A Generative Programming Approach To Developing DSL Compilers

Charles Consel¹, Fabien Latry¹, Laurent Réveillère¹, and Pierre Cointe²

¹ INRIA / LaBRI ² École des Mines de Nantes
F-33402 Talence, France F-44070 Nantes, France

Abstract. Domain-Specific Languages (DSLs) represent a proven approach to raising the abstraction level of programming. They offer high-level constructs and notations dedicated to a domain, structuring program design, easing program writing, masking the intricacies of underlying software layers, and guaranteeing critical properties.

On the one hand, DSLs facilitate a straightforward mapping between a conceptual model and a solution expressed in a specific programming language. On the other hand, DSLs complicate the compilation process because of the gap in the abstraction level between the source and target language. The nature of DSLs make their compilation very different from the compilation of common General-Purpose Languages (GPLs). In fact, a DSL compiler generally produces code written in a GPL; low-level compilation is left to the compiler of the target GPL. In essence, a DSL compiler defines some mapping of the high-level information and features of a DSL into the target GPL and underlying layers (*e.g.*, middleware, protocols, objects, ...).

This paper presents a methodology to develop DSL compilers, centered around the use of generative programming tools. Our approach enables the development of a DSL compiler to be structured on facets that represent dimensions of compilation. Each facet can then be implemented in a modular way, using aspects, annotations and specialization. Because these tools are high level, they match the needs of a DSL, facilitating the development of the DSL compiler, and making it modular and re-targetable.

We illustrate our approach with a DSL for telephony services. The structure of the DSL compiler is presented, as well as practical uses of generative tools for some compilation facets.

1 Introduction

“Generative software development aims at modeling and implementing system families in such a way that a given system can be automatically generated from a specification written in one or more textual or graphical domain-specific languages [7].”

At the design level, generative software development emphasizes the role of Domain-Specific Languages (DSLs) as a way to bridge the gap between high-level modeling and General-Purpose Languages (GPLs). Nevertheless, from the

programming language viewpoint, there is a lack of methodology and a lack of tools to support the development of DSL compilers. The purpose of this paper is to apply some of the mainstream techniques promoted by the generative programming community to develop DSL compilers.

What is a DSL? From a programmer viewpoint, a DSL is typically created to model a program family [4]. The commonalities and the variabilities found in the target program family suggest abstractions and notations that are domain specific [6]. In contrast with GPLs such as Java or C++, a DSL has a narrow application scope and must be readable for domain experts.

In general, a program in a textual DSL is a concise set of high-level declarations, focusing on what to compute, as opposed to how to compute it. As an illustration, consider the telephony domain, and more specifically service creation. Conceptually, telephony services represent variations in creating, modifying and terminating a communication between parties. While this program family in a GPL covers the full implementation of services, its counterpart in a DSL corresponds to variations of service logic, abstracting over implementation details.

From a language designer viewpoint, a DSL makes domain-specific information an integral part of the programming paradigm. As such the programmer can be viewed as being prompted by the language to provide domain-specific information. This information may take the form of domain-specific types, syntactic constructs and notations. It serves domain-specific concerns, such as library interfacing, optimization, instrumentation, profiling and verification of the generated code. As of now, there are neither specific methodology, nor dedicated support tools, well suited to handle the compilation of both the high-level nature of a DSL and the richness of the built-in domain-specific information.

Challenges in DSL compilation. When mapping a DSL to a GPL, the higher level the DSL is, the more program generation is needed to bridge the gap with the target execution environment. Concretely, one program line written in a DSL commonly compiles into many lines in GPL. For instance, we have developed a DSL for telephony services, named SPL [2], that compiles into Java. An SPL program is on average 4 times more concise than its Java counterpart.

Not surprisingly, GPL-translated programs include rather large program templates. The process of generating these templates can be quite complex, relying on various conditions, and requiring a number of instantiations by computing and inserting constants. Without any dedicated tool support, this process can be quite laborious and error-prone. The resulting generator is often cumbersome and hard to debug. Additionally, the lack of tool support makes it difficult to have a modular treatment of the domain-specific concerns exposed by a program.

This paper. We propose a methodology to develop DSL compilers. The key idea of this methodology is to rely on generative programming tools [8]. These tools enable modeling the high-level nature of DSLs and the richness of the built-in domain-specific information in terms of program generation. This modeling

can be done in the context of various generative programming approaches. Our methodology is composed of two steps: compiling program logic and performing generative programming.

First, the DSL program logic is translated into an abstract GPL representation. This representation is abstract because it includes operations whose interpretation is not necessarily defined yet; it may thus not be executable. Although abstract, this translation generates a representation that encodes domain-specific information and that is localized in code regions of interest for further compilation treatment. In doing so, the representation is amenable to generative tools that represent the second step of our methodology.

