

An Approach for the Systematic Development of Domain-Specific Languages

Mark Strembeck¹, Uwe Zdun²

¹*Institute of Information Systems, New Media Lab*

Vienna University of Economics and BA, Austria

mark.strembeck@wu-wien.ac.at

²*Distributed Systems Group, Information Systems Institute*

Vienna University of Technology, Austria

zdun@infosys.tuwien.ac.at

Abstract

Building tailored software systems for a particular application domain is a complex task. For this reason, domain-specific languages (DSLs) receive a constantly growing attention in recent years. So far the main focus of DSL research is on case studies and experience reports for the development of individual DSLs, design approaches and implementation techniques for DSLs, and the integration of DSLs with other software development approaches on a technical level. In this paper, we identify and describe the different activities that we conduct when engineering a DSL, and describe how these activities can be combined in order to define a tailored DSL engineering process. Our research results are based on the experiences we gained from multiple different DSL development projects and prototyping experiments.

1 Introduction

Software engineers have to deal with a wide variety of application domains originating from diverse economic and industrial sectors, such as airline companies, banks, car manufacturers, energy suppliers, hospitals, newspaper publishers, or telecommunication companies. Moreover, such business domains are subdivided in organizational domains, such as accounting, customer support, human resources, procurement, research, or sales. In addition to business and organizational domains, software engineers have to deal with orthogonal, technical domains concerning different functions of a software system, such as access control, quality of service, system backup, or workflow definition. While these enumerations of domains are far from being complete, the examples illustrate the diversity of application domains existing on different abstraction layers.

In each of these domains, human users have to fulfill certain tasks. Throughout this paper we use the term *domain expert* to refer to a human user who is a professional in a particular domain, such as a stock analyst in the investment banking domain, a physician in the health-care domain, or a network administrator in the domain of computer networks. In addition, in many domains there is a fine-grained differentiation between domain experts specializing in different sub-domains. For example, in the banking domain there are different types of customer consultants (e.g., for ordinary customers, for corporate customers, or for private wealth management), different types of experts for investment banking (e.g., for mergers and acquisitions, financial risk management, or stock and derivative trading), and different types of management experts (such as business process experts or legal experts).

Building software systems that are tailored to the needs of a particular type of domain expert is already a complex task. This task becomes even more complex due to the continuing specialization of professional knowledge in all kinds of application domains. Therefore, the construction of software systems demands systematic development approaches that enable software engineers to cope with this situation. However, in spite of the advances in different software engineering disciplines, the problem of building tailored software systems for domain experts – in arbitrary domains – is far from being solved. For this reason, *domain-specific languages* (DSLs) receive a constantly growing attention in recent years (see, e.g., [33, 44, 61, 65, 74, 75]). DSLs are specifically tailored for the needs of a particular problem or application domain. They promise that domain experts themselves can understand, validate, modify, test, and sometimes even develop DSL programs.

Unfortunately, the systematic development of DSLs is not fully tackled yet. Most existing approaches especially focus on technical facets of designing and implementing DSLs. That is, the main focus of DSL research is up to now on case studies and experience reports for the development of individual DSLs [6, 42, 68, 74], design approaches and implementation techniques for DSLs [22, 28, 35, 42, 44, 51, 60, 71], and the integration of DSLs with other software development approaches, such as programming languages for embedding DSLs [18, 28], model-driven software development [11, 23, 33, 61], or component-based software development [8]. However, the process of defining the DSL and the way this process is used in the context of a general purpose software development approach, such as the Unified Process (UP) [31] or agile approaches like Extreme Programming (XP) [4], are seldom addressed.

In this paper, we identify and describe the different activities that we conduct when engineering a DSL. The activities documented in this paper are a result of the experiences we gained from numerous different DSL development projects and from multiple prototyping experiments. In particular, we performed an in-depth analysis of 14 DSL projects we participated in. In addition, we conducted several experiments with students and applied qualitative research methods such as observations, conversations, and interviews with professional developers to generalize and validate our results on a broad basis. Details about lessons learned from our experiences in various projects can be found in Section 8.

Our experiences in building DSLs especially result from projects in the model-driven software development (MDSD) context, projects where we built DSLs as extensions for dynamic programming languages, and projects

where these two approaches overlap (see Sections 3 and 8). In addition, other approaches to building DSLs exist, such as the Grammarware approach [34] which is focusing on grammars and all software artifacts that directly depend on grammars. The work that is presented in this paper focuses on DSLs in an MDSD context and DSLs as extensions for dynamic programming languages; the Grammarware view on DSLs is not explicitly addressed in this paper.

Our analyses revealed the different activities that we conduct when engineering a DSL. These activities form a micro-process that can be tailored to various influencing factors of an actual DSL project. Examples of such influencing factors are the general purpose software development approach used in a certain project (in which the micro-process for DSL development has to be embedded), the availability of domain experts, or the budget of the project. In particular, we found that one of the main challenges for software engineers is to increase the participation of domain experts in DSL design activities, in order to avoid misinterpretations of the domain and/or designing a DSL detached from the needs of the domain experts. In this paper, we describe both, the general activities that we perform to engineer a DSL, as well as activities we apply to tailor the DSL engineering process.

The remainder of this paper is structured as follows. In Section 2 we discuss related work to illustrate the state-of-the-art and how our approach contributes to the state-of-the-art. Next, in Section 3, we give essential background information on DSLs and introduce the different DSL concepts that we refer to in subsequent sections. In Section 4, we give an overview of our approach for the systematic development of DSLs. Subsequently, Section 5 presents the details of our approach by illustrating the individual process activities for DSL engineering using one specific variant of the DSL engineering process. Next, in Section 6 we describe how the DSL development process can be tailored in other ways than the prototypical tailoring presented in Section 5. To exemplify our approach, Section 7 provides an illustrative example which describes the engineering of a DSL for role-based access control (RBAC). Section 8 describes lessons learned from various projects, before Section 9 concludes the paper.

2 Related Work

To the best of our knowledge, the work presented in this article is the first approach to identify the development activities and describe a generic process for DSL development. However, a number of previous approaches also address specific process aspects of DSL development. In this section, we summarize them and compare them to our approach.

2.1 Related Work on Patterns, Best Practices, and Lessons Learned in DSL Development

Several authors introduced patterns and pattern languages that can be applied in DSL development. This includes patterns for the design and implementation of DSLs [60] and patterns for evolving frameworks into DSLs [54, 55]. A pattern is a time-proven solution to a recurring design problem. The patterns do not only describe how a DSL is developed, but also why it is developed in a specific way. In this sense, they go beyond our approach and

can be applied in the context of our approach. In addition, patterns in pattern languages also describe the potential sequences in which the patterns can be applied. This includes a process aspect, similar to our work. However, the named pattern languages rather focus on specific design issues, whereas our approach focuses on describing the overall micro-process for DSL development.

The patterns for the design and implementation of DSLs by Spinellis [60] have been used by Mernik et al. [44], who provide a survey of decision factors for the decision, analysis, design, and implementation phases of DSL development. These decision factors can be considered during the DSL development. For example, the decision factor *Notation* deals with the consideration whether the DSL should provide a new or existing domain notation. For a few decision factors, Mernik et al. suggest implementation guidelines, which describe a process aspect. This process aspect focuses on the general design decision and implementation issue, and not the general DSL development micro-process as in our approach.

Some other works than our approach and the patterns approaches are based on observations from practical DSL development. For example, Luoma et al. [42] conducted a study including 23 industrial projects for the definition of domain-specific (graphical) modeling languages. As in our approach, a number of DSLs are systematically compared. However, only one kind of DSL development process, the language model driven process (introduced below), is used and they focus on the discussion of factors that drive language construct identification. Our approach also covers the integration of DSL concepts with a target platform and addresses the tailoring of the DSL development process.

Similar to patterns, lessons learned have been used as a vehicle to convey best practices of DSL development. For example, Wile reports on twelve lessons learned from three DSL experiments [74]. For each lesson he introduces a respective rule of thumb and gives an overview of the experiences that are the origin of the corresponding rule. Our approach is based on a conceptual model of DSLs which incorporates the concepts identified in the above mentioned approaches and adds a process-oriented perspective that can be combined with these approaches. However, in general the lessons learned reported by Wile can also be observed in our projects and are hence reflected in parts of our process models.

In summary, the related work on patterns, best practices, and lessons learned in DSL development have in common with our approach that they all derive from experiences and contain some process aspects, but none of them concentrates on describing the DSL development process explicitly. They all focus only on specific, mostly technical issues, whereas our approach provides a more holistic, process-centric view.

2.2 Related Work on Model-Driven and Dynamic Languages DSLs

In this article, we especially focus on DSLs in an MDSD context and embedded DSLs for dynamic languages, as our research method is based on the observation of our experiences and most of our own projects have been in these fields (see also Sections 3 and 8). A number of approaches in these fields confirm our observations and report on similar development activities as the ones we have observed in our own projects.

For example, supporting model-driven development using several DSLs implemented in Ruby [11] can be achieved by implicitly using a similar process of developing such DSLs as the one we observed in our projects. This has been demonstrated for four embedded DSLs for Ruby that support basic tasks in model-driven development: a DSL to create language models, a DSL to specify restrictions on models, a model-to-model transformation DSL, and a model-to-code language. This study independently confirms our observations for technical DSLs developed in dynamic programming languages.

Others examples that independently confirm our results are based on the idea to build DSLs from component building blocks that can be incrementally designed and composed [1, 41, 51, 68, 70]. This idea originates from approaches such as keyword-based programming [8], in which so called “keywords” serve as building blocks for DSLs. In particular, a number of (universal) keywords are suggested which are then glued together to compose DSLs. This approach was first envisioned in [38] and is akin to building embedded DSLs in dynamic languages (such as Ruby, Perl, Python, or Tcl for example). These approaches independently confirm some of the key artifacts we observed in our projects, as well as the process activities related to those artifacts.

DSLs developed in dynamic programming languages are usually build by tailoring an existing general-purpose language (GPL) by adding or changing methods, operators, and default actions [19]. This approach inherently adds activities to the DSL engineering process for either adding methods to the language, or finding proper abstractions in the GPL and changing them in order to best suit the needs of the DSL. Our DSL engineering process can integrate these additional activities. In contrast to the implicit process for building a DSL described in [19], our development process adds a systematic approach for engineering not only the appearance of the DSL in the GPL, but also the models defining the DSL.

Similar to the language-oriented approaches, we also find many independent confirmations of our findings in some model-based or model-driven approaches. For example, similar processes can be used to define domain-specific (UML-based) modeling languages [22]. In particular, that means to define domain-specific modeling languages via UML profiles [49], an approach that is commonly used when defining domain-specific modeling languages (see, e.g., [61]). This approach shares the key artifacts that we observed when an external modeling DSL is build.

Similar approaches have also been presented for non-standard modeling tools, such as MetaCase’s MetaEdit+ tool [33]. In this approach, the steps and considerations for developing DSMs are pretty similar to the processes observed in our work. However, the processes described in [33] are based on experiences only with the MetaEdit+ tool. As our approach is independent of a specific tool or modeling approach, it is more generic. That is, our approach can be tailored to develop DSM languages, and in addition it comprises other approaches for DSL development. However, concerning other facets of DSL development, the approach in [33] also address topics which are out of scope in our approach. For example, it also reports on experiences regarding organizational issues, such as running small proof of concept projects to convince executives, or on criteria for creating a team that builds a certain DSM solution.

A number of other independent confirmations of our finding can be found in the model-driven literature. In [36], Kurtev et al. identify typical problems that may arise when developing a DSL. Moreover, they describe how these problems can be addressed with a model-based DSL framework called AMMA (ATLAS Model Management Architecture, see [2]). AMMA was built on top of the Eclipse Modeling Framework (EMF) [13] and provides tool support for the definition of DSLs. The AMMA core consists of the KM3 meta-model definition language, the TCS language for the definition of concrete syntaxes, and the ATL model transformation language. AMMA does not provide means for defining a DSL’s (behavioral) semantics, but Kurtev et al. are experimenting with abstract state machines as a foundation for modeling behavioral DSL semantics with AMMA. Other examples for tools that support DSL definition are the Generic Modeling Environment (GME, see [20, 40]) and Microsoft’s Domain-Specific Language tools (see [10]). In [10], Cook et al. describe Microsoft’s approach for domain-specific development with Microsoft Visual Studio DSL tools. They suggest an approach for designing a DSL and provide an extensive tutorial for creating DSLs in Visual Studio. In [3], Balasubramanian et al. describe two domain-specific modeling languages that were built using GME. The process aspects of these approaches are mainly implicitly reported, e.g. via examples. These examples confirm that the named approaches follow one of the process tailorings we propose. This article additionally provides a systematic approach to describe the DSL engineering process – including different ways to tailor it.

2.3 Related Work on Grammarware DSLs

Grammarware is defined as a term that comprises grammars (such as context-free grammars or XML schemas) and grammar-dependent software (like parsers, program converters, or XML document processors) [34]. The DSLs we analyzed for this article are mostly built using dynamic languages or model-driven tools (see also Section 8). Hence, the grammarware approach has not been sufficiently reflected in our experiences, and it might be possible that the process models must be adapted to be used with the grammarware approach. However, as the approaches based on grammarware and model-driven development are quite close (for instance, many model-driven, textual DSL suites, such as XText of openArchitectureWare [50], use grammarware inside), we expect the necessary alterations to be rather limited.

