs

In the following, we report an encoding from LEMMA domain models to Jolie APIs that formalises and extends the mapping exemplified in Section 2.2. Figure 3 shows the encoding.

The encoding is split in three encoders: the *main* encoder $[\![\cdot]\!]$ walks through the structure of LEMMA domain models to generate Jolie APIs using the encoders for *operations* ($(\!(\cdot)\!)$) and for *structures* ($[\![\![\cdot]\!]\!]$), respectively.

The operations encoder $(\!(\cdot)\!)$ generates Jolie interfaces based on **procedure**s and **function**s in the given models by translating structure-specific operations into Jolie operations. This translation requires some care. On one hand, LEMMA's **procedure**s and **function**s are similar in nature to methods of OOP, since they operate on data stored in their defining structure. On the other hand, Jolie does not support objects in the OOP sense but rather separates data from code that can operate on it (operations). Therefore, the encoding needs to decouple **procedure**s and **function**s from their defining structures as illustrated in Section 2.2 by the mapping of the LEMMA domain concept *ParkingSpaceBooking* and its operation signature *priceInDollars* to the Jolie interface*ParkingSpaceBooking_ interface* with the operation *priceInDollars*.

Given a structure X, we extend the signature of its **procedure**s with a parameter for representing the structure they act on and a return type X for the new state of the structure, essentially turning them into functions that transform the enclosing structure. For instance, we regard a procedure with signature $(Y \times \cdots \times Z)$ in X as a function with type $X \times Y \times \cdots \times Z \to X$. This approach is not new and can be found also in modern languages like Rust [14,28] and Python [19]. The operation synthesised by the $(\!(\cdot)\!)$ encoder accepts the *id_type* generated by the $[\![\cdot]\!]$ encoder that, in turn, has a *self* leaf carrying the enclosing data structure (id_s). The encoding of **function**s follows a similar path. Note that, when encoding *self* leaves, we do not impose the constraint of providing one such instance (represented by the **?** cardinality), but rather allow clients to provide it (and leave the check of its presence to the API implementer).

The main encoder $[\![\cdot]\!]$ and the structure encoder $[\![\![\cdot]\!]\!]$ transform LEMMA types into Jolie types. **context**s translate into pairs of ///@*beginCtx*(*context_name*) and ///@*endCtx* Joliedoc comment annotations. All the other constructs translate into **type**s and their subparts. When translating **procedure**s and **function**s, the two encoders follow the complementary scheme of $(\!(\cdot)\!)$ and synthesise the types for the generated operations. The other rules are straightforward.

2.4 Applying the Encoding

This subsection illustrates the application of the encoding from Section 2.3 using the Booking Management Microservice (BMM) of a microservice-based Park and Charge Platform (PACP) modelled with LEMMA [22]. The PACP enables drivers of electric vehicles to offer their charging stations for use by others. Its BMM manages the corresponding bookings based on domain concepts that were designed following DDD principles [6] and expressed in LEMMA's DDML.

$$
\begin{aligned}
&\llbracket \textbf{context}\ id\ \{\overline{CT}\} \rrbracket &&= ///@beginCtx(id) \\
& && \quad \overline{\llbracket CT \rrbracket} \\
& && \quad ///@endCtx \\
&\llparenthesis \textbf{structure}\ id\ [\langle \overline{STRF} \rangle]\ \{\overline{FLD}\ \overline{OPS}\} \rrparenthesis &&= [\overline{///@STRF}]\ \textbf{interface}\ id_interface\ \{\overline{\llparenthesis OPS \rrparenthesis_{id}}\} \\
&\llparenthesis \textbf{procedure}\ id\ [\langle \overline{OPSF} \rangle]\ (\overline{FLD}) \rrparenthesis_{id_s} &&= \textbf{RequestResponse}:\ [\overline{///@OPSF}]\ id(id_type)(id_s) \\
&\llparenthesis \textbf{function}\ (S \mid id_r)\ id\ [\langle \overline{OPSF} \rangle]\ (\overline{FLD}) \rrparenthesis_{id_s} &&= \textbf{RequestResponse}:\ [\overline{///@OPSF}]\ id(id_type)((\llbracket S \rrbracket \mid id_r)) \\
&\llbracket \textbf{structure}\ id\ [\langle \overline{STRF} \rangle]\ \{\overline{FLD}\ \overline{OPS}\} \rrbracket &&= \textbf{type}\ \llbracket \textbf{structure}\ id\ [\langle \overline{STRF} \rangle]\ \{\overline{FLD}\} \rrbracket \\
& && \quad \overline{\llbracket OPS \rrbracket_{id}}\ \llparenthesis \textbf{structure}\ id\ [\langle \overline{STRF} \rangle]\ \{\overline{OPS}\} \rrparenthesis_{id} \\
&\llbracket \textbf{procedure}\ id\ [\langle \overline{OPSF} \rangle]\ (\overline{FLD}) \rrbracket_{id_s} &&= \textbf{type}\ id_type:\ \textbf{void}\ \{self?:\ id_s\ \overline{\llbracket FLD \rrbracket}\} \\
&\llbracket \textbf{function}\ (id_r \mid S)\ id\ [\langle \overline{OPSF} \rangle]\ (\overline{FLD}) \rrbracket_{id_s} &&= \textbf{type}\ id_type:\ \textbf{void}\ \{self?:\ id_s\ \overline{\llbracket FLD \rrbracket}\} \\
&\llbracket \textbf{collection}\ id\ \{(S \mid id_r)\} \rrbracket &&= \textbf{type}\ id:\ \textbf{void}\ \{\llbracket \textbf{collection}\ id\ \{(S \mid id_r)\} \rrbracket\} \\
&\llbracket \textbf{enum}\ id\ \{\overline{id}\} \rrbracket &&= \textbf{type}\ \llbracket \textbf{enum}\ id\ \{\overline{id}\} \rrbracket \\
&\llbracket \textbf{structure}\ id\ [\langle \overline{STRF} \rangle]\ \{\overline{FLD}\} \rrbracket &&= [\overline{///@STRF}]\ id:\ \textbf{void}\ \{\overline{\llbracket FLD \rrbracket}\} \\
&\llbracket S\ id\ [\langle \overline{FLDF} \rangle] \rrbracket &&= [\overline{///@FLDF}]\ id:\ \llbracket S \rrbracket \\
&\llbracket id_r\ id\ [\langle \overline{FLDF} \rangle] \rrbracket &&= [\overline{///@FLDF}]\ id:\ id_r \\
&\llbracket \textbf{collection}\ id\ \{S\} \rrbracket &&= id*:\ \llbracket S \rrbracket \\
&\llbracket \textbf{collection}\ id\ \{id_r\} \rrbracket &&= id*:\ id_r \\
&\llbracket \textbf{enum}\ id\ \{\overline{id}\} \rrbracket &&= id:\ \textbf{string}(enum(\overline{\text{``}id\text{''}})) \\
&\llbracket \textbf{int} \rrbracket &&= \textbf{int} \\
&\llbracket \textbf{unspecified} \rrbracket &&= \textbf{undefined}
\end{aligned}
$$