Generative programming tools are used to define the compilation of domain-specific facets of a DSL in terms of program generation processes. One category of facets addresses the mapping of a DSL program into a target execution environment. A second category of facets is devoted to the compilation of language abstractions. A third category of facets defines code generation processes that are specific to the subject DSL program. This approach modularizes the process of generating a DSL compiler. Furthermore, each generative programming approach provides a paradigm, associated abstractions and tools dedicated to a specific family of program generation. The compiler developer can thus choose the most appropriate generative programming approach for a given facet.

Our approach is illustrated by three generative programming approaches, namely AOP [13], annotations [3], and program specialization [5, 10]; conceptually other approaches can be considered. AOP is well-suited to introduce cross-cutting behaviors in GPL-translated programs (*e.g.*, prologue and epilogue code of API invocations). Annotations enable non-functional concerns to be introduced in the compilation process (*e.g.*, resource management). Program specialization can address optimization-oriented program generation (*e.g.*, customization of software components).

Refining facets further leads us to distinguish between functional and non-functional facets. On the one hand, the functional facets define program generation processes that make the GPL-translated program executable. On the other hand, non-functional facets enrich the language execution in terms of requirements like performance, reliability and security.

Contributions.

- **A methodology to develop DSL compilers.** We define a methodology to develop a DSL compiler. We introduce a GPL-translated abstract representation of the program logic. This representation is structured so that further compilation can produce a spectrum of implementation variants.
- **A novel use of generative programming tools.** Our methodology relies on a novel use of generative programming. This methodology facilitates the generative programming process required by DSL compilation: high-level generative programming paradigms can be used to modularly process domain-specific information and abstractions.
- **A case study.** We use our methodology to develop a compiler for a DSL dedicated to the creation of telephony services. We present aspects, annotations

and specialization opportunities that model various compilation dimensions of this DSL.

Outline. The rest of this paper is organized as follows. Section 2 introduces our case study, a DSL for telephony services. Section 3 presents our methodology to develop a DSL compiler. Section 4 focuses on compiling the program logic. Section 5 and 6 presents the compilation of functional and non-functional language units. Finally, Section 7 discusses our approach, and Section 8 concludes.

2 Case Study

Our approach to developing DSL compilers is illustrated by a DSL to create telephony services. We choose the Session Initiation Protocol (SIP) [14, 15] as the underlying signaling protocol. This protocol is used for Voice over IP (VoIP) and third generation mobile phones. It is standardized by the IETF¹ and adopted by the ITU.²

SIP is a rapidly emerging protocol that places telephony into mainstream computer science. It is combined with a number of existing protocols to handle various aspects of communications (*e.g.*, transport and real-time streaming). It relies on the client-server model. SIP platforms provide rich programming interfaces in languages like C# and Java. The general-purpose nature of these programming interfaces make them very large and intricate to use. This situation is demonstrated by JAIN SIP, a standardized Java interface to SIP [9, 16], which consists of 130 classes and more than 3000 methods.

The need for a DSL in this domain stems from three main reasons. First, modern telephony is a software intensive area because of the host of new functionalities it offers. Second, programming telephony services now requires extensive expertise in SIP and its companion protocols, distributed programming, networking, and SIP programming interfaces. Third, a study of existing telephony services shows that programming a service logic in a given platform with a GPL is quite a laborious and error-prone process [1]. It requires recurring code patterns as the prologue and epilogue of invocations of the SIP programming interface.

We have developed a DSL named Session Processing Language (SPL) that enables telephony services to be defined concisely, using high-level abstractions and notations [2]. SPL enables the programmer to concentrate on the service logic, abstracting over low-level intricacies such as the protocol details and the underlying platform programming interface.

An SPL program defines a telephony service to which users can subscribe. The *session* is a key notion in SPL; it structures the development of a telephony service. A session consists of a set of handlers and a state. A handler defines a treatment for a protocol request or a platform event. A handler may be omitted,

¹ IETF: Internet Engineering Task Force.

² ITU: International Telecommunications Union.

if no service logic is associated with it. A state allows some data to be maintained across a set of handlers. SPL offers different kinds of sessions that form a hierarchy. A simple counter service written in SPL, exhibiting this hierarchy, is displayed in Figure 1.

```

1  service Counter {
2    processing {
3      local void log (int);
4
5      registration {
6        int count;
7
8        response outgoing REGISTER() {
9          count = 0;
10         return forward;
11       }
12
13       void unregister() {
14         log (count);
15       }
16
17       dialog {
18         response incoming INVITE() {
19           response resp = forward;
20           if (resp != /SUCCESS) {
21             return forward 'sip:secretary@company.com';
22           } else {
23             count++;
24             return resp;
25           }
26       }
27     }
28   }
29 }