2.4 Approaches Dealing with DSL Evolution and Maintenance

One important aspect of our process models is that they consider DSL evolution and maintenance as an integral part of an incremental development process. In contrast, some approach specifically focus on evolution and maintenance as separate aspects. For example, Bierhoff et al. investigate the incremental development of a DSL based on existing example applications for a target domain [6]. This means, they chose an application and on top of that they (ex-post) develop a DSL that is expressive enough to describe the application’s functionality. This approach represents an innovative version of the “Extracting DSL from existing system” process variant that we identified in our work (see also Section 6). van Deursen and Klint describe report on the experiences they

gained with Risla [71], a DSL for financial engineering, and especially focus on DSL maintenance issues. In our approach, we have incorporated some of the process aspects of maintenance through feedback loops and iterative development. Wang and Gupta [72] present an approach to rapidly adapt the DSL development environment (consisting of components like compiler, interpreter, or debugger) to changes in the respective DSL. In particular, Wang and Gupta express DSL syntax and semantics using Horn logic. Because the corresponding horn logic expressions are executable, they automatically yield an interpreter for the respective DSL. Based on this DSL interpreter and an actual DSL program, Wang and Gupta then use partial evaluation [32] to automatically generate compiled code. Each time the DSL's syntax and semantics are changed, the corresponding interpreter can be automatically derived from the changed horn logic specification. Our approach deals with DSL evolution in terms of feedback loops and iterative development. Hence, the approach by Wang and Gupta can be seen as a specific concretization of some steps in our development process.

In our approach, in contrast to the above mentioned approaches, maintenance and evolution of DSLs are not considered as separate process aspects, but rather as an integral aspect of incremental DSL development. Apart from this aspect, all named maintenance and evolution approaches can be used within our DSL engineering process.

3 Domain-Specific Languages

A *domain-specific language* (DSL) is a tailor-made (computer) language for a specific problem domain. In this sense, DSLs differ from general purpose languages (GPL), such as C, C#, Java, Perl, Ruby, or Tcl, that can be applied to arbitrary problem domains. In particular, DSLs are often not Turing complete, and they only provide abstractions suitable for one particular problem domain. However, this specialization results in significant gains in expressiveness and ease of use in the DSL's application domain compared to a GPL (see, e.g., [42, 44]): Due to this expressiveness, a user typically needs to use less language instructions to achieve a certain result in a DSL, compared to achieving the same result in GPLs. A significantly smaller number of instructions leads to programs which are more easy to read and comprehend.

One of the very first papers describing the idea of defining tailored (programming) languages from a set of predefined primitives was presented by Landin [38]. In recent years, domain-specific languages received a constantly growing attention, especially in the area of model-driven software development (MDSD, see, e.g., [23, 33, 57, 58, 61]). The basic idea of domain-specific languages, however, already has a long history (see, e.g., [5, 24]). In the Unix context, for example, DSLs have a tradition as so called “little languages” or “mini languages” (see [5, 53]), and in the context of languages such as Common Lisp the development of “embedded languages” is promoted (see [21]). Other popular examples of DSLs are \LaTeX for Typesetting [37, 39], HTML for hypertext web pages [52], the (extended) Backus-Naur Form (BNF) for syntax specification [29], or SQL for database queries and manipulation [7, 30].

In general, DSLs can be designed and used on different abstraction layers, ranging from DSLs for technical

tasks to DSLs for tasks on the business-level. Thus, DSLs can be defined for non-technical stakeholders, such as business analysts or biologists, for example. In general, a DSL makes domain-knowledge explicit. That is, the DSL is built so that domain experts can understand and modify DSL code to phrase domain-specific statements that are understood by an information system. An important design principle of DSLs is therefore to keep the DSL as simple as possible while making it as powerful as needed.

For the purposes of this paper, we distinguish three main areas for the development of DSLs:

- DSLs in model-driven software development (MDS): here DSLs are often defined as domain-specific modeling languages that include an infrastructure for the automated transformation of DSL-based models to source code artifacts (see, e.g., [23, 33, 50, 57, 58, 61]).
- DSLs as extensions of (dynamic) programming languages: here DSLs are domain-specific extensions for GPLs (see, e.g., [19, 65, 76]). Such DSLs are defined as embedded languages that use the language infrastructure of their host language and extend the host language with domain-specific language abstractions.
- DSLs in the grammarware context: here DSLs are based on grammar formalisms and grammar notations and are often specified for the definition of special-purpose document types, for example defined via XML schemas (see, e.g., [34]).

In actual engineering projects, these three areas are not disjunct but may have significant overlaps (for instance, many model-driven, textual DSL suites, such as XText of openArchitectureWare [50], use grammarware inside). As mentioned above, our experiences in developing DSLs result from projects that either applied MDS techniques, defined an embedded DSL as an extension for a dynamic programming language, or combined these two approaches (see also Section 8). Therefore, we do not explicitly address the grammarware view on DSLs, even if some of our findings may also be applicable to this field of research as well.

Figure 1 depicts the main DSL artifacts and their relations¹. From a language user's point of view, the DSL consists of language elements. The definition of these language elements is provided by the language model. The DSL's *language model* specifies elements from the DSL's target domain. Some authors also call the language model the *abstract syntax* of the DSL (see, e.g., [23]). In essence, the language model is a composite model that consists of three sub-models (see Figure 1).

The *core language model* captures all relevant domain abstractions and specifies the relations between these abstractions. In particular, these abstractions refer to elements of the DSL's *target domain*. Examples of domain abstractions are account, bond, client, fund, stock, or stock order in the banking domain; pre-condition, post-condition, test case, or test result in the domain of software testing; role, subject, permission, or constraint in the

¹Please note that these artifacts are relevant for any DSL. For many existing DSLs, however, not each of these artifacts is necessarily realized as a separate artifact. For instance, the DSL's core language model is often not explicitly specified, but only existing in the developers' minds and/or buried in the code. In our approach, we treat the central DSL artifacts as explicitly specified artifacts in the development process.

domain of access control. The DSL's core language model thus formalizes domain-specific knowledge and must be validated by domain experts. The core language model can be defined using any suitable modeling language, such as the UML [49] and UML extensions. Depending on the type of DSL, the core language model can either be a meta-model or an ordinary model.

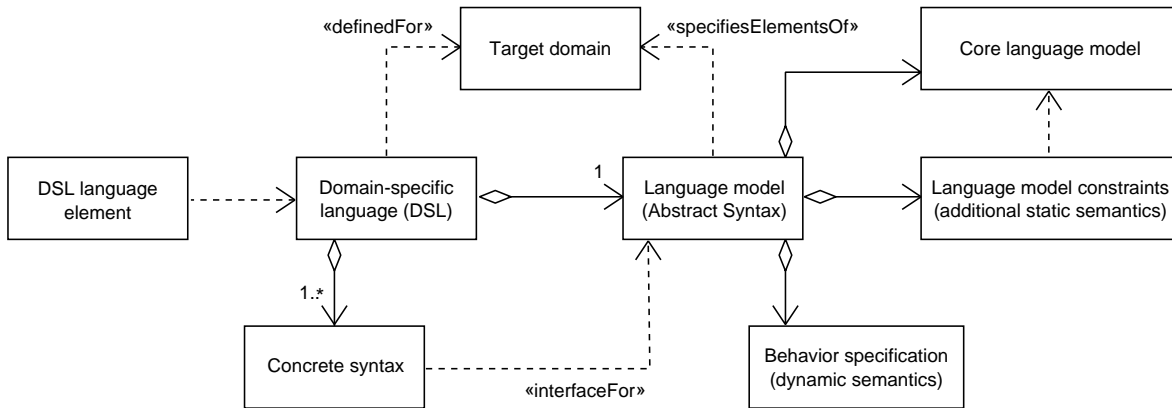


Figure 1. Domain-specific languages: Artifacts

The *language model constraints*, also referred to as (additional) *static semantics*, are essentially a part of the DSL language model (see Figure 1). They express invariants on elements of the core language model and/or on relations between those elements and thereby define semantics that cannot be expressed directly in the (graphical) core language model. An example from the access control domain would be an invariant which defines that two mutual exclusive roles must never be assigned to the same user (see also Section 7). Usually, constraints are provided in a formal constraint language. If the core language model is defined using the UML, for example, such constraints can be specified using the Object Constraint Language (OCL) [48]. Please note that some language model constraints can also be given as informal text, for instance if it is impossible to express a certain constraint with a particular formal constraint language. The downside of this option is that we then cannot use software tools for an automatic handling and checking of such informal constraints.

The *DSL behavior specification*, sometimes also referred to as *dynamic semantics*, is a part of the language model and defines the (behavioral) effects that result from using a DSL language element. Moreover, it defines how the DSL language elements can interact at runtime. The behavior can be specified in many different ways, ranging from high-level control flow models, over detailed behavioral models, to a precise textual specification.

In addition to the artifacts specifying the abstract syntax of the DSL, each DSL needs a concrete syntax to use the DSL in a certain system environment. A *concrete syntax* represents the abstractions defined through the DSL's abstract syntax, and each DSL can have multiple concrete syntaxes, e.g. a graphical syntax and a textual syntax, or a syntax for human users and a syntax for interacting software components.

The concrete syntax serves as the DSL's interface. Thus, the definition of the concrete syntax(es) is especially important from the DSL user perspective. In case a DSL is primarily built for non-programmer human users,

usability properties of this interface, such as being simple and intuitive, are of central importance. In an ideal case, the concrete syntax is therefore both, convenient for human users and easy to process by software components. If a DSL is primarily built for technical experts and/or machine interaction, however, a more complex concrete syntax may be acceptable or useful.

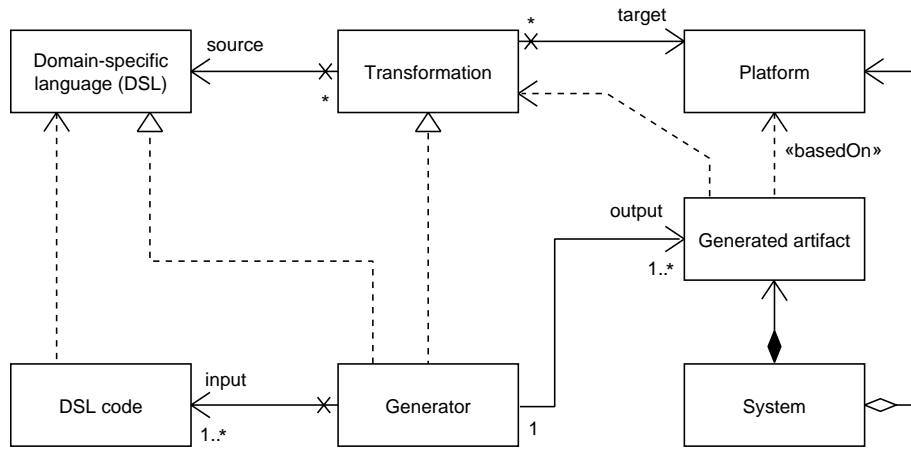


Figure 2. DSL transformations: Artifacts (loosely based on [61])

Transformations are defined to transform DSL code written in a concrete syntax to another model representation or to (programming language) code that can be executed on a specific platform (see Figure 2). In general, a *transformation* is a directive that defines how one (model) format is to be transformed to another (model) format (see, e.g., [43, 59]).

A *platform* consists of software building blocks that provide functions to implement the DSL’s semantics in a specific system environment. The platform consists of generic platform artifacts, such as programming languages and frameworks (for example based on the Enterprise JavaBeans (EJB) technology or Microsoft .NET), as well as DSL-specific platform artifacts (i.e. parts of the software platform that need to be implemented or integrated into a generic platform just to support the DSL).

DSL-to-platform transformations are conducted by a *generator* component (see Figure 2) and result in *generated artifacts* (such as Java, C#, or Ruby code for example) based on the respective platform (see, e.g., [12, 33, 61]). Note that often repeated transformations to intermediate formats are conducted, before machine code is generated that can actually be executed on a certain processor hardware. However, after transforming DSL code written in a concrete syntax to platform-specific artifacts, the remaining transformations are typically executed by standard tools of the target platform, such as compilers or interpreters (examples are a C compiler or a Tcl interpreter).

Two different styles of DSLs can be distinguished with regard to the implementation approach of the DSL (see, e.g., [18, 60]):

- An *embedded DSL* (also called *internal DSL*) is defined as an extension to an existing GPL and uses the syntactic elements of the underlying (host) language. In [44], the approach of building an embedded DSL

is called “language exploitation”. An embedded DSL can directly access all features of the host language including libraries, frameworks, or other platform-specific components. Moreover, the tools (e.g., editor, debugger, compiler, or interpreter) available for the host language (or platform) can directly be used when working with the DSL. Thus, it is often not necessary to define transformations or to build an extra generator for an embedded DSL, even though it is possible, of course. In particular, the compiler or interpreter of the host language serves as DSL generator or it is an essential part of the DSL generator. Typical examples of embedded DSLs are the DSLs provided in the Ruby on Rails framework, the Ruby DSLs described in [19], and the Haskell example presented in [28].

- An *external DSL* is defined in a different format than the intended target language(s) and can use all kinds of syntactical elements (independent of any other language). In [44], the approach of building an external DSL is called “language invention”. External DSLs come with the advantage that DSL designers may define any possible syntax, be it textual or graphical, without considering the syntactical particularities of a given host language. Thus, an external DSL is not bound to a certain host language or platform but can be mapped to different target platforms via transformations. In an external DSL, only the language elements exposed by the concrete syntax are available to the DSL user. Hence it is impossible to use a feature accidentally that is not part of the DSL (as it could happen in embedded DSLs). Typical examples of textual external DSLs are DSLs developed using openArchitectureWare’s XText framework [50] or Microsofts’ OSLO framework [45], and many examples of graphical DSLs such as the domain-specific modeling languages presented in [33].

4 Approach Overview

In the previous section, we have discussed different facets of DSLs and essential DSL concepts. While these facets are more or less well understood, a major issue remains: What are the steps necessary to systematically create a DSL? We conducted many projects where we built different types of DSLs (see Section 8). From these experiences, we can say that the development process of our different DSLs usually consisted of more or less the same realization activities. Hence, our approach is to describe these activities and to identify their interdependencies. In particular, we identified the four main activities summarized in Table 1 (see also Figure 3). Details on these main activities and the corresponding sub-activities are described in Section 5.