Figure 3. Salient parts of the Jolie encoding for LEMMA's domain modelling concepts.

In the following paragraphs, unless indicated, the encoded Jolie APIs respect the DDD constraints expressed by the considered features.

Aggregate and Part In DDD, aggregates prescribe object graphs, whose parts must maintain a consistent state [6]. Aggregates are always loaded from and stored to a database in a consistent state and within one transaction. A DDD aggregate consists of at least an entity or value object (see below). The following left listing shows the *PSB* aggregate in the LEMMA domain model for the BMM.

```
structure PSB
⟨ aggregate ⟩ {
 TimeSlot timeSlot ⟨ part ⟩,
 double priceInEuro
}
structure TimeSlot { ... }
                                  LEMMA
```

```
///@aggregate
type PSB {
 ///@part
 timeSlot: TimeSlot
 priceInEuro: double
}
type TimeSlot { ... }
                                  Jolie
```

PSB is a **structure**d domain concept with the **aggregate** feature (cf. Section 2.1) and it clusters the field *timeSlot*, which has a structured type and is

a **part** of the aggregate. Notice that for this domain model, LEMMA's DDML would emit warnings, because (i) a DDD aggregate must specify a root entity; and (ii) a part should either be an entity or value object [6]. We extend the *PSB* aggregate below to gradually fix these issues, thereby explaining the semantics of DDD entities and value objects.

In the Jolie encoding (on the right), we have as many **type** definitions as we have **structure**s in the LEMMA model.

Entity and Identifier Instances of DDD entities are distinguishable by a domain-specific identity [6], e.g., a unique ID. The following left listing extends the *PSB* aggregate with the **entity** feature and an **identifier** field.

```
structure PSB
⟨ aggregate, entity ⟩ {
 long bookingID ⟨ identifier ⟩,
 TimeSlot timeSlot ⟨ part ⟩,
 double priceInEuro
}
                                   LEMMA
```

```
///@aggregate
///@entity
type PSB {
 ///@identifier
 bookingID: long
 ///@part
 timeSlot: TimeSlot
 priceInEuro: double
}
                                   Jolie
```

LEMMA's DDML requires the **entity** feature on an aggregate to signal that its fields prescribe the structure of its root entity. The **identifier** feature can be used to mark those fields that determine the identity of an entity instances. In the example above, the value of *bookingID* is marked to identify *PSB*s.

The Jolie encoding of **entity** and **identifier** fields is straightforward.

Next to fields, DDML supports the **identifier** feature on a single **function** of an entity to enable identity calculation at runtime. To illustrate this approach, the following listing models the *bookingID* of the *PSB* root entity as a function.

```
structure PSB ⟨ entity ⟩ {
 function long bookingID
 ⟨ identifier ⟩ (),
 ...
}
                                   LEMMA
```

```
///@entity
type PSB { ... }
type bookingID_type { self?: PSB }
interface PSB_interface {
 RequestResponse:
  ///@identifier
  bookingID(bookingID_type)(long)
}
                                   Jolie
```

Following our encoding (cf. Section 2.3), we create the Jolie **type** *bookingID_type* for the *bookingID* **identifier function**. The **type**'s *self* leaf enables implementers to access the fields of the *PSB* and define how to compute the identifier.