```

Fig. 1. The counter service in SPL

The innermost session is the *dialog*: it manages a communication between parties. A dialog is created by the INVITE request, confirmed by the ACK request, and terminated by the BYE request. A dialog session defines handlers for the SIP requests and platform events pertaining to the communication management. In the counter service, an INVITE handler is defined to process incoming calls. This handler systematically *forwards* (i.e., routes) a call to the user who subscribed to this service. If this forward does not succeed (e.g., the user is busy), the call is forwarded to a secretary. Otherwise, the call is accepted, a counter is incremented, and the response is returned to the caller. The session variable **count** used in this handler is defined in the session surrounding a dialog, namely the *registration* session. Such a session is created for each user of the counter service that registers on the SIP platform by sending the REGISTER request. This session defines a state that consists of a unique variable (**count**), initialized in the REGISTER handler. A user is unregistered either when his registration lease expires or when the REGISTER request contains a zero-lease. Both situations are viewed by SPL as a unique event named **unregister**, for which a handler can be defined in the SPL program. In our example, the **unregister** handler invokes a function that logs the counter, when the session is terminated. At the top of the session hierarchy is the *service* session. Such a session is created when the service is deployed on the platform and deleted when it is undeployed.

SPL abstracts over a number of protocol and platform issues among which is the statefulness of a transaction. Let us explain this issue because it illustrates the gap between SPL and a SIP platform. This issue will later be used in examples. A transaction consists of a request and the response it triggers. Existing SIP platforms provide separate interface entries for processing requests and responses. A telephony program may process a request and its response independently; such a transaction is said to be *stateless*. Alternatively, a telephony program may need to match a response with the original request to continue some treatment initiated when the request got processed. Such a transaction is said to be *stateful* because it requires the platform to retain enough information to link a response to the original request. SPL abstracts over these implementation details. As shown in the `INVITE` handler, when the `INVITE` request is forwarded, the response is treated within the same handler. Statefulness of the transaction is determined by analyzing the service. In fact, SPL offers statefulness throughout the session hierarchy. In our example, the `count` variable is used in both the registration and dialog sessions. When a SIP message is processed by an SPL program, its GPL-translated version executes some code to trigger the corresponding handler and to extract the appropriate state.

3 A Methodology to Develop DSL Compilers

Our starting point in developing a DSL compiler is to define a direct translation of the *program logic* into a GPL representation. This translation is direct in that it generates abstract GPL code, not necessarily executable, where DSL mechanisms are not yet expanded but rather left uninterpreted. Although abstract, this translation encodes domain-specific information into the GPL representation, in a context where it can later be interpreted.

In our telephony case study, for example, while the routing of a SIP message in SPL is not concerned with statefulness, this property is explicitly encoded in the GPL-translated version.

The GPL-translated program logic can then be the input to various interpretations guiding two compilation dimensions: *functional* and *non-functional language units*. On the one hand, the functional units of the language are intended to complement the GPL-translated program logic to make it executable. On the other hand, the non-functional units of the language address implementation refinements or enrichments. The compilation of functional units varies with respect to both the target GPL and the target execution environment. In addition to these dimensions, the compilation of non-functional units may vary with respect to requirements like performance, reliability and security.

Whether or not functional, the compilation of DSL units can be decomposed into treatments that are inherent to either the target execution environment, the language or the program. These three categories of treatment are named *facets*.

The goal of the *execution environment facet* is to bridge the gap between the DSL execution model, possibly implicit, and the target execution environment. The DSL execution model is supposed to be high level and portable,

abstracting over the intricacies of the target execution environment. The execution environment facet is intended to generate the necessary code to interface the GPL-translated program logic with the underlying layers. Considering the SPL case, the execution environment facet bridges the gap between the explicit event-handling architecture of JAIN SIP and the implicit SPL model based on request handlers. For example, default request handlers need to be introduced to create and delete session states, whenever an SPL program does not define handlers for the corresponding requests.

The *language facet* is concerned with the interpretation and expansion of language mechanisms, whether or not explicit in the GPL-translated program logic. Such a facet generates recurring code patterns intrinsic to the language. For example, routing operations of all SPL services need to be analyzed to determine their statefulness. This information is made explicit in the GPL-translated program logic by adding some information in each invocation of the routing operations.

Lastly, the *program facet* invokes compiler treatments that are specific to the subject program. For instance, in the case of SPL, the request handler necessitates specific state extraction operations, depending on what session variables occur. Such a treatment is specific to a given SPL service and thus part of the program facet.

4 Compiling The Program Logic

The goal of compiling the logic of a program is to produce a representation that abstracts over implementation details while being amenable to generative programming tools. To be applied successfully, generative programming tools generally introduce some constraints on the program structure and may require program generation parameters. Program structure is needed to make the GPL-translated program logic match the granularity of the program entities manipulated by the generative programming tools. The program generation parameters are typically needed to customize the program generation process or the execution generated program. These parameters may correspond to compilation data that are domain-specific information and are encoded in the translated program or/and in generative programming declarations (*e.g.*, an aspect declaration).