In our experience, in most cases DSL development is an explorative, iterative process. Unfortunately, for different DSLs developed in different application contexts, the order in which these activities need to be performed and the exact steps that must be executed to perform the activities can change significantly (see Section 6). Influencing factors are, for example, the size or impact of the DSL, the size of the corresponding development project, the involvement of the various DSL stakeholders (especially the domain experts), and the type of development process that is used in the DSL project (see Section 8 for examples how these factors have influenced the DSLs

<i>Main Activity</i>	<i>Sub-Activities</i>	<i>Typical Input</i>	<i>Typical Output</i>
Defining the DSL's core language model	Identify domain abstraction, Add domain abstraction to language model, Define language model constraints, Check language model	Depends on the actual process tailoring, see Sections 5 and 6	DSL core language model, DSL language model constraints
Defining the behavior of DSL language elements	Select language model element(s), Define language model element(s) behavior, Check DSL behavior	DSL core language model, DSL language model constraints, DSL concrete syntax	DSL behavior definition
Defining the DSL's concrete syntax(es)	Define symbols for language model elements, Define DSL production/composition rules, Define DSL concrete syntax	DSL language model	DSL concrete syntax
Integrating DSL artifacts with the platform/infrastructure	Map DSL artifacts to platform features, Extend/adapt platform, Define DSL-to-Platform transformations, Test DSL, Check Integrated DSL	DSL language model, DSL concrete syntax	Transformations, DSL test suite, platform extensions

Table 1. Overview: Main activities in our DSL development approach

we have built in our projects). As DSLs are usually developed in the context of larger projects or existing software products, often these influencing factors are predetermined and cannot easily be changed or adapted for the DSL development. However, in any case a systematic process, including all required activities, is useful to guide DSL developers. For these reasons, our general approach to DSL development consists of two major tasks:

1. Tailor the DSL engineering process to the project's context.
2. Apply the tailored DSL engineering process to develop the DSL.

It often makes sense to further refine a tailored DSL engineering process during its application phase in order to improve the process and/or adapt it to the environment in which it is performed (see Section 6). Some prototypical examples of tailored DSL engineering processes that we identified in our projects are sketched below:

- *Language Model Driven*: In this type of DSL engineering process, the language model definition drives the DSL development. That is, first the core language model is defined to reflect all relevant domain abstractions, then the concrete syntax is defined along with the DSL's behavior, and finally the DSL is mapped to the platform/infrastructure on which the DSL runs. Like all other process variants, the language model driven process is working incrementally. This means, incremental refinements, iterations, and feedback loops are applied whenever necessary, for instance to correct mistakes or to refine the results of prior activities. The language model driven process variant can be further tailored by specifying how large the increment defined in an iteration is – ranging from larger language model chunks that are defined in one batch, to an own iteration of the process for every single language model element.
- *Mockup Language*: To raise the participation of domain experts, it is also possible to start DSL development with the concrete syntax design and then distill the abstract syntax and semantics from this concrete syntax.

In particular, the concrete syntax is developed together with domain experts by utilizing the domain experts' domain knowledge. Next, a first prototype implementation of the resulting mockup language is developed. The prototype then serves as a means to perform acceptance testing activities, and to further refine and tailor the concrete syntax to the domain experts' needs.

- *Extract DSL from Existing System:* Sometimes an existing software system should (ex post) be provided with a DSL. That is, the domain abstractions for the DSL can be derived directly from the existing system. In this case, it makes sense to first elicit the language model elements from the abstractions given in that system. If the architecture of the corresponding software system is documented (e.g. using a graphical modeling language such as the UML), this architecture description can be a valuable source for the elicitation of domain abstractions.

In our experience, the Language Model Driven Process is well suited for explanatory purposes – it proceeds in a top-down fashion which is easy to understand and follow in a documentation. Thus, even if the DSL is defined using some other tailored process variant, it makes sense to document or explain the DSL with the language model driven process. In our experience, the language model driven process is also well-suited to be applied for small projects, where the DSL developers are domain experts themselves and which are developed from scratch (for example when developing technical DSLs for software developers).

In Section 5, we will use the language model driven process variant to explain the different activities of DSL development. In addition, we describe interdependencies of these activities in terms of inputs required and outputs provided by the activities.

Our approach mainly aims at the engineering activities for developing a DSL – in particular, we describe the recurring development activities that we identified in our DSL projects. Nevertheless, other development activities must also be performed, such as defining/selecting the software development tools and techniques, or project management activities to involve and coordinate different stakeholders. However, in our experiences such aspects are often determined through the general purpose software development process that is applied in a particular project, such as the Unified Process [31] or agile processes such as Extreme Programming [4]. Therefore, such general purpose development activities are not in focus of our engineering approach. That is, our activities for DSL engineering do not define a complete software development process, but rather a micro-process that can be plugged into these general purpose processes. The (considerable) differences between general purpose software development processes are an important reason why our micro-process must be tailored before it can be applied in a concrete project (see Section 6).

5 Process Activities for DSL Engineering – Illustrated Using a Language Model Driven Process

5.1 Main Activity Flow

Figure 3 shows an example tailoring of the DSL engineering process. In particular, the process in Figure 3 shows the main control and object flow of the “Language Model Driven Process” variant (see Section 4). Note that it is of course possible to move back in the process at any time, for example, if an error or an incompleteness in one of the models is detected. For the sake of simplicity, however, we chose to omit the corresponding reverse links in Figure 3 and focus on the main flows. The purpose of this and the following figures in this section is two-fold:

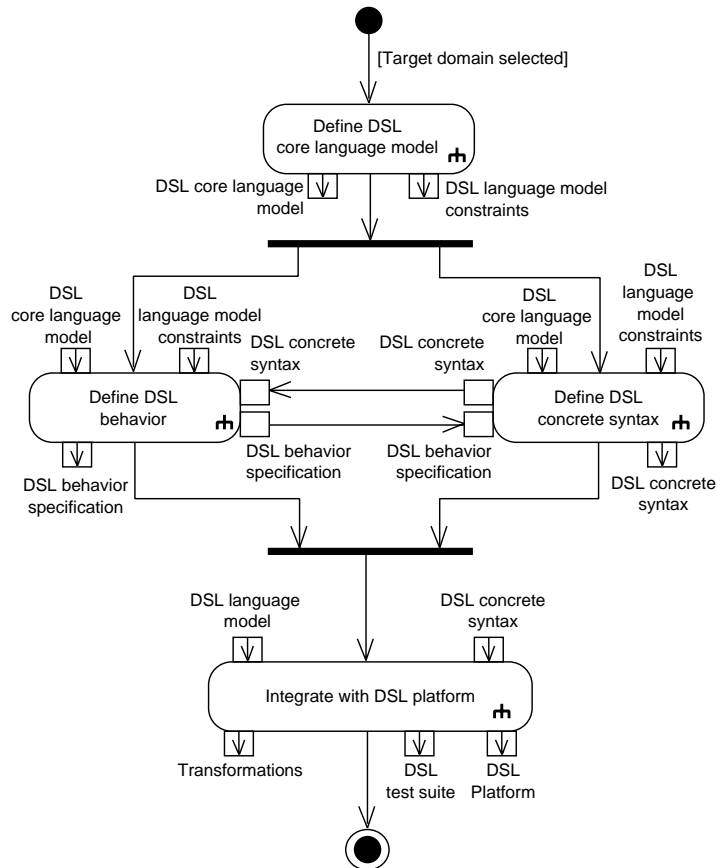


Figure 3. Language Model Driven DSL Engineering Process (main control flow)

- First, the figures specify each action (or sub-activity) in terms of the inputs required and the outputs provided by this particular action (or sub-activity).
- Second, the figures show one possible tailoring of the process as a control and object flow. The purpose is to show to the reader one possible way of how to combine the DSL engineering activities and traverse through a DSL process. In particular, we use the “Language Model Driven Process” for explanatory purposes. Other

tailored processes can be defined by combining the process activities using a different control and object flow (see Sections 4 and 6).

The “Language Model Driven Process” is initiated by selecting the target domain of the DSL. The first activity is to define an (initial) core language model and corresponding language model constraints for this target domain. Subsequently, the process continues with two parallel activities: the definition of the DSL’s behavior and the definition of at least one concrete syntax (see Figure 3). We suggest to conduct these activities in parallel because the behavior and the concrete syntax of a DSL may mutually affect each other. Performing the two activities in parallel is especially useful for embedded DSLs (see Section 3) because syntax and behavior of the host language have a considerable influence on the concrete syntax and behavior of the DSL. When realizing the DSL as an external DSL, these two activities may also be conducted in succession (i.e. the DSL development process can be tailored in this way). However, even if the DSL is build as an external DSL, conducting both activities in parallel is often useful because intermediate results can be exchanged to produce a concrete syntax and corresponding behavior definitions which are consistent, complete, and well integrated with each other.

After defining the DSL’s behavior and concrete syntax, the fourth activity in the engineering process maps the different artifacts to the target platform (see Figure 3). This activity produces transformations, corresponding integration tests, and corresponding platform extensions for the DSL.

Note that an important facet of the engineering process is testing and each of the main activities includes sub-activities to check the respective DSL artifacts on different levels of abstraction. For example, whenever a running piece of code is produced, it should be unit tested. In general, there are many kinds of code artifacts that could be produced during a DSL engineering process, such as a language model element with an API, a concrete syntax element, a grammar, a transformation template, or a method for producing a part of an intermediary language. In addition to unit testing, DSL integration testing checks if different DSL artifacts interact as specified. In Section 5.5, we introduce integration tests as part of the DSL-to-platform integration activity. This is because integration tests can only be run successfully if all parts of the platform that are related to the respective integration test are available. Integration tests are especially derived from the DSL’s behavior specification (see Section 5.3).

In the following sections, we describe each of the four main activities in detail.

5.2 Subprocess: Define DSL Core Language Model

In this section, we explain the refinement of the main activity “Define DSL Core Language Model” (see Figure 3). The subprocess starts with the identification of domain abstractions (see Figure 4). For example, if we are engineering a DSL for the access control domain, corresponding domain abstractions could be: subject, role, or permission; examples for domain abstractions in the banking domain are account, bond, stock, and stock order. For complex domains, the identification of domain abstractions is usually non-trivial. Hence, we generally recommend to follow a domain analysis method, such as domain-driven design for the identification of domain abstractions (see, e.g., [9, 14, 17, 73]).

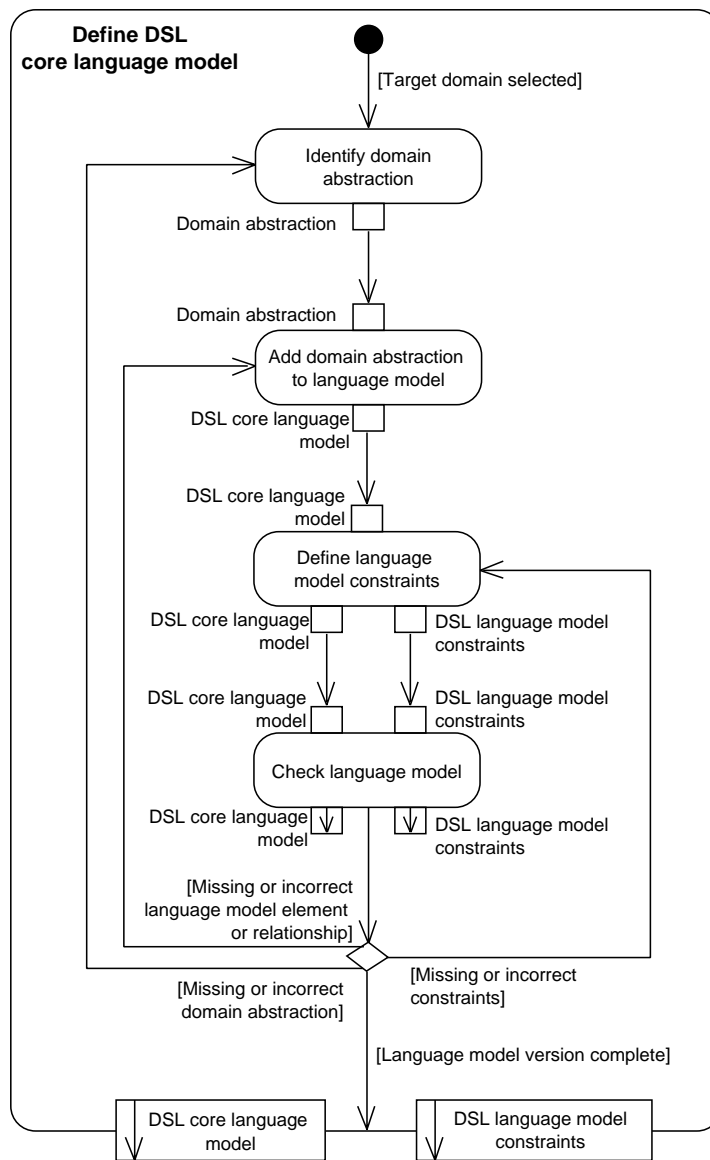


Figure 4. Subprocess: Define DSL Core Language Model

Subsequently, the domain abstractions are integrated with each other to build a version of the DSL core language model. Most often language models are defined using a modeling language, such as UML [49], or a modeling framework, such as EMF [13]. In the most simple case, adding a domain abstraction means to add an additional entity (e.g. adding a class to a class model) for the respective domain abstraction and then connect it to existing entities using relationships, such as associations or generalizations. However, adding domain abstractions can also mean to refactor the core language model. For instance, a new abstraction can also be added by modifying existing entities or relationships.

Next, it is checked whether we need to define additional language model constraints or modify existing ones. Subsequent to defining the core language model and its constraints, the language model is checked for completeness and correctness from a domain-oriented perspective (not in a formal computer science sense). As a thorough

core language model is essential for engineering a DSL, the software engineers must check the core language model in collaboration with domain experts. If these checks reveal an incorrectness or a missing domain abstraction, the software engineers move back to the corresponding activity to adapt the model accordingly. This cycle is repeated until the language model is accepted by the software engineers and the domain experts (see Figure 4). We generally advise to follow the best practices described in [14] to stepwise bring the domain understanding and the technical understanding together.