Factory DDD factories make the creation of objects with complex consistency requirements explicit [6]. LEMMA's DDML considers factories to constitute **function**s that return instances of aggregates, entities, or value objects. The

following left listing illustrates the usage of factories by specifying the **factory** function *create* as part of the *PSB* aggregate. This function shall create *PSB* instances for a given time slot *timeSlot* and a *priceInEuro*.

```
structure PSB ⟨ ... ⟩{
 TimeSlot timeSlot,
 double priceInEuro,
 function PSB create⟨factory⟩(
  TimeSlot timeSlot,
  double priceInEuro
 )
}
LEMMA
```

```
type PSB { ... }
///@factory
type create_type {
 timeSlot: TimeSlot
 priceInEuro: double
}
interface PSBFactory_interface {
 RequestResponse:
  create(create_type)(PSB)
}
Jolie
```

As opposed to the encoding for LEMMA **identifier function**s (see above), we do not encode a *self* leaf in Jolie **type**s such as *create_type* for LEMMA **factory function**s. Since the semantics of factories is that of generating an instance of the enclosing **structure**, it would not make sense to pass to it one of those instances as a *self* leaf. Consequently, we could include a rule in Figure 3 which avoids the generation of said *self* leaf (this is more an issue of minimality of the generated code, since we set the leaf as optional (**?**)). Additionally, we can enforce a check on Jolie operations like *create* following immediately after ///@*factory*-commented **type**s by making sure their input **type**s do not contain the produced **type**, e.g., *PSB*. Complementary, we can also check that the response type of Jolie-encoded factory operations coincides with the produced **type**.

Specification and Validator DDD specifications are domain concepts that make business rules, policies, or consistency specifications for aggregates explicit [6]. A specification must comprise one or more validators, which are functions with a boolean return type that reify the specification's predicates.

LEMMA's DDML provides the features **specification** and **validator** to mark structures as specifications and identify their validators. The following left listing extends the BMM's domain model with the *BookingExpiration* specification. Its *isExpired* validator returns **true** if a parking space booking in the form a *PSB* instance has expired.

```
structure PSB ⟨ ... ⟩ { ... }
structure BookingExpiration
⟨ specification ⟩ {
 function boolean isExpired
 ⟨ validator ⟩ (PSB p)
}
LEMMA
```

```
type PSB { ... }
///@specification
type isExpired_type { p: PSB }
interface BookingExpiration_interface {
 RequestResponse:
  ///@validator
  isExpired(isExpired_type)(bool)
}
Jolie
```

Since the specification is a field-less **structure**, we do not create a corresponding **type** *BookingExpiration* as it would be empty. Instead, and as per our encoding (cf. Section 2.3), we create the ///@*specification*-annotated **type** *isExpired_type* for the *isExpired* **validator** within the **interface** *BookingExpiration_interface*. From the point of view of the consistency of the annotations, following the namespace convention from Figure 3, we can check that the ///@*validator* actually accepts the related **structure**. To do this, we follow the "breadcrumbs" left by our encoders. First, we find a ///@*validator*-commented **RequestResponse** (e.g., *isExpired*) and we make sure its response type is **bool**. Then, we follow the request type (e.g., *isExpired_type*) to make sure that: *i*) the ///@*validator* has an associated ///@*specification* (e.g., *isExpired_type*) **type** and *ii*) the **type** has one leaf, which is the **structure** the **validator** validates.

Notice that, here, we lose the enclosing relation between *BookingExpiration* and *PSB*. This might introduce subtle bugs (e.g., due to typos), since we have only one place (the request type of the ///@*validator*) that states the relation between the **specification** and the **validator**. To strengthen our checks, we might include a rule in Figure 3 which would include the reference to *PSB* in the annotation comment of the **validator**, e.g., ///@*validator(PSB)*. In this way, we can assert the correspondence between the validated **type** and the type of the other leaf in the request of the ///@*validator*.

Value Object and Domain Event As opposed to entities, DDD value objects cluster data and logic, which are not dependent on objects' identity [6]. Thus, value objects serve as DTOs for data exchange between microservices [17]. In asynchronous communication scenarios, value objects can model domain events emitted by a bounded context during runtime [6]. For example, all PACP microservices interact with each other via domain events [22].

LEMMA's DDML supports the **valueObject** and **domainEvent** features to mark structured domain concepts as value objects and possibly as domain events. The following left listing illustrates the usage of the **valueObject** feature.

```
context BookingManagement {
 structure PSB ⟨ ... ⟩ {
  TimeSlot timeSlot,
  double priceInEuro
 }
 structure PSB_VO⟨ valueObject ⟩ {
  TimeSlot timeSlot,
  double price,
  string currency
 }
 structure TimeSlot⟨ valueObject ⟩ {
  ...
 }
}
                                LEMMA
```

```
///@beginCtx(BookingManagement)
type PSB {
 timeSlot: TimeSlot
 priceInEuro: double
}
///@valueObject
type PSB_VO {
 timeSlot: TimeSlot
 price: double
 currency: string
}
///@valueObject
type TimeSlot { ... }
///@endCtx
                                Jolie
```

Above, we extend the BMM's domain model with the *PSB_ VO* value object: a DTO for the *PSB* aggregate that slightly changes it type to make its representation more general. Namely, *PSB_ VO* makes the *currency* explicit and separates it from the value of the *priceInEuro*, which we store in the field *price*. The *timeSlot* field remains the same, but we make sure it is also a **valueObject**.

The LEMMA domain model also shows the definition of bounded **context**s in the DDML. All three structures *PSB*, *PSB_ VO*, and *TimeSlot* are enclosed by the *BookingManagement* **context** on which the BMM operates exclusively.

The encoding from LEMMA to Jolie follows Figure 3 without exceptions. Notice, in particular, the "opening" ///@*beginCtx*(*BookingManagement*) and "closing" ///@*endCtx* comments for the context. With those comments, we are declaring that the types (and interfaces) that appear between them belong to the context *BookingManagement*. In LEMMA, contexts indicate a boundary within which (complex) types belonging in the same context can co-exist and interact (e.g., by being part of the inputs and output of **procedure**s and **function**s). Then, as seen above, **valueObject**s exist to allow data to cross boundaries, by defining data types (e.g., **structure**s) purposed to act as DTOs.