Our approach enables the program generation process to be modularized: each module is responsible for a “slice” of program generation needed to compile a given DSL. This modularization is particularly well-fitted to explore the implementation scope offered by the high-level nature of a DSL.

An example of a GPL-translated program logic is displayed in Figure 2. It is the SPL program (showed in Figure 1) compiled into Java for a JAIN SIP interface [16]. This interface is typical of a client-server model in that it requires a telephony service to implement a **SipListener** interface, providing **processRequest** and **processResponse** methods. Additionally, a method is declared to handle various platform timeouts. Also, private methods have been introduced to handle the registration and un-registration of the service owner, as well as the forward-

ing of call invitations (responses are treated in `processResponse`). As can be noticed in `handler_INVITE`, for example, the service logic has been straightforwardly translated into Java: expressions manipulating variables are reproduced verbatim; message routing is translated into invocations of `sendRequest` for the SPL forward and `sendResponse` for the returning responses.

```

1 public class Counter implements SipListener {
2     [...]
3     private void handler_REGISTER (Request rq, String method) {
4         count = 0;
5         sendRequest (false, rq_request); /** SPL, line 10 **/
6     }
7     private void handler_unregister (Request rq, String method) {
8         local.log (count); /** SPL, line 14 **/
9         sendRequest (false, rq_request); /** SPL: default behavior **/
10    }
11    private void handler_INVITE (Request rq, String method) {
12        sendRequest (true, rq_request); /** SPL, line 19 **/
13    }
14
15    public void processRequest (RequestEvent requestEvent) {
16        String method = rq_request.getMethod();
17        if (method.equals (Request.REGISTER)) { /** SPL, line 5 **/
18            if (!registrar.hasExpiresZero (rq_request)) {
19                if (!registrar.hasRegistration (rq_request)) { /** SPL, line 8 **/
20                    handler_REGISTER (rq_request, method);
21                } else { /** SPL: default behavior **/
22                    sendRequest (false, rq_request);
23                }
24            } else if (registrar.hasRegistration (rq_request)) { /** SPL, line 13 **/
25                handler_unregister (rq_request, method);
26            } else { /** SPL: default behavior **/
27                sendRequest (false, rq_request);
28            }
29        } else if (method.equals (Request.INVITE)) { /** SPL, line 18 **/
30            handler_INVITE (rq_request, method);
31        } else { /** SPL: default behavior **/
32            sendRequest (false, rq_request);
33        }
34    }
35
36    public void processResponse (ResponseEvent responseEvent) {
37        String method = rs_request.getMethod();
38        rs_responseCode = rs_response.getStatusCode();
39        if (method.equals (Request.INVITE)) { /** SPL, line 18 **/
40            if (rs_responseCode >= 300) { /** SPL, line 20 **/
41                AddressFactory addressFactory = getAddressFactory();
42                SipURI sipURI = addressFactory.createSipURI ("secretary", "company.com");
43                rs_request.setRequestURI (sipURI);
44                sendRequest (false, rs_request); /** SPL, line 21 **/
45            } else {
46                count++;
47                sendResponse (true, rs_response); /** SPL, line 24 **/
48            }
49        } else { /** SPL: default behavior **/
50            sendResponse (false, rs_response);
51        }
52    }
53
54    public void processTimeout (TimeoutEvent timeoutEvent) {[...]}
55 }

```

Fig. 2. Java-translated program logic

Let us now examine how to compile the program logic using generative programming approaches.

4.1 Aspect-Oriented Programming

The AOP approach consists of a join point language, to identify locations in the program execution, and an advice language, to define additional behavior at these locations [11].

Our goal is to produce a GPL-translated program logic that matches the expressivity of the join point language. That is, to structure the translated program so that it matches the granularity of the join point identification and the semantics of the associated advice. Conceptually, to fit the AOP approach, DSL compilation must introduce specific code structuring and communicate compilation data.

The GPL-translated program logic needs some structuring when some code is to be inserted at some program point. Identifying the target program point requires language entities such as a method definition or invocation, variable occurrences, and declarations. As a result, a sequence of commands may need to be placed into a method in a translated program to enable code to be inserted before, after, or around its execution. The top of Figure 2 shows private methods that have been introduced by the DSL compiler to easily insert code as the prologue and epilogue of invocations of SPL request handlers.

Compiling the program logic also involves collecting or computing information about the DSL program that needs to be made explicit in the translated program. This strategy separates the production of the information and its exploitation, leaving its interpretation to other compilation passes. In the context of AOP, passing information can either be made by adding extra arguments to operations or by introducing specific instance variables. The latter case is further discussed in Section 5.2.

In principle, functional DSL units correspond to aspects that can be performed statically since, intuitively, these aspects refine the static compilation process. In contrast, non-functional units may compile into both static and dynamic aspects. Static aspects may be used to expand the implementation of specific operations, whereas dynamic aspects may define conditional monitoring actions.