5.3 Subprocess: Define DSL Behavior

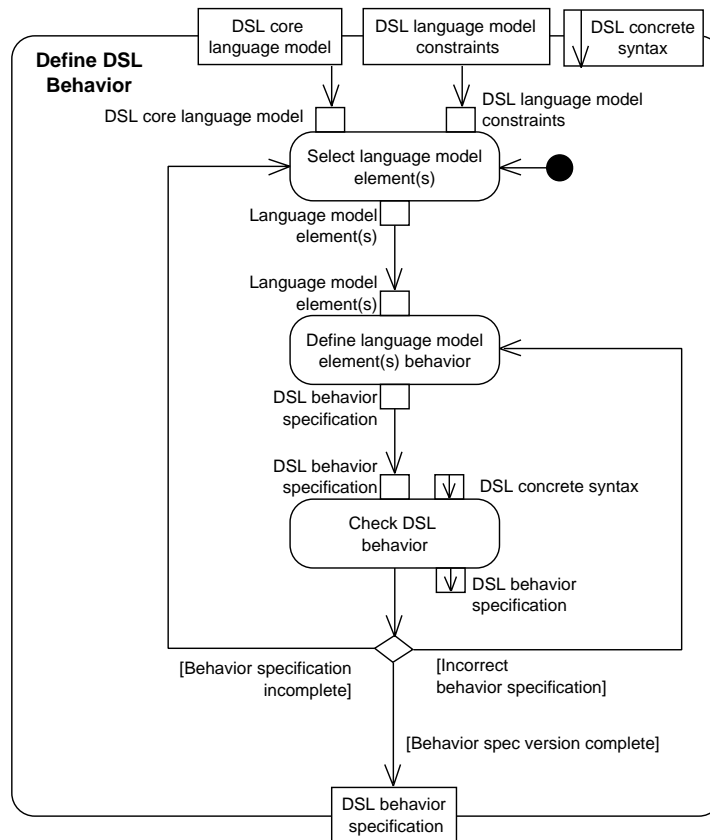


Figure 5. Subprocess: Define DSL Behavior

Even though the language elements of a DSL are defined for a narrow target domain and often appear to be “intuitively clear” to DSL designers and users, intuition may be misleading. In our experience, one of the most common sources of errors is that the software engineers believe they understand the application domain, while this is not yet the case. Therefore, it is most often useful to explicitly define the behavior of a DSL.

The behavior definition of a DSL determines how the language elements of the DSL interact to produce the behavior intended by the DSL designers. This is not only important to the users (domain experts) of the DSL but also to the software engineers. In particular, explicitly specified behavior allows for a correct mapping of the (platform independent) DSL specifications to a certain software platform which must behave exactly as specified.

The sub-process for the definition of DSL behavior (see Figure 5) starts with selecting the core language model elements for which we like to specify its behavior. In our experience, the selection is typically either a single element of the core language model or a number of related core language model elements. Only in rare cases the whole core language model is selected at once. Next, the behavior of the selected elements is defined. In general, the behavior can be specified in many different ways, ranging from high-level control flow models, over detailed behavioral models which are used for model-driven generation of the corresponding code, to a precise (executable) textual specification. Subsequently, the software engineers and domain experts check the behavior specification for domain-correctness and for software-technical correctness. If an error is detected, or if the behavior definition is incomplete, the software engineers move to the respective activity to adapt the definitions accordingly.

5.4 Subprocess: Define DSL Concrete Syntax

The definition of the concrete syntax(es) of a DSL is of central importance for DSL users because it represents the “user interface” of the language. Hence, it should be comprehensible and convenient to use for domain experts – even if actual DSL code is often written by software engineers (see also Section 3).

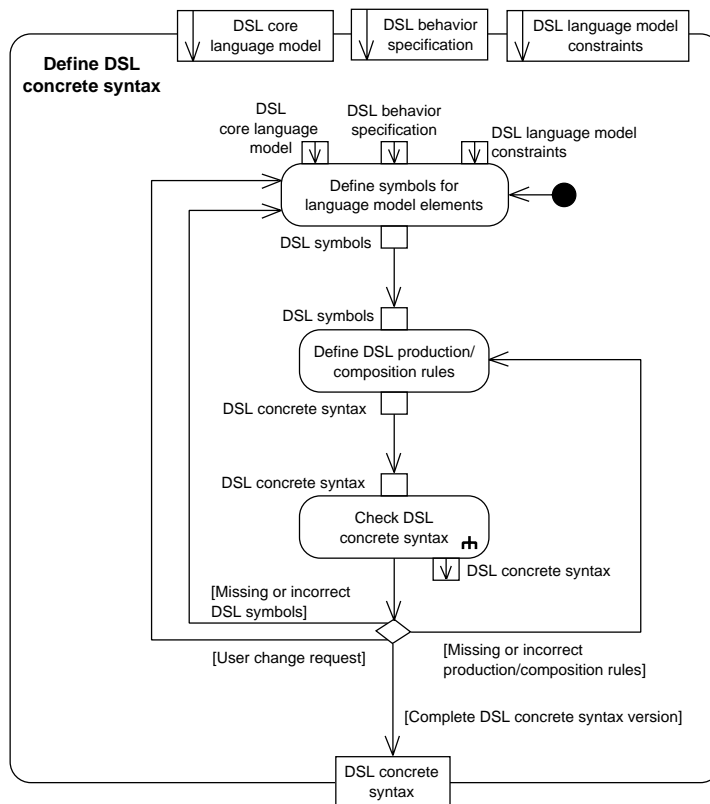


Figure 6. Subprocess: Define DSL Concrete Syntax

As discussed in Section 3, the concrete syntax of a DSL can be either graphical or textual. In addition, we further distinguish embedded DSLs and external DSLs. Moreover, some DSLs have only one concrete syntax while others

have more than one (e.g. a graphical and a textual syntax). Hence, the realization techniques for implementing concrete DSL syntaxes include implementing a GUI editor, implementing a grammar and a parser, or extending an interpreter, for example. Even though these realization alternatives are highly different on a technical level, we can find similar engineering activities that are necessary to define the DSL’s concrete syntax.

The sub-process for defining the concrete DSL syntax (see Figure 6) starts by defining symbols for each element of the DSL. We use the generic term “symbol” to refer to arbitrary concrete representations of DSL elements, for example graphical symbols or textual symbols of a grammar. With regard to these symbols, DSLs should, on the one hand, try to use standard programming language conventions where possible, e.g. to define how comments, strings, and numbers are specified. On the other hand, symbols and terms familiar to the different DSL stakeholders should be used.

Next, the software engineers must define the composition rules or production rules of the syntax. These rules define how the symbols can be composed to phrase legal DSL expressions. In principle, all languages have a grammar and hence production rules. However, we also use the more general term “composition rules” here, because in some cases the grammar is not formally specified in terms of formal production rules, but rather the composition rules are just implicitly given, for instance via the implementation of a graphical tool.

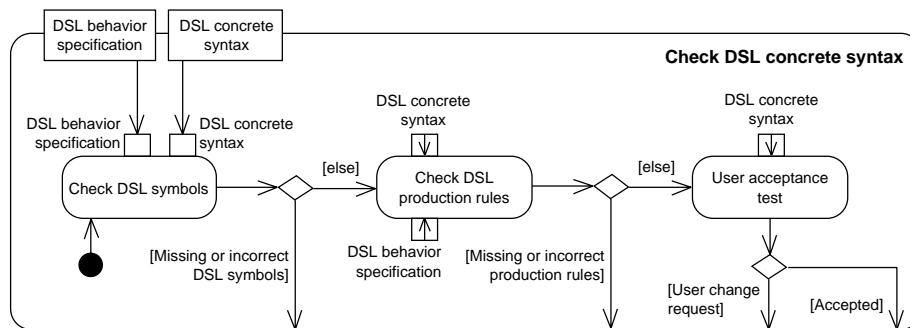


Figure 7. Checking concrete syntax artifacts

The next activity then checks the concrete syntax for correctness and completeness from the point of view of the domain experts (i.e. correctness and completeness are not meant in a “formal” sense here). This is done by interviewing and discussing the domain experts, as well as using observations and studies on their work with the DSL prototype. Moreover, domain experts check if the concrete syntax is convenient to understand and use. Figure 7 shows this activity in detail. If the concrete syntax is found to be incomplete or incorrect, the technical experts move back to the corresponding preceding activity.

5.5 Subprocess: Integrate with DSL Platform

The target platform of the DSL consists of two main parts: generic platform artifacts, such as programming languages and frameworks, and DSL-specific platform artifacts. Examples for generic platform parts could be Enterprise JavaBeans (EJB) or Microsoft’s .NET infrastructure. The DSL artifacts, such as core language model,

language model constraints, behavior definition, and concrete syntax, must be *mapped* to the corresponding platform (see Section 3).

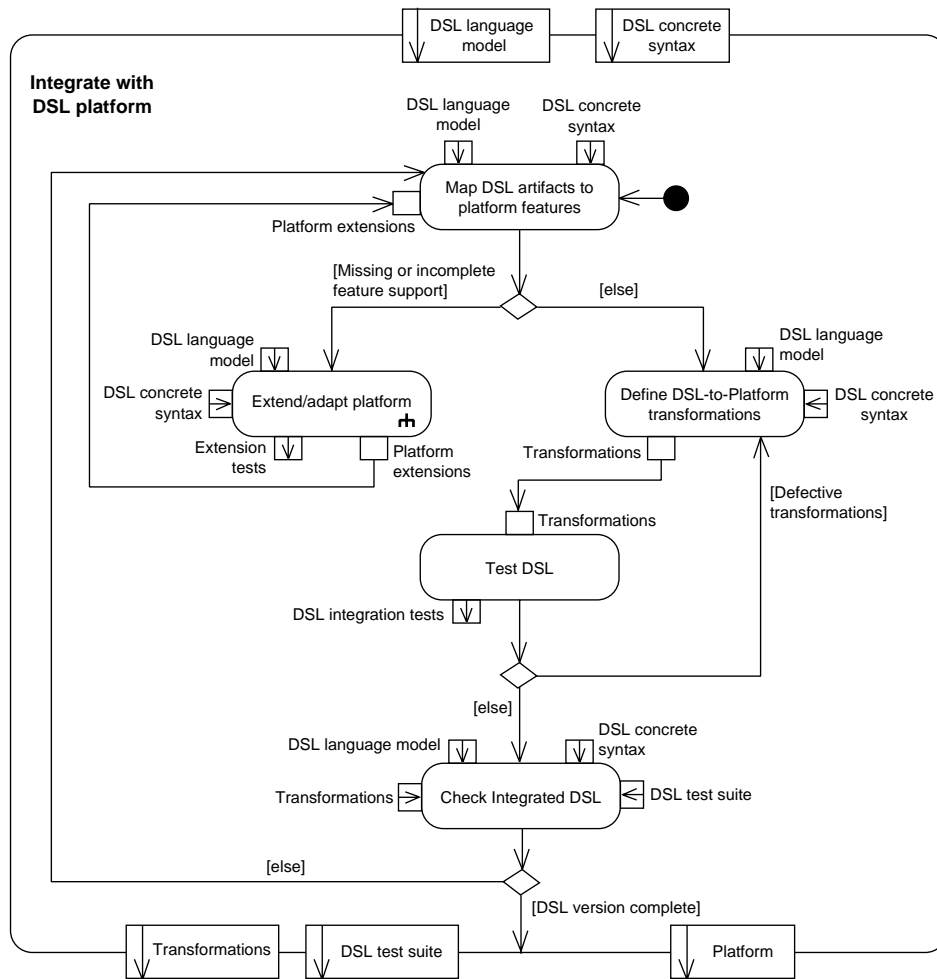


Figure 8. Subprocess: Integrate with DSL Platform

The first activity in the subprocess is to map the DSL artifacts to the features of the platform (see Figure 8). That is, we first have to decide which existing platform parts can be used to realize the respective DSL artifact. Often the platform needs to be extended or adapted because feature support is missing or incomplete. This is done as a next step, if necessary – the corresponding subprocess is shown in Figure 9. In case platform extensions must be defined, we also have to specify test cases for these platform extensions (see Figure 9). When testing of the platform extensions is completed, the platform extensions (and corresponding regression tests) are ready to use.

Now the software engineers check again if the target platform provides all features needed by the DSL. If the platform is still missing features, additional platform extensions are defined. In many cases it is best to define the platform’s interface (and the corresponding extensions) through a clear API. In some cases, a low-level DSL can be used to specify the mapping between the API and the DSL under development.

After all features needed for the platform are identified and implemented, DSL-to-platform transformations are

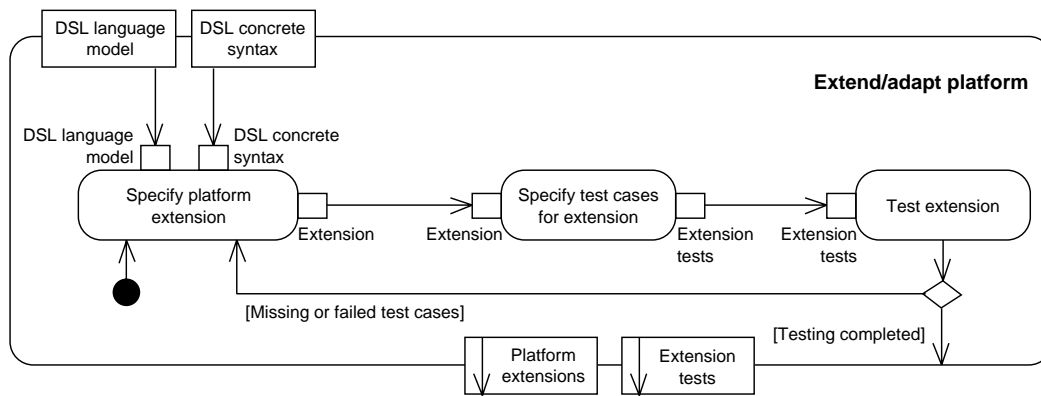


Figure 9. Subprocess: Extend/adapt platform

defined. In case the language is defined as an embedded DSL, it is likely that this activity is very simple (see also Section 3). For external DSLs, the transformations can for instance be realized by defining transformation templates or transformation rules. Of course, both approaches can also be combined (for instance, transformation templates can also be used to generate code for an embedded DSL).