While the encoding from LEMMA's DDML ensures that, at the API level, the anti-corruption invariants defined by **context**s and **valueObject**s are preserved (e.g., there exists no **type** with leaves whose types belong in different contexts nor **interface**s belonging in a context that accept types from another context, unless ///@*valueObject*s), this is not the case for behaviour, which can arbitrarily combine data structures and operators.

A possible way to ensure the enforcement of LEMMA's DDML anti-corruption invariants also in behaviours (which will be subject to future work), is through the definition of static checks that trace the contexts in which values belong—from the types of the operations that generated them, via receptions—and prohibit mixing values that belong in different contexts (e.g., by forbidding to use them with operations belonging in different contexts, although their types might be compatible). This static check would also handle the exception of values whose types are annotated as ///@*valueObject*s, which are the only ones allowed to be used in a mixed way (i.e., in operations that take or produce ///@*valueObject*-annotated types.).

Additional Features Our encoding captures the **repository** and **closure** features of LEMMA's DDML (cf. Figure 1) without exceptions. Checks regarding the **sideEffectFree** feature follow the same considerations of the **valueObject** feature: we need to inspect a service behaviour to make sure its does not modify the values obtained from ///@*sideEffectFree*-commented operations. **service**s are a generalisation of the **specification** feature, where we have a structure that contains only **function**s and **procedure**s—LEMMA further refines **service**s into, **domainService**s, **infrastructureService**s, **applicationService**s, which are subject to future works.

3 LEMMA2Jolie: A Code Generator to Derive Jolie APIs from LEMMA Domain Models

This section presents our LEMMA2Jolie tool which makes the encoding presented in Section 2 practically applicable. In the sense of MDE, LEMMA2Jolie is a *model-to-text transformation* [3] that generates Jolie APIs from LEMMA domain models. Section 3.1 describes LEMMA2Jolie's architecture and Section 3.2 gives an overview of its implementation. More details about our implementation and the usage of LEMMA2Jolie are given in Appendix A.

3.1 Architecture

As depicted in Figure 4, LEMMA2Jolie consists of three phases to derive Jolie APIs from LEMMA domain models.

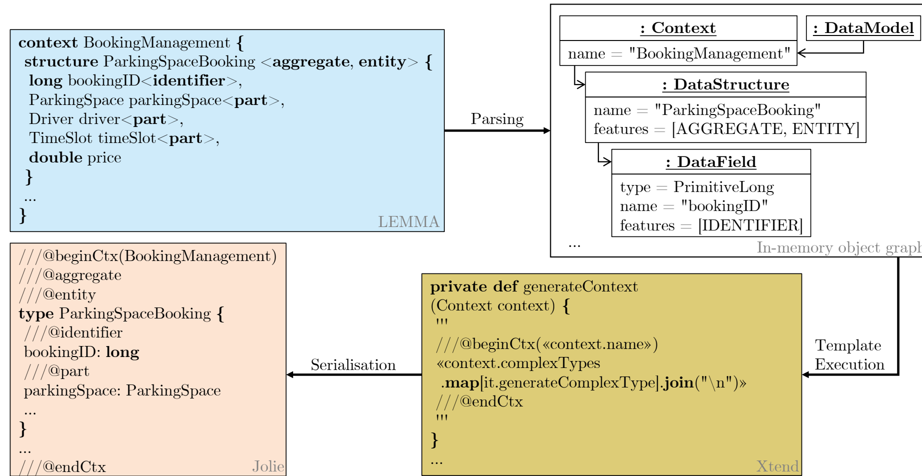

Figure 4. LEMMA2Jolie phases to generate Jolie APIs from LEMMA domain models.

In the Parsing phase, LEMMA2Jolie instantiates an in-memory object graph conforming to the metamodel of the DDML [21] from a given LEMMA domain model. The object graph allows systematic traversal of the model elements to map them to the corresponding Jolie code (cf. Section 2.3) in the following Template Execution phase. As the phase name indicates, LEMMA2Jolie relies on *template-based code generation* [3] to transform in-memory LEMMA domain models to Jolie. That is, we prescribe the target blocks of a Jolie program as strings involving static Jolie statements and dynamic variables which are evaluated at runtime to complement the prescribed target blocks with context-dependent information, e.g., the name of a bounded context in a specific LEMMA domain model. After template execution, the Serialisation phase stores the evaluated templates to physical files with valid Jolie code.

3.2 Implementation Overview

We implemented LEMMA2Jolie in Xtend[5], which is a Java dialect that integrates a sophisticated templating language (see below). Furthermore, LEMMA2Jolie relies on LEMMA's Java-based Model Processing Framework[6], which aims to facilitate the development of model processors such as code generators. To this end, the framework provides built-in support for parsing models constructed with languages that are based on the Eclipse Modelling Framework [26]—as is the case for all LEMMA modelling language including the DDML. Additionally, the framework prescribes a certain workflow for model processing and enables implementers to integrate with it using Java annotations.

Listing 1.1 describes the implementation of LEMMA2Jolie's code generation module which integrates with the Code Generation phase of LEMMA's Model Processing Framework. The module is responsible for template execution and the eventual serialisation of Jolie code (cf. Section 3.1).

A code generation module is a Java class with the `@CodeGenerationModule` annotation that extends the `AbstractCodeGenerationModule` class (Lines 1 and 2). LEMMA's Model Processing Framework delegates to a code generation module after it parsed an input model in the modelling language supported by the module. To specify the supported language, a code generation module overrides the inherited `getLanguageNamespace` method to return the language's namespace, which in the case of LEMMA2Jolie is that of LEMMA's DDML (Line 4).