A key advantage of our approach is to permit a DSL compiler to be designed and structured in terms of modules, that is, aspects. Each aspect defines a specific DSL behavior whose cross-cutting nature makes it an ideal target for AOP.

4.2 Annotations

Annotations can be included in the translated program to make some information explicit as the DSL program gets mapped into a lower-level representation. Annotations are traditionally used as extra information describing non-functional language issues.

Like aspects, annotations modularize the compilation of a DSL in that they introduce information that are later interpreted with respect to a given annotation processor.

Different sets of annotations have different goals. There may be annotations to make resource management explicit, or annotations intended to trigger check pointing of some state. When compiling the program logic, the aim is to generate annotations that are as self-contained as possible to minimize the work of the annotation processor occurring at a later phase. To do so, the location of an annotation is a key parameter. Also, arguments to the annotation may be needed as additional input to the annotation process.

4.3 Program Specialization

Program specialization is another use of a generative programming tool for DSL compilation. It addresses the optimization of DSL implementation building blocks. The idea is that a DSL can often be seen as a language gluing software components together. Because these components can be glued in a variety of contexts they must be highly generic. While this approach has obvious software engineering advantages, in practice, it may entail a significant performance penalty. To alleviate this problem, program specialization enables generic software components to be customized with respect to the context in which they are used. Because software components are invoked by compiler-generated code, the customization contexts can be determined by definition of the DSL compiler. Furthermore, since components are fixed, or slowly evolving, their *specializability* can be determined precisely. As a result, this program transformation can be made fully predictable, which is not the case for arbitrary program transformations like most program optimizations.

5 Compiling Functional Units

Compilation of functional units fits particularly well with AOP: it incrementally refines the semantics of language mechanisms, whether or not explicit in the source program. We do various forms of code expansion triggered by method names and instance variables.

We use a wide-spread and well-tested AOP tool, AspectJ [12], to define the compilation of various functional units needed for SPL. Limitations discussed in the following examples are not intrinsic to the AOP approach but are rather specific to AspectJ tool. These limitations are further discussed in Section 7.

5.1 Execution Environment Facet

The goal of functional execution environment facets is to bridge the gap between the DSL execution model and the target execution environment. In our case study, these facets are intended to generate the necessary code to interface the program logic with the underlying JAIN SIP platform.

```

public aspect Environment {

    public Request Counter.rq_request;
    public SipProvider Counter.rq_sipProvider;

    pointcut processRequest(): execution(public void Counter.processRequest (RequestEvent));
    before(RequestEvent rq_Event, Counter obj): processRequest() && args(rq_Event) && target(obj) {
        obj.rq_request = rq_Event.getRequest();
        obj.rq_sipProvider = (SipProvider) rq_Event.getSource();
    }
    [...]
}

```

Fig. 3. An aspect for an execution environment facet

The aspect program presented in Figure 3 introduces code that implements the SPL model in terms of the JAIN SIP event-handling architecture. To do so, the inserted code extracts the current request message from the `event` object generated by JAIN SIP. Dispatch over the type of request can then be done to invoke the appropriate Java-translated SPL handler. In addition, code to extract the `sipProvider` object is generated to enable further processing of the SIP message (*e.g.*, message headers and transaction creation).

Thanks to the functional execution environment facet presented above, the Java-translated program logic does not have to deal with the intricacies of the underlying JAIN SIP platform. Such a strategy helps to isolate target-dependent program generation in compiler modules, corresponding to aspects.

5.2 Language Facet

Functional language facets are concerned with the interpretation and expansion of language mechanisms. In our case study for example, state management and statefulness of routing operations require generating recurring code patterns. The corresponding aspects are discussed below.

State management. As illustrated in Figure 2, the Java-translated program logic does not include code for attaching a state to a session, and managing this state across the session life-cycle. For example, the state associated with a registration session needs to be created when processing a `REGISTER` request and deleted either when the request failed or the session ends. We have developed a language facet in AOP dedicated to state management. An excerpt of this facet is depicted in Figure 4.

The first advice specifies the code to execute for creating a registration state when processing a `REGISTER` request. However, since the handler for the `REGISTER` request is not mandatory in SPL, the corresponding method may not be present in the Java-translated program. If so, some code must be inserted to catch a `REGISTER` request and create the registration state. However, the pattern matching capability of the pointcut language does not permit to test the non-existence of a method invocation. To remedy this problem, this information is encoded as flags introduced into the aspect program. These flags enable the ap-

```

public aspect Language {

    private boolean handler_REGISTER = true;
    private boolean handler_REREGISTER = false;
    private boolean handler_unregister = true;

    pointcut register(): execution(private void Counter.handler_REGISTER (Request, String));
    before(Request rq, String method, Counter obj): register() && args(rq, method) && target(obj) {
        State state = new State();
        int ident = obj.env.getId (obj.rq_request);
        obj.env.setEnv (ident, state);
        obj.state = state;
    }

    pointcut processRequest(): execution(public void Counter.processRequest (RequestEvent));
    before(Counter obj): processRequest() && target(obj) {
        if (!handler_REGISTER) {
            String method = obj.rq_request.getMethod();
            if (method.equals (Request.REGISTER)) {
                if (!obj.registrar.hasExpiresZero (obj.rq_request)) {
                    if (!obj.registrar.hasRegistration (obj.rq_request)) {
                        State state = new State();
                        int ident = obj.lib.env.getId (obj.rq_request);
                        obj.lib.env.setEnv (ident, state);
                        obj.state = state;
                    }
                }
            }
        }
        [...]
    }
}