For external DSLs, the software engineers define transformations that convert models, which conform to the DSL’s language model, to the platform (see also Section 3). These transformations should not directly rely on the concrete syntax. That is, if more than one concrete syntax is defined, the same (generic) transformations should work for all concrete syntaxes. For instance, the generator openArchitectureWare [50] requires all concrete syntaxes to be converted to EMF models, and its transformation language XPand allows developers to specify transformation templates that convert any legal EMF model to the target platform.

After defining the transformations, we specify corresponding transformation tests to assure that the generated artifacts resulting from the transformations actually behave as specified. Moreover, we define integration tests for the DSL to check that all DSL artifacts work together properly. The complete test suite of the DSL then includes various unit tests defined throughout the process, the integration tests defined in this activity, and (possibly) tests for platform extensions. In addition, the test suite may include the criteria for user acceptance tests of the concrete syntax (see Section 5.4).

Eventually, all DSL artifacts are subject to a final check conducted by technical experts and domain experts. This final check includes the repeated verification and validation of all DSL artifacts. If the language is found to be complete, the engineering process is finished and a version of the DSL is ready for use (see Figures 3 and 8).

6 Tailoring the DSL Development Process

The micro-process for DSL development is generic in the sense that it can be used in the context of different types of general purpose development approaches, such as the Unified Process or agile development approaches. However, to apply the DSL engineering process in a particular project, it has to be tailored to fit into the corre-

sponding organization’s standard development approach or the specific approach of its project, and to meet the needs of the engineers conducting the DSL engineering process.

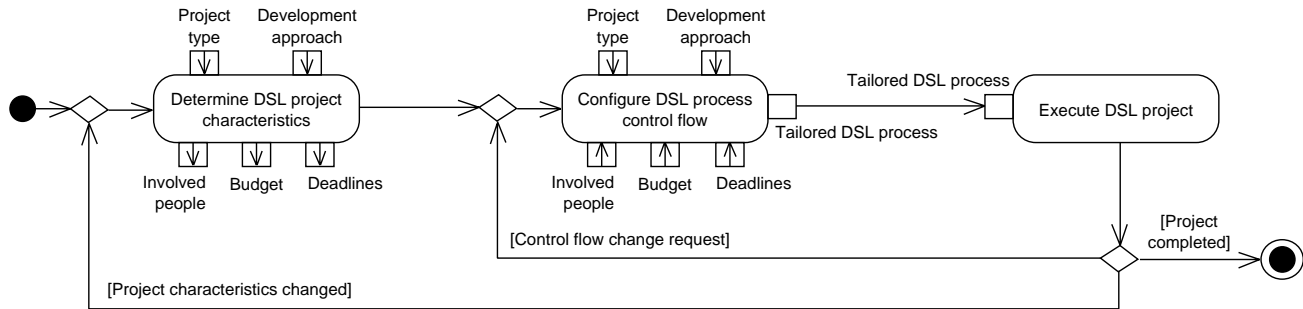


Figure 10. Project-specific tailoring of the DSL engineering process

Figure 10 shows the activities that need to be conducted for a project-specific tailoring of the DSL engineering process. In general, we first determine the characteristics of a particular DSL project. Next, we define a tailored control flow for this particular project, before we can execute the tailored engineering process. It is also possible to dynamically adapt a tailored DSL engineering process, for example if the characteristics of a project change. The characteristics of a particular DSL project are determined by the project type, the development approach, the people (or roles) involved in the project, the (monetary) budget, and the deadlines that must be met. In the second tailoring activity, these factors then define the surrounding conditions of a particular DSL engineering project and influence the control flow of the corresponding tailored process.

Figure 11 shows two examples of tailored DSL engineering processes. In particular, Figure 11 a) shows a process variant that starts with the engineering of a mockup language, and Figure 11 b) shows a process for the extraction of a DSL from an existing system (see also Section 4). Similar to the process variant shown in Figure 3, the tailored processes in Figure 11 show the main control and object flow of the corresponding variants. In addition to the links depicted in Figures 11 a) and b) it is of course possible to move back in the process at any time. However, for the sake of simplicity we focus on the main flows.

In the following two sections, we provide more details on the two process tailoring examples.

6.1 Mockup Language Driven DSL Development

In case a domain is very complex and cannot be easily understood by the software engineers (e.g., air traffic control, definition of financial products, or power-plant control) DSL engineering requires the domain experts to be highly available for the development team. In such a project, it is sensible to raise the participation of the domain experts through the initial definition of a *mockup language*.

In particular, we start DSL development with the concrete syntax design and then distill the abstract syntax and semantics from this concrete syntax, as shown in Figure 11 a). This means, the concrete syntax is developed together with a domain expert by utilizing the domain expert’s domain knowledge. Next, a first prototype im-

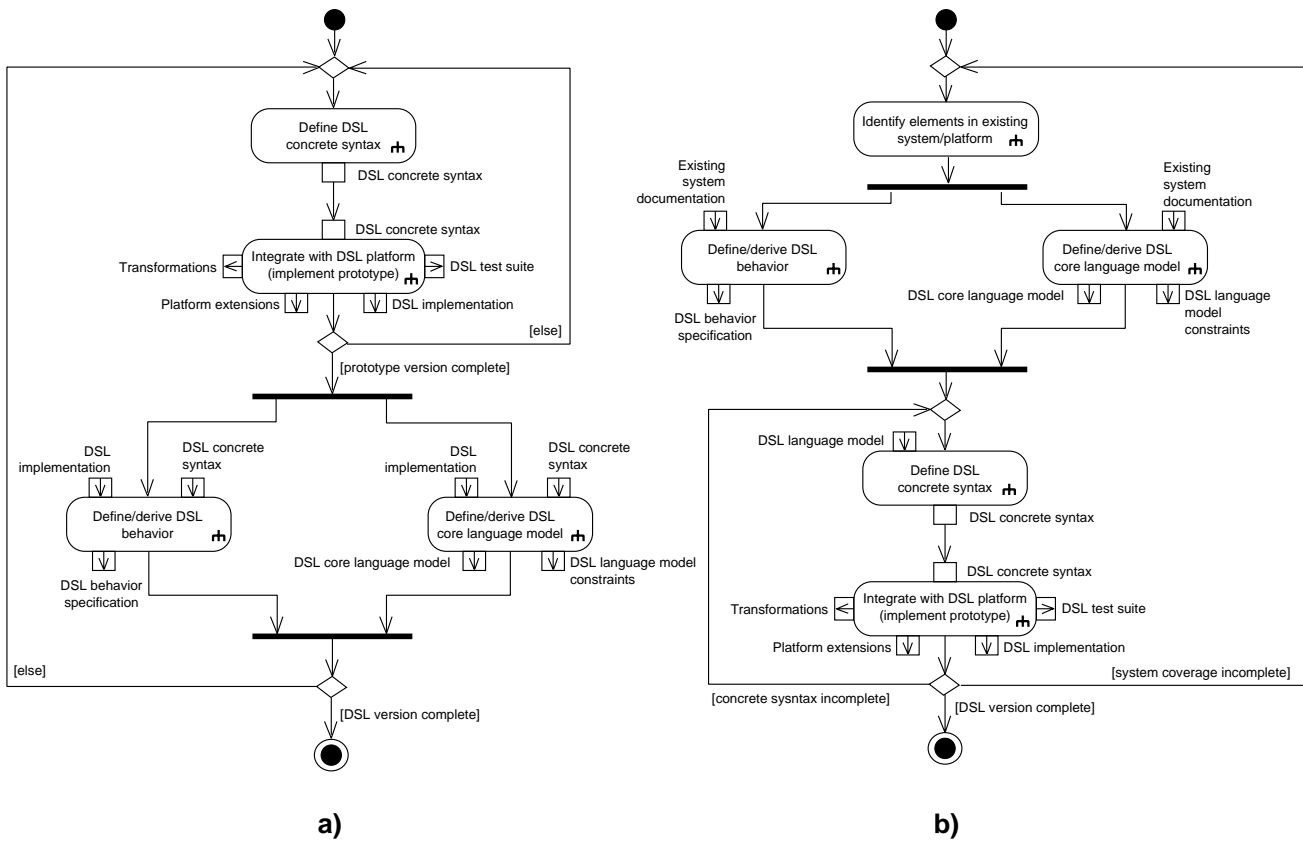


Figure 11. Examples of tailored process variants for DSL engineering

plementation of the resulting mockup language is developed. The prototype then serves as a means to perform acceptance testing activities, to further refine and tailor the concrete syntax to the domain expert's needs. This is done in multiple iterations to get closer to the domain abstractions in a stepwise manner.

Please note that this approach is not driven by minor syntactical considerations, such as which types of parentheses to use. Instead, the focus of the mockup language design is to identify the domain abstractions and terminology. Hence, it makes sense to combine the mockup language driven process with a domain engineering approach, such as domain-driven design [14]. In domain-driven design, the abstraction level targeted by the mockup language driven process is called the ubiquitous language of the domain.

After a version of the prototype is completed, the corresponding DSL core language model and DSL behavior definition that result from this version of the prototype should be documented. Subsequently, the process continues with the next iteration to further enhance the draft version of the DSL. This co-evolution of the different DSL artifacts continues until the domain experts and software engineers define the DSL as complete (see Figure 11 a).

This process tailoring puts emphasis on checking the language model together with the domain experts. In particular, the subprocess for definition of the concrete syntax is guiding the whole DSL engineering process. This subprocess can either be performed iteratively or a syntax can be completely defined for a version of the language model. For instance, in case a second concrete syntax is added to a DSL, it often makes sense to specify the whole syntax at once. In our experience, a more iterative approach makes sense in case the language is used for incrementally illustrating and sketching domain concepts while the DSL is developed. This, of course, is dependent on the availability of the domain experts. Mockup language driven DSL development is well-suited for projects that follow an agile development approach [25, 26], which advocate on-site domain experts (customers).

6.2 Extracting a DSL from an Existing System

Sometimes the goal of an DSL development project is to define a DSL as an (additional) user interface for an existing software system. Therefore, a tailored engineering process for such a project adds an "Identify elements in existing system/platform" analysis activity which identifies one or more elements of the system/platform (see Figure 11 b). Examples for such elements are system components or interfaces that can be mapped to language model elements. That is, the domain abstractions for the DSL can directly be derived from the existing system. In seldom cases (e.g. for small systems or very well understood domains), it makes sense to map the whole existing system/platform to the DSL at once. However, in most real world cases this type of engineering process is also iterative (see Figure 11 b).

If the architecture of the corresponding software system is documented (e.g. using a graphical modeling language such as the UML), this architecture description can be a valuable source for the elicitation of domain abstractions and DSL behavior. Here, domain experts either take part in the elicitation process, or they review the resulting language model. The definition of the DSL's behavior and core language model can be further tailored, e.g. by how many DSL elements are specified in an iteration of the process, or by defining how often the domain

expert gets involved in the sub-process. This depends on the availability of the domain expert and how iterative the development process is. In less iterative processes, we usually specify the behavior of multiple elements at once. After defining the DSL core language model and behavior, the process proceeds with the definition of a concrete syntax and the integration with the platform. Integration with the platform here primarily means the definition of transformations from models (that conform to the language model) to code that is executable on the target platform.

7 An Illustrative Example: Engineering a DSL for RBAC

In this section, we describe an example application of our DSL engineering process for the development of a role-based access control (RBAC) DSL. In particular, we applied the language model driven variant of our process (see Section 5) to develop this RBAC DSL. Below we present a selected, yet representative, subset of the RBAC DSL specification to demonstrate the engineering tasks described in Section 5. A complete definition of the RBAC DSL, including all formal models and specification details, would be beyond the scope of this article.

7.1 Defining the DSL's Core Language Model

We have chosen to use the language model driven variant of our process because the RBAC DSL is placed in a domain well known to the developers (see, e.g., [46, 47, 62, 63, 64, 65]). Therefore, it was sensible for us to start with a conceptual language model which can be incrementally refined as the DSL evolves.

Figure 12 depicts our core language model of the role-based access control domain as an UML2 class model. In the RBAC context (see [16, 56]), an access control *subject* is an active entity (e.g. a human user or a software agent) that is (or should be) able to access objects (e.g. a file like an XML document, or a hardware resource such as a printer) in a particular information system. Each subject receives its permissions through the roles assigned to this particular subject. A *role* is an abstract entity which represents a certain type of subject in terms of the permissions granted to this type of subject. When modeled to represent human users, roles most often reflect the work profiles of a certain organization (see [47, 63]). In its basic form, a *permission* consists of an ⟨operation, object⟩ tuple. It represents the right to perform the respective operation on the corresponding object. Moreover, permissions can be associated with *context constraints*. Context constraints define predicates that must evaluate to true in order to grant a certain access request. They allow for the consideration of context information in access decisions and enable the definition of additional conditions on permissions, such as time constraints (for details see [64]).

An access request is defined through a ternary association and consists of a ⟨subject, operation, object⟩ triple (see Figure 12). Access requests are granted if the respective subject owns a corresponding permission.

In RBAC, the relations between subjects and roles, as well as between roles and permissions are called *assignment* relations (see [16, 56]). As both assignment relations are many-to-many relations, each role can be assigned to different subjects and each permission can be assigned to different roles.