The entrypoint for code generation logic is the `execute` method of a respective code generation module. It can access the in-memory object graph of a parsed model via the inherited `resource` attribute. Lines 6 to 13 show the `execute` method of LEMMA2Jolie's code generation module. In Line 7, we retrieve the root of the model as an instance of the `DataModel` concept of the DDML's metamodel (cf. Figure 4). Next, we call the template method `generateContext` (see below) for each parsed `Context` instance under the domain model root and gather the generated Jolie code as a list of strings in the `generatedContexts` variable (Line 8). In Lines 9 and 10, we then determine the path of the generated Jolie file, which will be created in the given target folder and with the same base name as the input LEMMA domain model but with Jolie's extension "ol". Line 11 triggers the serialisation of the generated Jolie code via the inherited `withCharset` method.

Lines 15 to 19 show the implementation of the template method `generateContext`. It expects an instance of the metamodel concept `Context` as input (cf. Figure 4) and represents the starting point of each template execution since bounded contexts are the top-level elements in LEMMA domain models. An Xtend template is realized between a pair of three consecutive apostrophes within which it is whitespace-sensitive and preserves indentation. Within opening and closing guillemets, Xtend templates enable access to variables and computing operations, whose evaluation shall replace a certain template portion. Consequently, the expression `«context.name»` in the template string in Line 16 is at runtime replaced

[5] `https://www.eclipse.org/xtend`

[6] `https://github.com/SeelabFhdo/lemma/tree/main/de.fhdo.lemma.model_processing`

Listing 1.1. Xtend excerpt of LEMMA2Jolie's code generation module.

```
1  @CodeGenerationModule(name="main")
2  class GenerationModule extends AbstractCodeGenerationModule {
3   ...
4   override getLanguageNamespace() { return DataPackage.eNS_URI }
5
6   override execute(...) {
7    val model = resource.contents.get(0) as DataModel
8    val generatedContexts = model.contexts.map[it.generateContext]
9    val baseFileName = FilenameUtils.getBaseName(modelFile)
10   val targetFile = '''«targetFolder»«File.separator»«baseFileName».ol'''
11   return withCharset(#{targetFile -> generatedContexts.join("\n")},
12    StandardCharsets.UTF_8.name)
13  }
14
15  private def generateContext(Context context) {'''
16   ///@beginCtx(«context.name»)
17   «context.complexTypes.map[it.generateComplexType].join("\n")»
18   ///@endCtx
19  ''' }
20
21  private def dispatch generateComplexType(DataStructure structure) {'''
22   «structure.generateType»
23   «IF !structure.operations.empty»
24    «structure.generateInterface»
25   «ENDIF»
26  ''' }
27 }
```

by the name of the bounded context passed to `generateContext`. For a bounded context with name "BookingManagement", Line 16 of the template will thus result in the generated Jolie code `///@beginCtx(BookingManagement)` (cf. Figure 4).

To foster its overview and maintainability, we decomposed our template for Jolie APIs into several template methods following the specification of our encoding (cf. Sect. 2.3). As a result, the generation of Jolie code covering the internals of modelled bounded contexts happens in overloaded methods called `generateComplexType`. Each of these methods derives Jolie code for a certain kind of LEMMA complex type, i.e., data structure, list, or enumeration. In Line 17, the template delegates to the version of `generateComplexType` for LEMMA data structures. Following our encoding, the method implements a template to map data structures to Jolie types (Line 22) and interfaces in case the LEMMA data structure exhibits operation signatures (Lines 23 to 25).

The LEMMA2Jolie source code is available on GitHub[7]. In addition, we provide a publicly downloadable video illustrating LEMMA2Jolie's practical capabilities[8].

[7] `https://github.com/frademacher/lemma2jolie`
[8] `https://bit.ly/3rTGysX`

4 Related and Future Work

Related Work The maturity of MDE in research and practice as well as its ability to effectively support the engineering of complex software systems [7] has fostered the development of a variety of tools similar to LEMMA2Jolie [23,13,12,27,11]. That is, they constitute code generators in the sense of MDE [3] and are capable to generate artefacts relevant to MSA engineering. For this purpose, the tools process models constructed in a certain modelling language.

However, and by contrast to LEMMA2Jolie, the majority of related code generators focuses on Java as target technology [23,27,11] and thus not on a programming language specifically tailored to the challenges of microservice implementation. Reducing the semantic gap between the concepts of microservices and implementation languages is the reason for which new service-oriented languages like Ballerina and Jolie have been developed. Furthermore, the modelling languages supported by related tools and hence the generated code address only single concerns in MSA engineering, i.e., domain modelling [23,13] or the implementation and provisioning of service APIs [12,27,11]. By contrast, LEMMA's modelling languages offer an integrated solution to multi-concern modelling in MSA engineering, by providing modelling languages dedicated to various viewpoints on microservice architectures (e.g., domain, service, and deployment) [21].

Future Work The specified encoding (cf. Section 2.3) and its implementation (cf. Section 3) show the feasibility to integrate the LEMMA and Jolie ecosystems. In future works we plan to extend this integration in several ways.