```

Fig. 4. An aspect for a language facet (1)

appropriate advice to be selected, as illustrated by the use of the `handler_REGISTER` flag.

Statefulness. In the Java-translated program logic, the first argument of the `sendRequest` method determines the statefulness of the routing operation. This information has been made explicit in the Java-translated program logic thanks to an analysis of the SPL program. The aspect program shown in Figure 5 describes the recurrent code fragment that needs to be executed when processing such an operation. In this example, calls to `sendRequest` are compiled into either stateless or stateful routing operations depending on the value of its first argument.

```

public aspect Language {
    [...]
    pointcut rq_sendRequest(): call(public void Counter.sendRequest (boolean, Request)) &&
        (withincode(public void Counter.processRequest (...)) || withincode(* Counter.handler_* (...)));
    void around(boolean b, Request r, Counter obj): rq_sendRequest() && args(b,r) && target(obj) {
        try {
            if (b) {
                ClientTransaction ct = obj.rq_sipProvider.getNewClientTransaction (r);
                ct.sendRequest();
                return;
            } else {
                obj.rq_sipProvider.sendRequest (r);
                return;
            }
        } catch (Exception ex) {}
    }
    [...]
}

```

Fig. 5. An aspect for a language facet (2)

The use of functional language facets enables the translated program to be closer to the program logic. For example, state management requires a precise understanding of the SIP protocol to determine when sessions need to be created and deleted. Such intricacies are factorized into the language facet, preventing the compiler writer to address all the DSL implementation issues at once.

5.3 Program Facet

Functional program facets are concerned with the generation of code that is specific to a program. For example, creation and manipulation of the state attached to a session depends on the session variables and the handlers that use them.

Such a situation is illustrated by the SPL program shown in Figure 1, where the `count` variable is used in the `REGISTER` handler (line 9). This variable occurrence leads to the declaration of an aspect, displayed in Figure 6, that inserts the class definition `State` and the methods to access the instance variable. Furthermore, this aspect specifies that each occurrence of the `count` variable in the Java-translated program logic must be replaced by an access to the state.

```
public aspect Program {
    public int Counter.count;
    public class State implements Lib.State {
        private int count;
        void setCount (int x) { count = x; };
        int getCount () { return count; };
    }

    pointcut set_count(): set (int Counter.count);
    void around(int count, Counter obj): set_count() && args(count) && target(obj) {
        obj.state.setCount(count);
    }

    pointcut get_count(): get (int Counter.count);
    int around (Counter obj): get_count() && target(obj){
        return obj.state.getCount();
    }
}
```

Fig. 6. An aspect for a program facet

Functional program facets allow defining DSL compilation at the granularity of a program, using implicit or explicit information from the source program.

6 Compiling Non-Functional Units

Non-functional DSL units are compiled by exploiting information that refines or extends the resulting program implementation. Just like functional units, compilation of non-functional units cover all of the DSL facets, that is, execution environment, language and program.

6.1 Execution Environment Facet

Program specialization [5, 10] has been successfully used to customize an execution environment according to a specific usage context. In our case study, the JAIN SIP platform is responsible for invoking specific methods in the Java-translated program logic for processing requests and responses. When no method is defined in the Java program, the platform does not need to pass the SIP message to the SPL service. Instead, it can directly perform the default platform behavior. Such a filtering of messages is similar to packet filtering in networking where only packets of interest are channeled to the application layer. Program specialization has already been applied in this context [18] and has demonstrated its benefits. In our case study, a similar strategy would aim to specialize the message filtering of the JAIN platform with respect to request names of interest to a program logic.

This approach enables automatically and systematically specializing highly generic software components, such as JAIN SIP components, according to a customization context derived from the program logic, such as an SPL service. Beyond performance, specialization opens opportunity to reduce the foot-print of a software layer.

6.2 Language Facet

Chander *et al.* [3] have recently proposed an approach to resource-bounds checking. Their approach limits the resource usage accordingly to a policy, specifying the resources that a program can use, along with the corresponding usage bounds. The key idea is to ensure that for each operation that consumes resources, an adequate amount is still available. They propose to annotate a program with a *consume e* command specifying that *e* units of resource are used, and with an *acquire e* command specifying that *e* units of resources must be reserved. One of the advantages of this approach is to make it possible to use a theorem-prover to prove that adequate checks are being performed to guarantee correct resource usage for a given program. The verifier is composed by two components: a safety condition generator that extracts logical predicates (safety conditions) whose validity implies resource-usage safety and a prover that proves the predicates.