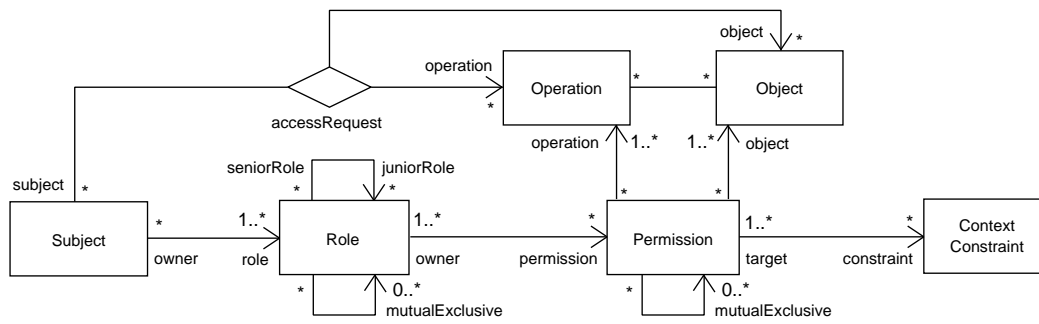


Figure 12. Core language model of the target domain: role-based access control

The core language model includes two recursive associations on Role. The first one defines that a role may have multiple junior and/or senior-roles. In RBAC, these junior and senior-relations are used to build role hierarchies. In a *role hierarchy*, senior-roles are more powerful than junior-roles which means that a senior-role inherits the permissions assigned to its junior-roles.

The second recursive association on Role defines that each role may be mutual exclusive to one or more other roles (see Figure 12). Mutual exclusion is used to enforce the “separation of duty” concept (see, e.g., [15, 16, 62]) which directly affects role-to-subject and/or permission-to-role assignment. In access control, separation of duty constraints enforce conflict of interest policies. Conflict of interest arises as a result of the simultaneous assignment of two mutual exclusive roles (or permissions) to the same subject. Mutual exclusive roles (or permissions) result from the division of powerful rights or responsibilities to prevent fraud and abuse. An example is the common practice to separate the “controller” role and the “chief buyer” role in medium-sized and large companies, or the separation of the “trader” and “internal revision” roles in investment banking.

In addition to the language model elements defined in Figure 12, there are a number of domain-specific constraints on core language model elements. Here, we use the Object Constraint Language (OCL) [48] for defining these constraints. Of course, one may also use other modeling languages than UML and OCL to define a DSL language model and corresponding constraints (see also Section 3).

The following OCL constraint defines that a subject can never own two mutual exclusive roles:

```
context Subject
inv: self.role->forall(r1,r2 |
    r1.mutualExclusive->select(mer|mer.name=r2.name)->isEmpty()
    and
    r2.mutualExclusive->select(mer|mer.name=r1.name)->isEmpty() )
```

Furthermore, similar to permissions, separation of duty constraints are subject to inheritance (see, e.g., [62]). Thus, we define the OCL constraints below to express that no role can be mutual exclusive to itself nor to one of its junior or senior roles:

```

context Role
inv: self.mutualExclusive->select(mer|mer.name=self.name)->isEmpty()
inv: self.juniorRole->forall(r|
    self.mutualExclusive->select(mer|mer.name=r.name)->isEmpty() )
inv: self.seniorRole->forall(r|
    self.mutualExclusive->select(mer|mer.name=r.name)->isEmpty() )

```

Analogous to roles, mutual exclusion can be applied to individual permissions (see Figure 12). This means, that two mutual exclusive permissions may never be assigned to the same role and that no junior-role is allowed to possess a permission which is mutual exclusive to the permissions assigned to one of its senior-roles. This is defined by the following OCL constraints:

```

context Role
inv: self.permission->forall(p1,p2|
    p1.mutualExclusive->select(mep|mep.name=p2.name)->isEmpty()
    and
    p2.mutualExclusive->select(mep|mep.name=p1.name)->isEmpty() )
inv: self.juniorRole->forall(r|
    self.permission->forall(p1|
        r.permission->forall(p2|
            p1.mutualExclusive->select(mep|mep.name=p2.name)->isEmpty() )))

```

Finally, for the sake of completeness, we define that no permission can be mutual exclusive to itself:

```

context Permission
inv: self.mutualExclusive->select(mep|mep.name=self.name)->isEmpty()

```

The elements of the core language model (i.e. the classes and associations) determine the different language elements of our DSL. In addition, the core language model already determines and/or constrains the behavior of DSL elements to a certain degree. For example, the OCL constraints defined above specify invariants that must never be violated by an implementation (respectively the target platform) of the DSL. However, in addition to such invariants, a number of behavioral choices exist which are not defined through the DSL's core language model. Therefore, we define the DSL's behavior as a next step.

7.2 Defining DSL Behavior

Because of our extensive experiences in the RBAC domain we were able to completely define the behavior for each feature of the DSL before it got implemented (respectively mapped to the target platform). When, during implementation or other process stages, we observed issues with the DSL behavior specifications, we incrementally refined them.

For the RBAC DSL, we especially used extended UML activity diagrams (see [66]) to model DSL behavior and the interaction of different concerns in the DSL. In a subsequent step, these activity models are further refined through UML interaction diagrams. The activity models then specify the high-level control flows that are independent of the implementation on concrete platforms, while the interaction models describe the platform-specific details (see also Section 7.4).

Figure 13 depicts an activity diagram that shows the primary control flow for authorization decisions. Furthermore, it shows the interdependencies of the Authorization concern, the Role concern, the Permission concern, and the ContextConstraint concern.

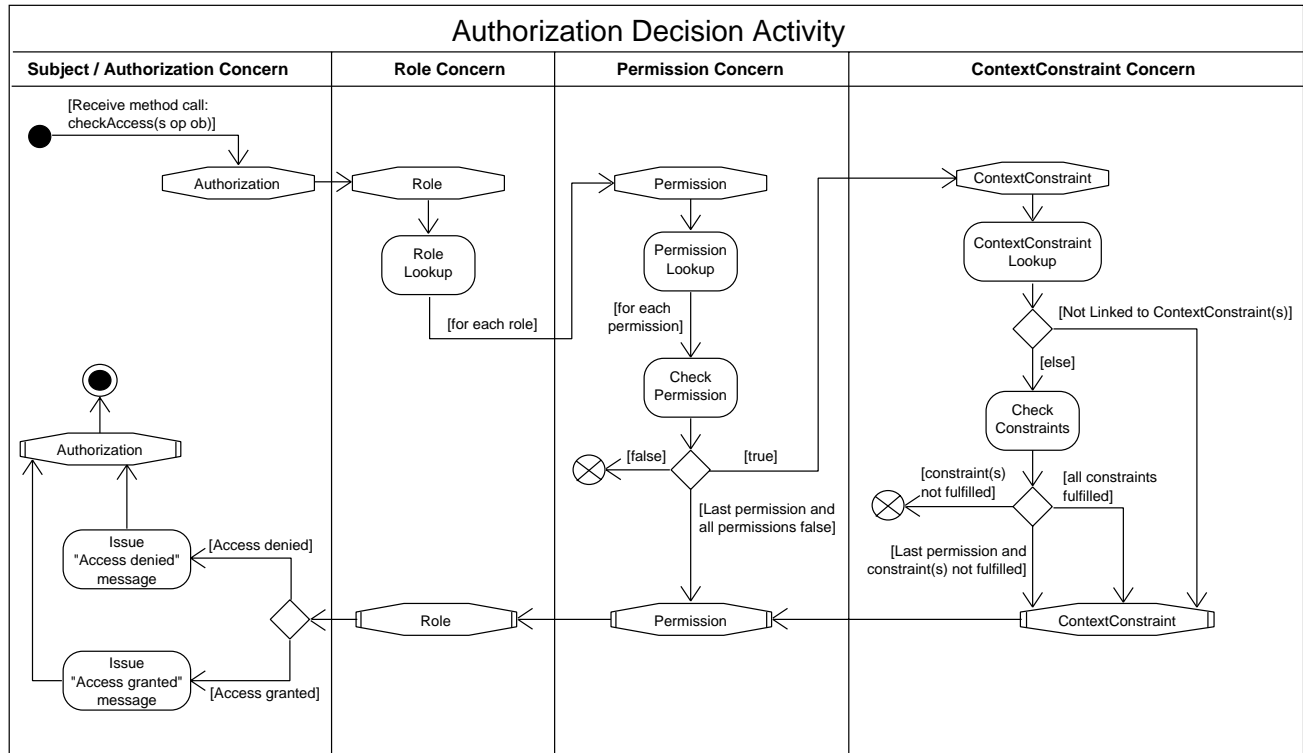


Figure 13. Control flow of the authorization concern [66]

The `checkAccess` language element is derived from the “accessRequest” association shown in Figure 12. The `checkAccess` element probably provides the most important functionality of the RBAC DSL and is applied to check if a certain access can be granted or must be denied, i.e. if a certain subject `s` is allowed to perform operation `op` on object `ob` (see Figures 12 and 13). After receiving a `checkAccess` message, the Authorization concern is responsible to make an access decision for the corresponding access request. In order to reach this decision, it has to interact with the Role concern. The Role concern then performs a role lookup procedure to determine the roles assigned to the respective subject. Subsequently, the Permission concern takes over to check the permissions which are assigned to the corresponding role objects.

However, as mentioned above, to grant a certain access request it is not sufficient to own the appropriate per-

mission – all context constraints associated with the corresponding permission must be fulfilled at the same time. Thus, if a certain permission actually grants the access request (indicated by returning “true”) the ContextConstraint concern intercepts the control flow to check the constraints linked to this particular permission object. Depending on the result of the permission and constraint checking procedures an “access granted” or an “access denied” message is issued (see Figure 13).

7.3 Defining a Concrete Syntax

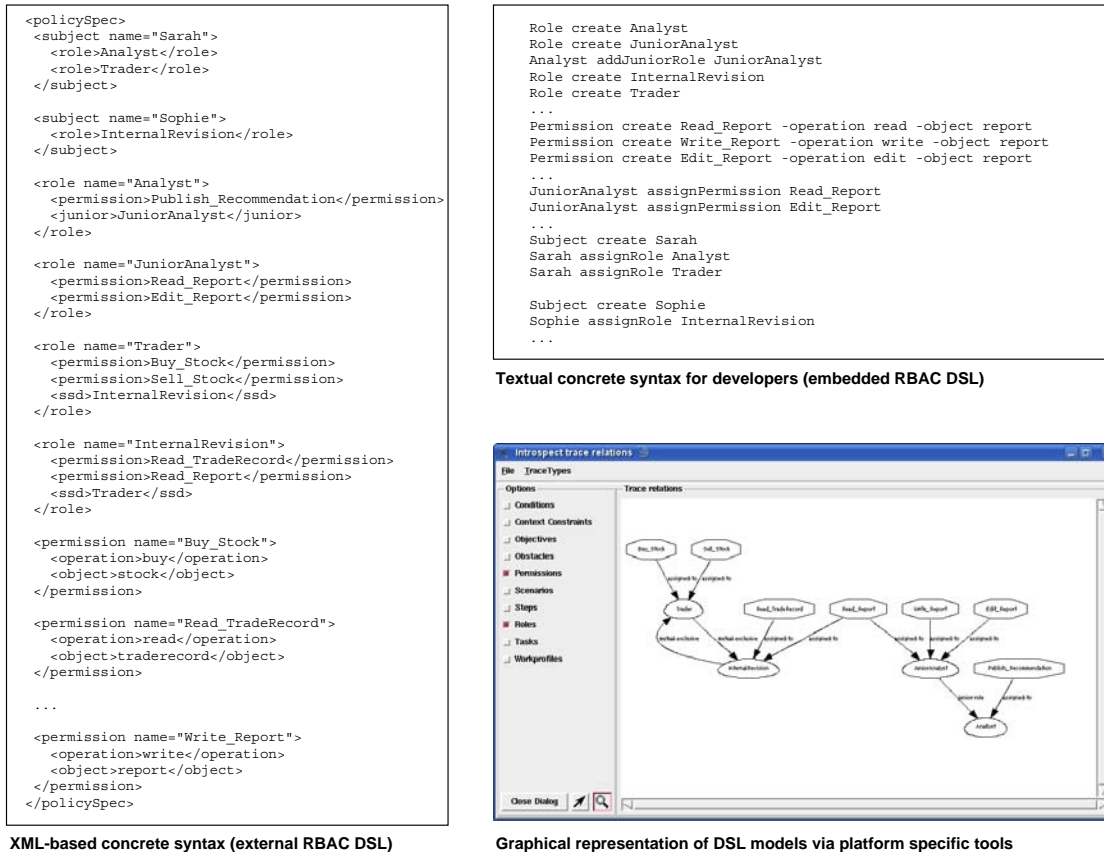


Figure 14. Concrete syntaxes of our RBAC DSL

In the next step, we define a concrete syntax for our RBAC DSL. At this point we like to point out that we conducted two projects to build RBAC DSLs (see also Section 8). In the first project, we defined a declarative XML-based concrete syntax that was primarily intended as a transfer syntax for exchanging RBAC DSL models over network connections. In the second project, we defined an additional textual syntax for an embedded RBAC DSL that is based on the same language model (see Sections 7.1 and 7.2). This embedded DSL is intended to be used by software developers. Moreover, our target platform for the RBAC DSL (the xoRBAC platform) provides GUI tools to create and manage RBAC models (see [46, 63, 64]). This means, we can use the GUI tools of our target platform to manage the resulting RBAC models. In other words, with the xoRBAC target platform we get

a graphical tool “for free” and do not need to define an additional graphical syntax. Figure 14 shows the concrete syntaxes of our RBAC DSL. Below, we describe the development of the XML-based syntax in detail. In particular, we implemented a generator for the RBAC DSL that uses an XML parser to generate a DOM tree [27] which is again transformed to a valid runtime model for our target platform (see Section 7.4)

Figure 15 shows a tree view of the corresponding XML schema (actually it shows an excerpt for the most relevant elements). The schema defines structures to represent the different entities defined via the language model, such as a `subject` element, a `role` element, and so forth. Furthermore, to depict the role-to-subject assignment relation (see Section 7.1) the `subject` element includes a nested `role` element. However, to define a valid RBAC model with our DSL, the `role` element nested in `subject` must refer to a role defined via the top-level `role` element (see Figures 12 and 15). To enforce this constraint we use the `key` and `keyref` elements of XML schema (for details see [69]). Other relations, such as permission-to-role assignment, or mutual exclusion (separation of duties) are realized accordingly (see Figure 15).