First, we plan to investigate the possibility of round-trip engineering (RTE), i.e., the bidirectional synchronisation of changes between LEMMA models and Jolie code. This would enable, for example, domain experts and microservice developers to interact by using their views of interest (model vs implementation) but without risking that they fall out of sync. While domain experts could continue to capture domain knowledge about a microservice architecture in conceptual DDD domain models, developers could adapt data types and APIs derived from those models using Jolie as their primary language. Based on RTE, changes in Jolie code could then automatically be reflected in DDD domain models and vice versa, with the option to immediately resolve potential conflicts in domain understanding.

Second, we see potential for LEMMA2Jolie to cover all phases in MSA engineering, from domain-driven service design to implementation and deployment. For example, we would like to extend LEMMA2Jolie to deal also with the definition of access points (communication endpoints that define how APIs can be accessed), behaviours (implementations of services written in Jolie that accompany LEMMA models), and the generation of deployment configurations (e.g., configuration of infrastructural services, containerisation, and deployment plans for Kubernetes). This potential is specifically fostered by both LEMMA and Jolie constituting *language-based approaches to MSA engineering*, which facilitates their integration. For example, we could extend LEMMA to include Jolie implementation code in service models.

References

1. Ameller, D., Burgués, X., Collell, O., Costal, D., Franch, X., Papazoglou, M.P.: Development of service-oriented architectures using model-driven development: A mapping study. Information and Software Technology **62**, 42–66 (2015), elsevier
2. Bogner, J., Fritzsch, J., Wagner, S., Zimmermann, A.: Microservices in industry: Insights into technologies, characteristics, and software quality. In: 2019 IEEE International Conference on Software Architecture Companion (ICSA-C). pp. 187–195. IEEE (Mar 2019). https://doi.org/10.1109/ICSA-C.2019.00041
3. Combemale, B., France, R.B., Jézéquel, J.M., Rumpe, B., Steel, J., Vojtisek, D.: Engineering Modeling Languages: Turning Domain Knowledge into Tools. CRC Press (2017)
4. Daigneau, R.: Service Design Patterns. Addison-Wesley (2012)
5. Dragoni, N., Giallorenzo, S., Lafuente, A.L., Mazzara, M., Montesi, F., Mustafin, R., Safina, L.: Microservices: Yesterday, today, and tomorrow. In: Mazzara, M., Meyer, B. (eds.) Present and Ulterior Software Engineering, pp. 195–216. Springer (2017)
6. Evans, E.: Domain-Driven Design. Addison-Wesley (2004)
7. France, R., Rumpe, B.: Model-driven development of complex software: A research roadmap. In: 2007 Future of Software Engineering. pp. 37–54. IEEE (2007)
8. Freeman, T., Pfenning, F.: Refinement types for ML. In: Proc. of the 1991 Conf. on Programming Language Design and Implementation. pp. 268–277 (1991)
9. Giallorenzo, S., Montesi, F., Peressotti, M., Rademacher, F., Sachweh, S.: Jolie and LEMMA: Model-driven engineering and programming languages meet on microservices. In: Coordination Models and Languages. pp. 276–284. Springer (2021)
10. Guidi, C., Maschio, B.: A jolie based platform for speeding-up the digitalization of system integration processes. In: Proceedings of the Second International Conference on Microservices (Microservices 2019) (2019), https://www.conf-micro.services/2019/papers/Microservices_2019_paper_6.pdf
11. JHipster: Jhipster domain language (jdl) (2022-14-02), https://www.jhipster.tech/jdl
12. Kapferer, S., Zimmermann, O.: Domain-driven service design. In: Service-Oriented Computing. pp. 189–208. Springer (2020)
13. Kapferer, S., Zimmermann, O.: Domain-specific language and tools for strategic Domain-driven Design, context mapping and bounded context modeling. In: Proceedings of the 8th International Conference on Model-Driven Engineering and Software Development - Volume 1: MODELSWARD. pp. 299–306. INSTICC, SciTePress (2020). https://doi.org/10.5220/0008910502990306
14. Klabnik, S., Nichols, C.: The Rust Programming Language (Covers Rust 2018). No Starch Press (2019)
15. Milner, R.: The tower of informatic models. From semantics to Computer Science (2009)
16. Montesi, F., Guidi, C., Zavattaro, G.: Service-oriented programming with jolie. In: Bouguettaya, A., Sheng, Q.Z., Daniel, F. (eds.) Web Services Foundations, pp. 81–107. Springer (2014). https://doi.org/10.1007/978-1-4614-7518-7_4, https://doi.org/10.1007/978-1-4614-7518-7_4
17. Newman, S.: Building Microservices: Designing Fine-Grained Systems. O'Reilly (2015)

18. Oram, A.: Ballerina: A Language for Network-Distributed Applications. O'Reilly (2019)
19. Python Software Foundation: The Python Language Reference (2021), https://docs.python.org/3/reference/index.html
20. Rademacher, F., Sachweh, S., Zündorf, A.: Deriving microservice code from underspecified domain models using DevOps-enabled modeling languages and model transformations. In: 2020 46th Euromicro Conf. on Software Engineering and Advanced Applications (SEAA). pp. 229–236. IEEE (2020)
21. Rademacher, F., Sorgalla, J., Wizenty, P., Sachweh, S., Zündorf, A.: Graphical and textual model-driven microservice development. In: Microservices: Science and Engineering. pp. 147–179. Springer (2020)
22. Rademacher, F., Sorgalla, J., Wizenty, P., Trebbau, S.: Towards holistic modeling of microservice architectures using LEMMA. In: Companion Proc. of the 15th Europ. Conf. on Software Architecture. CEUR-WS (2021)
23. Sculptor Team: Sculptor–generating Java code from DDD-inspired textual DSL (2022-14-02), https://www.sculptorgenerator.org
24. Soldani, J., Tamburri, D.A., 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), elsevier
25. Sorgalla, J., Wizenty, P., Rademacher, F., Sachweh, S., Zündorf, A.: Applying model-driven engineering to stimulate the adoption of devops processes in small and medium-sized development organizations. SN Computer Science **2**(6), 459 (2021)
26. Steinberg, D., Budinsky, F., Paternostro, M., Merks, E.: EMF: Eclipse Modeling Framework. Addison-Wesley, second edn. (2008)
27. Terzić, B., Dimitrieski, V., Kordić, S., Milosavljević, G., 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), taylor & Francis
28. The Rust Foundation: The Rust Reference (2021), https://doc.rust-lang.org/reference/