Figure 7 illustrates the use of this approach. The left of this figure shows an SPL program where a list of *callers* is defined: three persons authorized to contact the service owner. If not authorized, a call is redirected to a list of four *operators* until one of them picks up the phone. In the telephony domain, the forwarding action represents a critical resource because it triggers a chain of operations that may be costly in the telephony platform.

The Java-translated program logic is displayed in the right of Figure 7. Annotations have been introduced for routing operations and the loop command. Through annotation analysis, we can determine that this example requires, at worst, the reservation of four resource units (*acquire 4*), corresponding to the case where no operator picks up the phone. As a result, we can assure that

<pre> service limit_forward { processing { uri<4> operators = <...>; uri<3> callers = <...>; registration { dialog { response incoming INVITE() { foreach (caller in callers) { if (FROM == caller) return forward; } return forward operators; } } } } } </pre>	<pre> public void processRequest(RequestEvent requestEvent) { [...] if (method.equals (Request.INVITE)) { //@ acquire 4 Header f_h = rq_request.getHeader(FromHeader.NAME); String FROM = f_h.toString(); int i = 0; int size = callers.size(); while (i < size) { String caller = (String)callers.get(i); if (FROM.compareTo (caller) == 0) { //@ consume 1 rq_sipProvider.sendRequest(rq_request); return; } i++; } //@ inv(i <= 3, 3 - i) //@ consume 4 rq_sipProvider.sendRequest(operators, rq_request); return; } [...] } </pre>
---	--

Fig. 7. Annotations for a language facet

enough resources have been reserved before the program execution. Note that our *consume* annotations are made explicit in the translated program but they could just as well be incorporated in a library by extending the existing JAIN API.

This example illustrates the use of an existing tool to perform some non-functional processing of DSL programs. That is, resource-bounds checking is introduced by re-using an existing annotation language and underlying tools. In doing so, the DSL compiler writer is guided by the annotation language to determine the required non-functional information to be provided in the translated program. This strategy prevents him from re-doing an analysis of this non-functional domain. Furthermore, annotations allow a modularization of the compiler in that their processing is left to a later phase, performed by dedicated existing tools.

6.3 Program Facet

Unlike functional *program facets* that define compiler treatments specific to a subject program, non-functional *program facets* correspond to information collected on a given program. If we consider the example shown in Figure 7, an extension of the annotation approach could collect the *acquire* annotations for each SPL handler, to compute the maximum number of routing operations performed by the SPL program. This number could then be used by a security policy of the execution environment. In doing so, some restrictions could be enforced on the number of routing operations performed by an SPL program, to preserve the platform performance. This enforcement can occur prior to the program execution with respect to currently available resources. This process could be seen

as admission control. The approach could be extended to all of the resources consumed by a program. As a result, a telephony service would be accompanied by a list of the resources required for its execution.

Note that this program facet is different from the language facet, presented in Section 6.2. Indeed, the language facet did not address resource consumption globally to the service; it only ensured that a consumed resource had been previously allocated.

An important issue in the telephony domain is service billing. By examining the kind of compilation treatments on billing operations, one can observe that it amounts to defining non-functional aspects, analogous to monitoring activities (*e.g.*, logging). As an example, some timer could be enabled in the handler that starts a call session, and disabled in the handler terminating a call.

7 Discussion

Our methodology for DSL compiler development is to translate the logic of a program into a GPL representation that is amenable to generative programming approaches. One of these approaches relies on AOP to introduce specific behaviors at some locations in the GPL-translated program.

Our methodology for DSL compiler development heavily relies on generative programming approaches and corresponding tools. In doing so, their features, or even limitations, need to be taken into account when developing the compiler. Concretely, limitations regarding AOP were discussed earlier. In fact, they did not concern the approach but rather the AspectJ tool used for our experiments. For example, variable introduction is only possible at the level of a class, not inside a method. Moreover, AspectJ aspects are context-insensitive in that they cannot directly manipulate the variables that are in the scope of a cross-cutting point. This feature would permit defining finer grained advice and improve the quality of the generated code.

The generality of meta-programming [17] makes it an alternative approach to AOP, as well as most other generative programming approaches. In this context, a DSL compiler resembles an interpreter annotated so as to execute the static language actions and to produce code for the dynamic language actions. Meta-programming tools give a fine-grained control over program generation, which can occur at any program point. A key issue that needs to be explored is how to modularize DSL compilation with meta-programming, to mimic what we do with AOP, annotations and program specialization.