In the example from Figure 14, we have defined two subjects, Sarah and Sophie. Sarah owns the roles Analyst and Trader, while Sophie is assigned to the role InternalRevision. The Analyst role has a directly assigned permission Publish_Recommendation and inherits the permissions Read_Report and Edit_Report from its junior-role Junior_Analyst. The Trader role has two directly assigned permissions and is defined as mutual exclusive to the InternalRevision role (note that `ssd` is an abbreviation of “static separation of duty”, see also [62]). Finally, each permission consists of an \langle operation, object \rangle pair - for example, the Read_TradeRecord permission defines the right to perform the operation `read` on objects of type `traderecord`.

7.4 Platform Integration

7.4.1 Mapping to the xoRBAC Platform

As mentioned above, xoRBAC is used as the target platform for the DSL. xoRBAC provides a sophisticated RBAC service including a policy decision point and a policy repository for RBAC policies (for details on the xoRBAC platform see [46, 62, 64, 77]). To enable a thorough definition of DSL-to-platform-transformations, we thus need to map our DSL artifacts (including the behavior definition) to the functions provided by xoRBAC.

Figure 16 shows an excerpt of Figure 13 that includes \ll *concernSpec* \gg stereotypes for the ConcernStart nodes of the Role and the Permission concerns (see also [66]). In particular, these concernSpecs define that the behavior of the corresponding concerns is implemented through the Role and Permission classes of xoRBAC, respectively. Moreover, the `enterOperation` of both concerns is implemented via the `checkAccess` operation of the corresponding classes.

In Figure 17, we see an interaction diagram that models the invocation sequence of the Role concern in xoRBAC. After receiving a `checkAccess` message, the corresponding Role object invokes the `checkAccess` method on the permission objects assigned to this particular role. If one of the permissions grants the access by

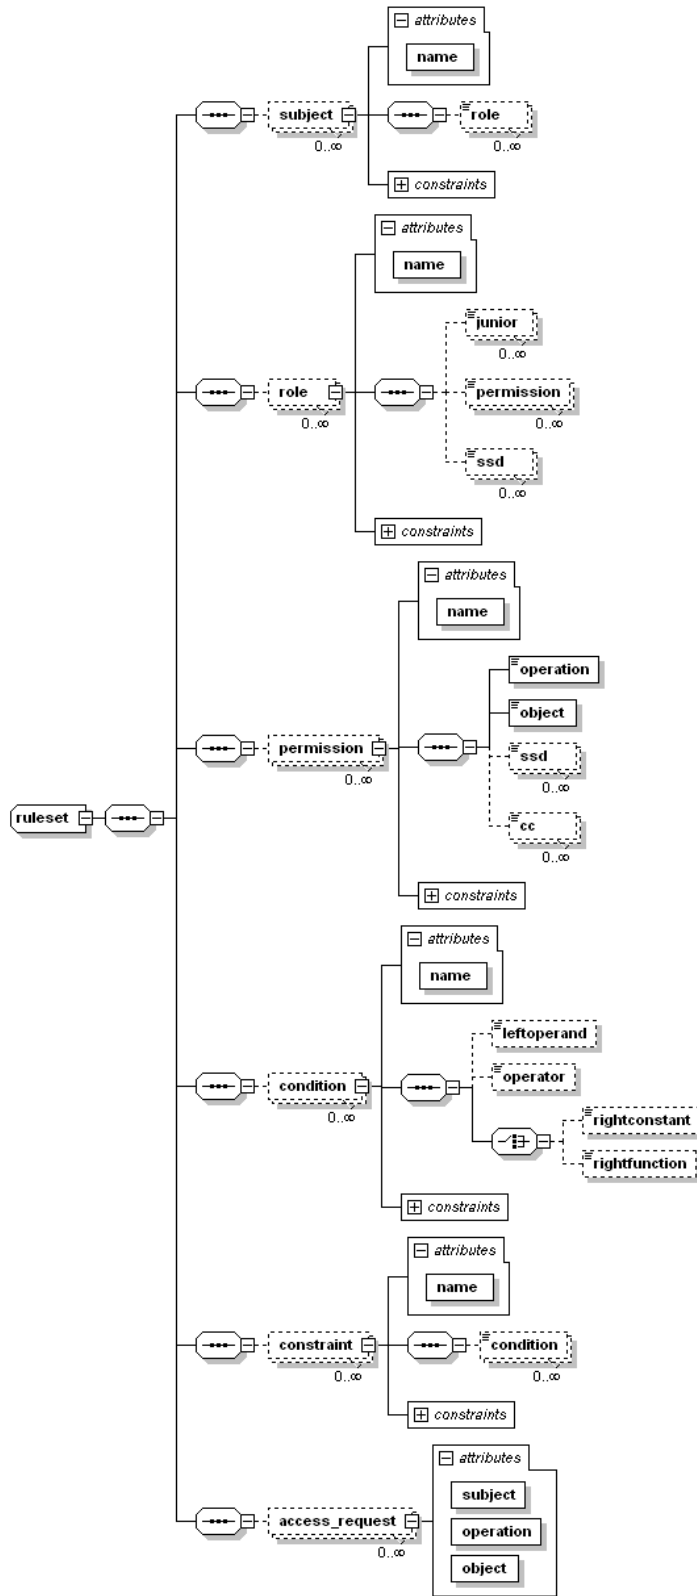


Figure 15. XML schema (excerpt) of a concrete syntax for the RBAC DSL

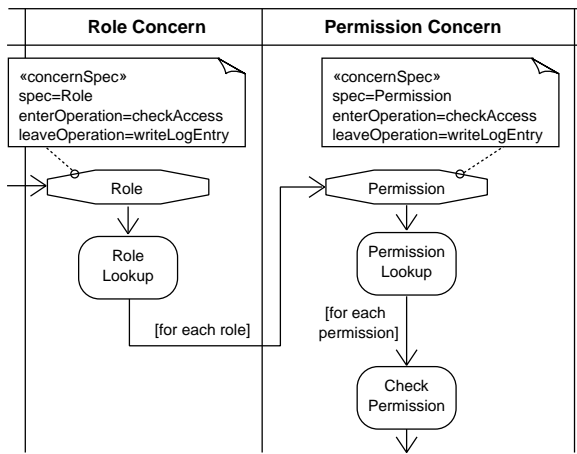


Figure 16. The Role and Permission Concerns with concernSpec stereotypes [66]

returning “true”, the `Role` object immediately stops checking its permissions (see the “break” InteractionOperator in Figure 17), writes a corresponding log entry, and returns the access decision.

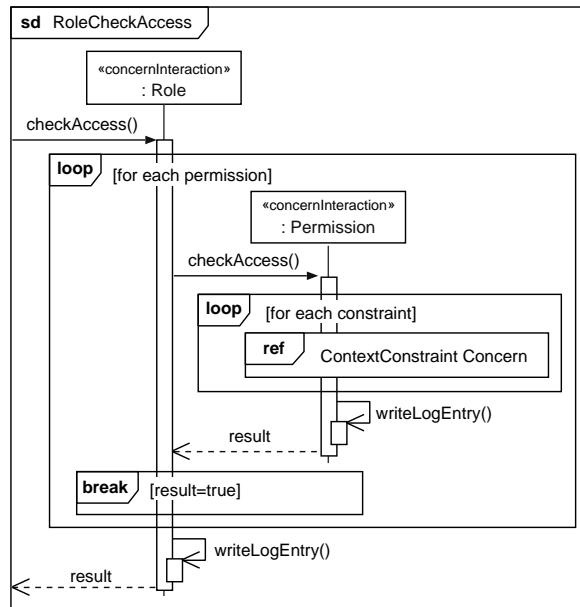


Figure 17. Interaction model for the Role concern [66]

7.4.2 Transformation of DSL Language Elements

After mapping the different DSL artifacts to our target platform, we define DSL-to-platform transformations that actually generate executable artifacts which “produce” the desired behavior on the target platform.

In our RBAC DSL, we used XML schema to define a concrete textual syntax (see Section 7.3). Now we have a number of different choices how to define transformations from XML-based DSL expressions to executable xORBAC models. One option is to define a respective XSLT template. However, while XSLT is well-suited to

define transformations between relatively simple (and static) structures, the templates soon become very complicated for more complex and dynamic target domains. As a consequence such templates are hard to maintain in case the DSL and/or the target platform evolve.

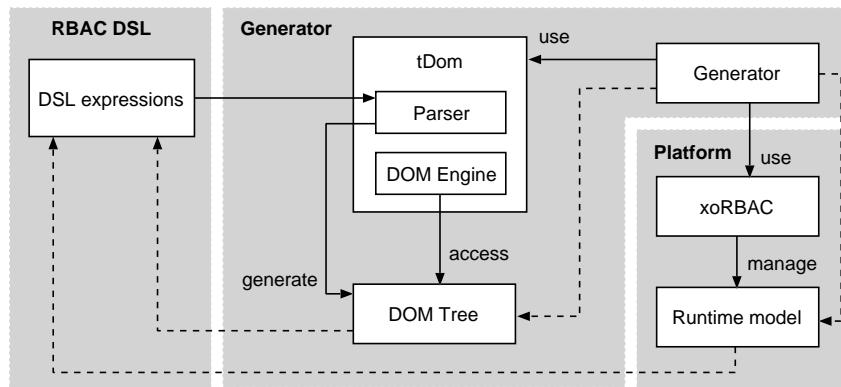


Figure 18. Transformation of DSL expressions: Structural overview

Moreover, for the RBAC DSL it is not sufficient to simply transform a certain XML-based DSL expression to a corresponding executable model element of the xoRBAC platform. Instead, the transformations, or, to be more precise, the generator applying the transformations, must have a certain knowledge of the RBAC domain to produce executable models that conform to the semantics specified for the DSL (see also [64, 65, 66]).

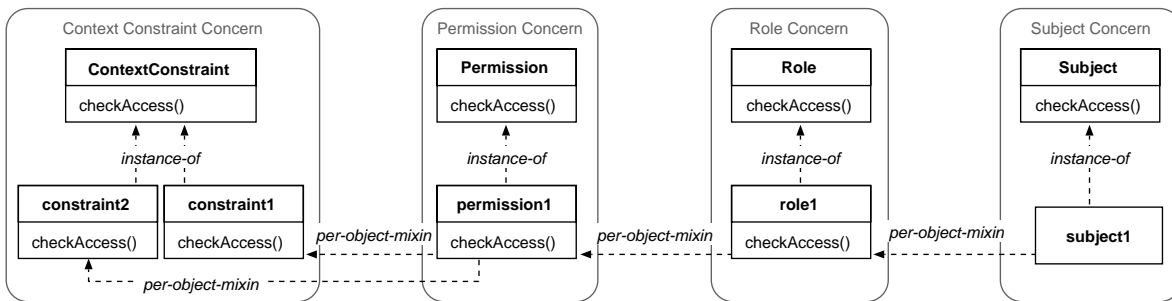


Figure 19. Example of an executable model generated from the RBAC DSL

For the above reasons, we decided to implement a generator component which reads the concrete syntax of our RBAC DSL and produces conforming executable xoRBAC models. Figure 18 shows the conceptual structure of this generator. In particular, the generator uses tDOM [67] to parse XML-based DSL expressions and to access the resulting DOM tree [27]. The DOM tree provides a runtime representation of the XML input and allows to flexibly access the different DSL expressions. The generator accesses the DOM tree via tDOM and generates a conforming executable xoRBAC model (see Figure 18). Figure 19 shows an example of such an xoRBAC runtime model.

Figure 20 shows an activity model that defines the corresponding action sequence in detail: After receiving a document with DSL expressions, the generator dispatches the corresponding XML document to tDOM and

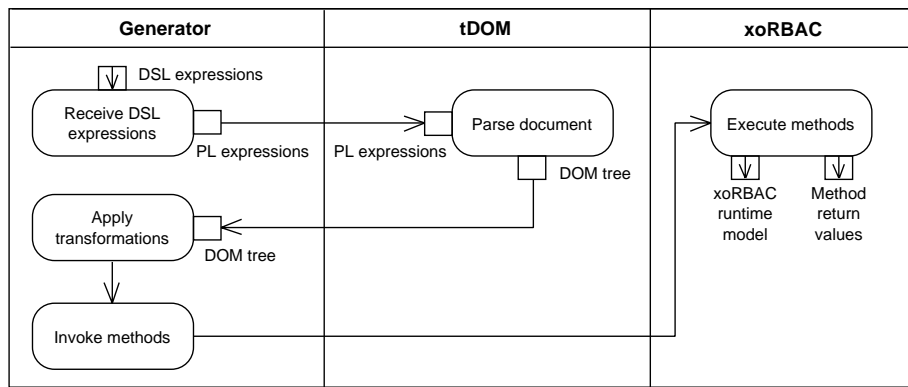


Figure 20. Transformation of DSL expressions: Activity model

invokes the “Parse document” action. The tDOM component parses the document, builds the respective DOM tree, and returns this DOM tree (respectively a reference to the DOM tree) back to the generator. The generator then applies the DSL-to-platform transformations and invokes the corresponding xoRBAC methods. As a result of these method calls xoRBAC produces a conforming executable model and respective method return values which may be processed by the generator or other services.