A Implementation and Usage of LEMMA2Jolie in Detail

Compared to Section 3, this appendix describes the implementation of LEMMA2-Jolie with LEMMA's Model Processing Framework in detail (cf. Appendix A.1). Furthermore, it explains the tool's usage (cf. Appendix A.2).

A.1 Implementation

Like all LEMMA modelling languages, the DDML's implementation is based on the Eclipse Modelling Framework (EMF) [26]. More precisely, the DDML's abstract syntax is implemented with Ecore, which is the metamodelling framework of EMF. The DDML's grammar for practical domain model construction was specified with the Xtext framework[9] for textual modelling languages. Xtext integrates with EMF and allows parser generation from grammar specifications. A generated Xtext parser is capable of translating textual model files to in-memory object graphs that constitute instantiations of the respective modelling language's metamodel. Such in-memory object graphs are then traversable for subsequent model processing steps using the EMF API.

LEMMA's Model Processing Framework makes the invocation of parsers for EMF-based modelling languages opaque to implementers, who instead gain direct access to in-memory object graphs as parsing results. The following paragraphs describe the implementation of each of LEMMA2Jolie's phases (cf. Section 3.1) with LEMMA's Model Processing Framework.

Parsing In the Parsing phase, LEMMA2Jolie parses a given LEMMA domain model into an in-memory object graph conforming to the metamodel of the DDML [21]. Using LEMMA's Model Processing Framework, we realised LEMMA-2Jolie's programmatic entrypoint as shown in Listing 1.2.

Listing 1.2. Programmatic entrypoint of LEMMA2Jolie written in Xtend.

```
1 class Lemma2Jolie extends AbstractModelProcessor {
2   new() { super("lemma2jolie") }
3   def static void main(String[] args) { new Lemma2Jolie().run(args) }
4 }
```

LEMMA's Model Processing Framework supports the development of model processors as standalone executable Java applications based on the `AbstractModelProcessor` class. Model processors extend this class (Line 1) and pass the name of a package to the framework in their constructor (Line 2). At runtime, the framework scans the passed package for annotated classes to invoke during model processor execution. Next to a constructor, entrypoints of model processors must implement a `main` method and delegate execution to the model

[9] https://www.eclipse.org/Xtext

processing framework by invoking the inherited `run` method with the given program arguments (Line 3).

Next to an entrypoint, LEMMA model processors must implement a *language description provider* to provide the model processing framework with information about the supported modelling language. Listing 1.3 shows LEMMA2Jolie's language description provider.

Listing 1.3. Xtend excerpt of LEMMA2Jolie's language description provider.

```
1  @LanguageDescriptionProvider
2  class LangDescriptionProvider implements LanguageDescriptionProviderI {
3   override getLanguageDescription(..., String namespaceOrExt) {
4    return switch (namespaceOrExt) {
5     case "data": new XtextLanguageDescription(DataPackage.eINSTANCE,
6      new DataDslStandaloneSetup)
7      ...
8    }
9   }
10 }
```

A language description provider is a class with the annotation `@LanguageDescriptionProvider`. Furthermore, the class must implement the interface `LanguageDescriptionProviderI` (Lines 1 and 2) and override its `getLanguageDescription` method (Lines 3 to 9), which LEMMA's Model Processing Framework will invoke to retrieve information about a supported modelling language. Model processors like LEMMA2Jolie can use the `namespaceOrExt` parameter to recognise the language of a given model file. Depending on the file format, the framework currently integrates language recognition based on XML namespaces or file extensions. Since our DDML is an Xtext-based modelling language and Xtext detects modelling languages from model files' extensions, LEMMA2Jolie checks the `namespaceOrExt` parameter for the value "data" (Line 5), which is the file extension for LEMMA domain models.

Upon modelling language recognition, a model processor must return an instance of the `LanguageDescription` class. It informs the model processing framework about the parsing mechanism to use for an input model in a certain language. In Lines 5 and 6 of its language description provider, LEMMA2Jolie returns a language description for Xtext-based modelling languages to the framework. Specifically, the corresponding `XtextLanguageDescription` clusters an instance of the `DataDslStandaloneSetup` class generated by Xtext. From this class, the model processing framework is able to trigger and control the parsing of LEMMA domain models into DDML-conform in-memory object graphs (cf. Section 3.1).

Template Execution In this phase, LEMMA2Jolie executes a template for Jolie APIs on the in-memory object graph of a parsed LEMMA domain model. We

used the integrated templating language of Xtend to formulate the template. Listing 1.4 shows an excerpt of the template.