8 Conclusion and Future Work

In this paper, we present a new methodology to develop DSL compilers. Our methodology is composed of two steps: compiling program logic and performing generative programming. Compiling a program logic produces a GPL-translated representation that abstracts over implementation details while being amenable

to generative programming tools. These tools allow to model the high-level nature of DSLs, and the richness of the built-in domain-specific information, in terms of program generation.

Our approach modularize the program generation process of a DSL compiler. Each generative programming approach provides a paradigm, associated abstractions and tools, dedicated to a specific family of program generation. The compiler developer can thus choose the most appropriate generative programming approach for a given compilation dimension.

We have used our methodology to develop a compiler for a DSL dedicated to the creation of telephony services. We present aspects, annotations, and specialization opportunities that model various compilation dimensions of this DSL. This case study demonstrates that a DSL can make use of generative programming approaches and techniques in very effective ways. In essence, a DSL exposes information about programs that can be mapped by our approach into the realm of generative programming.

Compared to traditional approaches for compiler development, our methodology enables to have a modular treatment of the domain-specific concerns exposed by a program. The resulting compilation process of a DSL is made simpler and less error-prone.

References

1. L. Burgy, L. Caillot, C. Consel, F. Latry, and L. Réveillère. A comparative study of SIP programming interfaces. In *Proceedings of the ninth International Conference on Intelligence in service delivery Networks (ICIN 2004)*, Bordeaux, France, October 2004.
2. L. Burgy, C. Consel, F. Latry, J. Lawall, N. Palix, and L. Réveillère. Telephony Software Engineering: A Domain-Specific Approach. Research Report RR-5548, INRIA, Bordeaux, France, April 2005.
3. A. Chander, D. Espinosa, N. Islam, P. Lee, and G. C. Necula. Enforcing resource bounds via static verification of dynamic checks. In *Proceedings of the 14th European Symposium on Programming (ESOP)*, volume 3444, pages 311–325. Springer-Verlag, 2005.
4. C. Consel. *Domain-Specific Program Generation; International Seminar, Dagstuhl Castle*, chapter From A Program Family To A Domain-Specific Language, pages 19–29. Number 3016 in Lecture Notes in Computer Science, State-of-the-Art Survey. Springer-Verlag, 2004.
5. C. Consel and O. Danvy. Tutorial notes on partial evaluation. In *Conference Record of the Twentieth Annual ACM SIGPLAN-SIGACT Symposium on Principles Of Programming Languages*, pages 493–501, Charleston, SC, USA, January 1993. ACM Press.
6. C. Consel and R. Marlet. Architecturing software using a methodology for language development. In C. Palamidessi, H. Glaser, and K. Meinke, editors, *Proceedings of the 10th International Symposium on Programming Language Implementation and Logic Programming*, volume 1490 of *Lecture Notes in Computer Science*, pages 170–194, Pisa, Italy, September 1998.

7. K. Czarnecki. Overview of generative software development. In *Proceeding of the Unconventional Programming Paradigm*. To appear as Springer LNCS (Fradet, P. editor), September 2005.
8. M. Czarnecki and U. Eisenecker. *Generative Programming: Methods, Tools, and Applications*. Addison-Wesley, 2000.
9. J. Deruelle, M. Ranganathan, and D. Montgomery. Programmable active services for JAIN SIP. Technical report, National Institute of Standards and Technology, June 2004.
10. N.D. Jones, C. Gomard, and P. Sestoft. *Partial Evaluation and Automatic Program Generation*. International Series in Computer Science. Prentice-Hall, June 1993.
11. G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, and W. G. Griswold. Getting started with AspectJ. *Communications of the ACM*, 44(10):59–65, October 2001.
12. G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, and W. G. Griswold. An overview of AspectJ. In *Proceedings of the 15th European Conference on Object-Oriented Programming (ECOOP)*, pages 327–353. Springer-Verlag, 2001.
13. G. Kiczales, J. Lamping, A. Menhdhekar, C. Maeda, C. Lopes, J.M. Loingtier, and J. Irwin. Aspect-oriented programming. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, volume 1241, pages 220–242. Springer-Verlag, 1997.
14. A. B. Roach. Session Initiation Protocol (SIP)-Specific Event Notification. RFC 3265, IETF, June 2002.
15. J. Rosenberg, et al. SIP : Session Initiation Protocol. RFC 3261, IETF, June 2002.
16. Sun Microsystems. The JAIN SIP API specification v1.1. Technical report, Sun Microsystems, June 2003.
17. W. Taha. *Domain-Specific Program Generation; International Seminar, Dagstuhl Castle*, chapter A Gentle Introduction to Multi-stage Programming, pages 30 – 50. Number 3016 in Lecture Notes in Computer Science, State-of-the-Art Survey. Springer-Verlag, 2004.
18. S. Thibault, C. Consel, J. Lawall, R. Marlet, and G. Muller. Static and dynamic program compilation by interpreter specialization. *Higher-Order and Symbolic Computation*, 13(3):161–178, September 2000.