Listing 1.4. Excerpt of our template for Jolie APIs in Xtend's templating language.

```
1  private def generateContext(Context context) {'''
2    ///@beginCtx(«context.name»)
3    «context.complexTypes.map[it.generateComplexType].join("\n")»
4    ///@endCtx
5   ''' }
6
7  private def dispatch generateComplexType(DataStructure structure) {'''
8    «structure.generateType»
9    «IF !structure.operations.empty»
10     «structure. generateInterface»
11   «ENDIF»
12  ''' }
```

Lines 1 to 5 show the implementation of the template method `generate-Context`. It expects an instance of the metamodel concept `Context` as input (cf. Figure 4) and represents the starting point of each template execution since bounded contexts are the top-level elements in LEMMA domain models. An Xtend template is realized between a pair of three consecutive apostrophes within which it is whitespace-sensitive and preserves indentation. Within opening and closing guillemets, Xtend templates enable access to variables and computing operations, whose evaluation shall replace a certain template portion. Consequently, the expression `«context.name»` in the template string in Line 2 is at runtime replaced by the name of the bounded context passed to `generateContext`. For a bounded context with name "BookingManagement", Line 2 of the template will thus result in the generated Jolie code `///@beginCtx(BookingManagement)` (cf. Figure 4).

To foster its overview and maintainability, we decomposed our template for Jolie APIs into several template methods following the specification of our encoding (cf. Sect. 2.3). As a result, the generation of Jolie code covering the internals of modelled bounded contexts happens in overloaded methods called `generateComplexType`. Each of these methods derives Jolie code for a certain kind of LEMMA complex type, i.e., data structure, list, or enumeration. In Line 3, the template in Listing 1.4 delegates to the version of `generateComplexType` for LEMMA data structures. Following our encoding, the method implements a template to map data structures to Jolie types (Line 8) and interfaces in case the LEMMA data structure exhibits operation signatures (Lines 9 to 11).

Serialisation In its last phase, LEMMA2Jolie serialises the results from template execution to physical files with Jolie code. To this end, we leveraged LEMMA's Model Processing Framework to implement a code generation module. Listing 1.5 shows an excerpt of its Xtend implementation.

Listing 1.5. Xtend excerpt of LEMMA2Jolie's code generation module.

```
1  @CodeGenerationModule(name="main")
2  class GenerationModule extends AbstractCodeGenerationModule {
3    ...
4    override getLanguageNamespace() { return DataPackage.eNS_URI }
5
6    override execute(...) {
7      val model = resource.contents.get(0) as DataModel
8      val generatedContexts = model.contexts.map[it.generateContext]
9      val baseFileName = FilenameUtils.getBaseName(modelFile)
10     val targetFile = '''«targetFolder»«File.separator»«baseFileName».ol'''
11     return withCharset(#{targetFile -> generatedContexts.join("\n")},
12       StandardCharsets.UTF_8.name)
13   }
14
15   /* cf. Listing 1.4 */
16   private def generateContext(Context context) { ... }
17   private def dispatch generateComplexType(DataStructure structure) { ... }
18 }
```

A code generation module in the sense of LEMMA's Model Processing framework is a Java class with the `@CodeGenerationModule` annotation and extending the `AbstractCodeGenerationModule` class (Lines 1 and 2). The model processing framework delegates to a code generation module after it parsed an input model with the namespace of the modelling language supported by the module. As LEMMA2Jolie parses LEMMA domain models (cf. Figure 4), the code generation module returns the namespace of LEMMA's DDML to the framework (Line 4).

The entrypoint for code generation logic is the `execute` method of a respective code generation module. The in-memory object graph of the parsed model is accessible via the inherited `resource` attribute. Lines 6 to 13 show the `execute` method in the code generation module of LEMMA2Jolie. In Line 7, we retrieve the root of the model as an instance of the `DataModel` concept of the DDML's metamodel (cf. Figure 4). Next, we call the template method `generateContext` (cf. Listing 1.4) for each parsed `Context` instance under the domain model root and gather the generated Jolie code as a list of strings in the `generatedContexts` variable (Line 8).

Finally, we determine the path of the file for the generated Jolie code, which will be created in the given target folder and with the same base name as the input LEMMA domain model but with the extension "ol" (Lines 9 and 10). The serialisation of the generated Jolie code is triggered by invoking the inherited `withCharset` method and returning its results to the framework. The method expects a map of file paths and contents, and the target encoding as argument. For LEMMA2Jolie's code generation module, the first argument associates the previously assembled path of the file for the generated Jolie code with the generated code concatenated in a string separated by line breaks (Line 11). As

the second argument of `withCharset`, we pass an identifier for UTF-8 encoding (Line 12).

A.2 Usage

LEMMA2Jolie integrates a commandline interface which is executable with Java 11 or greater. LEMMA2Jolie can be compiled to a standalone Java archive from its GitHub sources and run on physical hardware. Alternatively, it is possible to execute LEMMA2Jolie in a Docker container[10] for which we provide a dedicated Dockerfile on GitHub. In either case, the commandline invocation follows the pattern shown in Listing 1.6.

Listing 1.6. Commandline pattern for invoking LEMMA2Jolie.

```
1  java -jar lemma2jolie.jar -s <LEMMA_MODEL> -t <JOLIE_FOLDER>
```

The commandline option `-s` expects the path of a LEMMA domain model. Hence, LEMMA2Jolie assumes upfront construction of a LEMMA domain model, e.g., with LEMMA's editor plugins for the Eclipse IDE which provide sophisticated modelling support including syntax highlighting, code completion and cross-referencing [21]. The commandline option `-t` then points to the target folder for the generated Jolie file (cf. Sect. 3.2)

[10] `https://www.docker.com`