8 Lessons Learned from the Various Projects

The research presented in this paper is based on our experiences from various projects. In each of the corresponding projects we analyzed and verified many details of our DSL engineering process, and changed it, when mismatches occurred. In the course of the projects, at first, we started to construct an informal DSL engineering process guide and added all DSL development activities which we observed in our data. However, we then realized that this first (intermediate) result did not completely fit our observed data from the projects. Hence, we refined the result by explicitly describing the activities in engineering a DSL, including their respective inputs and outputs. This evolved version was still unsatisfactory, as it neglected the implicit process aspects which we observed in our analysis. For this reason, in the third iteration, we added the tailoring step reported in Section 6 which better reflects the observed DSL projects and other available data. In a similar way, all details of our results were also refined in an iterative manner.

In addition to many other projects, we observed and analyzed 14 DSL projects in-depth, which are summarized in Table 2. To illustrate the main lessons learned in terms of the selection and tailoring of the DSL development process, as well as the different influencing factors, we summarize the main considerations for each of these projects:

- *Bibliography management*: This project was a small scale project with only one developer in a University institute’s Web site and information system administration setting. Test-driven development was used as the software development approach. The developer understood the domain and was a DSL user himself.

<i>Domain</i>	<i>Target Platform/ Technology</i>	<i>DSL Type</i>	<i>Process Variant</i>	<i>Project Type</i>	<i>Our Role(s)</i>	<i>Project Duration</i>
Bibliography management	Frag, HTML, Latex	Embedded	Mockup language driven	Research project	DSL Designer, Developer	3 month
BPMN process modeling	BPEL process engine, WSDL	External	Language model driven	Research project	Architect, Ob- server	1 year
Compliance metadata	Java, BPEL, HTML	External	Language model driven	Research project	Architect, Ob- server	6 month
DSL editor specification	Eclipse, Java, Frag	Embedded	Mockup language driven	Research project	Architect, Ob- server	1 year
Data backup	Tcl/XOTcl	Embedded	Mockup language driven	Case study	DSL Designer, Developer	6 weeks
Document archiving	C-based custom platform	External	Extracting DSL from exist- ing system	Industry project prototype	DSL Designer, Developer, Observer	2 years
Multimedia home platform	Frag/Java MHP platform	Embedded	Mockup language driven	Industry project prototype	DSL Designer, Developer, Observer	1 year
Role-based access control	XML, xoRBAC	External	Language model driven	Research project	Domain Expert, DSL Designer, Developer	2 month
Role-based access control	Tcl/XOTcl	Embedded	Extracting DSL from exist- ing system	Case study	Domain Expert, DSL Designer, Developer	4 weeks
SOA architectural views	BPEL process engine, WSDL, Java	External	Language model driven	Research project	Architect, Ob- server	2 years
Software testing	Tcl/XOTcl	Embedded	Mockup language driven	Research project	DSL Designer, Developer	3 month
Software architectural knowledge	Frag, HTML, Latex	Embedded	Mockup language driven	Research project	DSL Designer, Developer	3 month
Software documentation	Frag, HTML, Latex	Embedded	Mockup language driven	Research project	DSL Designer, Developer	3 month
Workflow definition	Tcl/XOTcl	Embedded	Mockup language driven	Case study	Domain Expert, DSL Designer, Developer	2 month

Table 2. Overview of DSL projects we analyzed in-depth

A domain-model could rather easily be derived from the concrete syntax because the domain was well understood. Other potential users of the DSL where only initially involved – and hence a mockup language driven micro-process was chosen because mainly the syntax needed to be agreed upon. Only one iteration of the mockup language driven process was needed to develop an initial version of the DSL.

- *BPMN process modeling*: This project was a longer duration project embedded in a research project. Two people were involved in architecting the DSL. The development was mainly carried out by one developer, but access to the DSL users was very limited (only a few meetings). An agile process and test-driven development were chosen for the overall project. The main focus of the project was on the evolution of the models – the textual DSL syntax was mainly used as a convenient way to populate the models. A highly iterative version of the language model driven micro-process was chosen, as it was the most efficient way to facilitate the communication of architects and developers between the project’s regular releases.
- *Compliance metadata*: This project had a very similar setting as the BPMN process modeling and for the same reasons a language model driven micro-process was chosen here.
- *DSL editor specification*: This project also had a very similar setting as the previous two projects, but there was one major difference: the syntax of the language played an important role, as user acceptance was regarded as a crucial success factor for the project. Hence, a mockup language driven micro-process was chosen and executed in a highly iterative fashion. The mockup language design was frequently discussed with potential DSL users to get user input early on.
- *Data backup*: This project was a case study to develop a DSL for the specification of backup policies. The intention was to build an embedded DSL for XOTcl developers. Thus, a major requirement was the specification of a syntax that can easily be used by XOTcl developers. For this reason, we chose an agile development approach based on the mockup language driven process variant. In particular, we used our own XOTcl expertise and collected feedback from two other domain experts (XOTcl developers) to iteratively adapt the syntax. The backup DSL was then implemented as an embedded XOTcl DSL that can be used in other XOTcl programs.
- *Document archiving*: This project was an industry project with the goal to reengineer a large scale industry system. The system had a large set of developers. Little languages were used for various tasks in this system. The goal of the project was to architecturally improve the system. The existing little languages of the system were improved using a unified DSL concept. The domain concepts as implemented by the existing system were not in question. The main users of the DSL were the system developers and teams configuring the system for customers. For these reasons, it made sense to follow the “extracting DSL from existing system” micro-process – and then gradually improve the design in frequent meetings.

- *Multimedia home platform*: This project was an industry project with the goal to develop a proof-of-concept prototype. The main development work was carried out by one PhD student. The design was carried out in a larger group, and also frequently discussed with 10 potential customers. However, over the course of this project these meetings were rather infrequent. The project was then taken over by the developers of the company. The mockup language driven micro-process was chosen, as people on the customer sites were unfamiliar with modeling concepts such as the UML. The syntax-driven discussion of the DSL concepts, based on examples, helped to obtain as much as possible of the domain experience in a short time frame.
- *Role-based access control*: In this project we built an XML-based RBAC DSL. Because of our expertise in the RBAC domain we chose the language driven variant of the development process. Section 7 describes this project in detail.
- *Role-based access control*: This project was a follow-up case study of our first RBAC DSL project (see also Section 7). In particular, we aimed to build an embedded DSL for the xoRBAC platform. In this project, changing or extending the xoRBAC platform was not intended. Moreover, an API documentation of xoRBAC was available. Therefore, we chose to follow the “extracting DSL from existing system” variant of our process. Because of our experiences from the first RBAC DSL project, and because of our detailed knowledge of the xoRBAC platform, we were able to conduct the project in only four weeks. No additional domain experts were involved in the project.
- *SOA architectural views*: This project was a research project and development was carried out by 5 PhD students. Domain knowledge is provided by regular architecting sessions, patterns on the domain, and infrequent discussions and reviews with industrial DSL users. Because the different developers work on separate areas and DSLs with some overlaps, we decided to follow the language model driven micro-process, as it turns out to be efficient for the discussions. This was feasible because most of the domain experts were familiar with modeling abstractions.
- *Software testing*: This DSL project was part of a larger research project where we investigated the testing of dynamically changing runtime structures and dynamically changing runtime behavior of software programs. Our target platform was the dynamic object-oriented programming language XOTcl, and the intended users (i.e. domain experts) are XOTcl developers. Thus, the goal of this project was to build an embedded testing DSL in XOTcl. At the beginning of the project we only had basic knowledge of the testing domain. Therefore, we decided to choose the mockup language driven process variant and started with the definition of a concrete syntax. Two other XOTcl developers were consulted as domain experts to provide feedback on the embedded DSL. Both domain experts were available on short notice. Based on the discussions, the interim versions of the resulting mockup language were iteratively refined.
- *Software architectural knowledge*: This project was a small scale project that was conducted as part of

the infrastructure work needed for a research project. As it has a very similar setting as the bibliography management project, we used the mockup language driven micro-process for the same reasons.

- *Software documentation:* This project had a similar setting as the bibliography management project and hence we used the mockup language driven micro-process for the same reasons.
- *Workflow definition:* This project was a case study to develop an embedded DSL for the specification of workflows on the source code level. In particular, the DSL should provide language abstractions for the definition of processes that consist of several tasks which are performed sequentially and/or in parallel. Moreover, the tasks are assigned to roles and only owners of an adequate role are allowed to perform the corresponding tasks. From our experiences we knew that it is often convenient to develop embedded DSLs using the mockup language driven process variant. As always when using this process variant, the DSL was developed in an iterative fashion. In this project, we had one external domain expert who provided valuable insights in the process modeling domain. The domain expert was available on short notice which allowed for rapid feedback loops.

9 Concluding Remarks

We introduced an experience-based approach for the systematic development of DSLs. The approach provides guidance for DSL developers and facilitates the communication between engineers and domain experts. Moreover, it provides an integrated view on the respective domain and the different models defining the corresponding DSL. Our research is based on experiences from numerous DSL projects and a substantial amount of studies and experiments in the field (see Section 8). In addition to our own experiences, we applied techniques like observations and interviews with other professional developers to shape the approach, and we successfully applied the approach in follow-up projects.

In particular, we identified the activities that we conduct when engineering a DSL. We described the definition process of DSL language models and explained the activities to define a behavior definition. Moreover, we described how we specify a concrete syntax that conforms to the DSL's abstract syntax, and the integration of the different DSL models with the respective target platform. Our engineering process is intended to be performed in an incremental fashion to allow for an evolutionary enhancement of DSLs, including the definition of different concrete syntaxes for the same DSL. Moreover, to achieve a high flexibility we described how the engineering process can be tailored to the specifics of actual projects. This means, our approach allows for the adaption of the development process, e.g. by adapting the control flow of the process, or by adding or removing (sub-)activities. Furthermore, the approach can be used with any suitable modeling technique and does not prescribe specific modeling languages. Therefore, DSL developers are free to choose the most appropriate tools and techniques when engineering a DSL.

Our approach is in principle applicable whenever a DSL following a process similar to the micro-process vari-

ants presented should be developed. Tailoring allows it to be customized to specific project requirements. As the micro-processes are based on our experiences in developing model-driven and dynamic DSLs, they should be applicable in projects based on these DSL design and implementation techniques. If completely different or innovative techniques are used for designing and implementing the DSLs our micro-process variants might not be applicable without significant tailoring. The benefit of using our approach is that DSL designers and developers can rely on proven micro-processes and can focus on their work task at hand, instead of developing their own processes. Compared to general purpose processes, our approach has the benefit to focus on the specific activities and artifacts of DSL development. That is, it ensures that the specific tasks of DSL development are considered and not forgotten. The micro-processes can easily be used in most general purpose SW engineering processes, but tailoring them into these general purpose processes requires some effort. Compared to rigid process models, explaining step by step how to develop a DSL, our approach has the benefit to offer more freedom and design choices. However, as a drawback some effort for tailoring is required.

Acknowledgment

This work was partially supported by the European Union FP7 project COMPAS, grant no. 215175.

References

- [1] N. Allen, C. Shaffer, and L. Watson. Building Modeling Tools that Support Verification, Validation, and Testing for the Domain Expert. In *Proc. of the 37th Winter Simulation Conference (WSC)*, pages 419–426, December 2005.
- [2] The ATLAS Model Management Architecture (AMMA) Homepage. <http://www.sciences.univ-nantes.fr/lina/atl/AMMAROOT/>, 2008.
- [3] K. Balasubramanian, A. Gokhale, G. Karsai, J. Sztipanovits, and S. Neema. Developing Applications Using Model-Driven Design Environments. *IEEE Computer*, 39(2):33–40, February 2006.
- [4] K. Beck and M. Fowler. *Planning Extreme Programming*. Addison-Wesley, 2000.
- [5] J. Bentley. Programming Pearls – Little Languages. *Communications of the ACM*, 29(8):711–721, August 1986.
- [6] K. Bierhoff, E. Liongosari, and K. Swaminathan. Incremental Development of a Domain-Specific Language That Supports Multiple Application Styles. In *Proc. of the 6th OOPSLA Workshop on Domain-Specific Modeling*, October 2006.
- [7] D. Chamberlin and R. Boyce. Sequel: A structured english query language. In *Proc. of the ACM SIGFIDET (now SIGMOD) Workshop on Data Description, Access and Control*, pages 249–264, May 1974.
- [8] T. Cleenewerck. Component-based DSL Development. In *Proc. of the 2nd International Conference on Generative Programming and Component Engineering (GPCE)*, pages 245–264. Springer-Verlag, September 2003.
- [9] P. Clements and L. Northrop. *Software Product Lines – Patterns and Practices*. Addison-Wesley, 2002.
- [10] S. Cook, G. Jones, S. Kent, and A. Wills. *Domain-Specific Development with Visual Studio DSL Tools*. Addison-Wesley Professional, 2007.
- [11] J. S. Cuadrado and J. G. Molina. Building Domain-Specific Languages for Model-Driven Development. *IEEE Software*, 24(5):48–55, September/October 2007.
- [12] K. Czarnecki and U. Eisenecker. *Generative Programming – Methods, Tools, and Applications*. Addison-Wesley, 2000.

- [13] Eclipse Modeling Framework Project (EMF) – Homepage. <http://www.eclipse.org/modeling/emf/>, 2008.
- [14] E. Evans. *Domain-Driven Design – Tackling Complexity in the Heart of Software*. Addison-Wesley, 2004.
- [15] D. Ferraiolo, J. Barkley, and D. Kuhn. A Role-Based Access Control Model and Reference Implementation within a Corporate Intranet. *ACM Transactions on Information and System Security (TISSEC)*, 2(1):34–64, February 1999.
- [16] D. Ferraiolo, R. Kuhn, and R. Chandramouli. *Role-Based Access Control, Second Edition*. Artech House, 2007.
- [17] M. Fowler. *Analysis Patterns: Reusable Object Models*. Addison-Wesley, 1996.
- [18] M. Fowler. Language Workbenches: The Killer-App for Domain Specific Languages? <http://martinfowler.com/articles/languageWorkbench.html>, June 2005.
- [19] J. Freeze. Creating DSLs with Ruby. *Ruby Code & Style*, March 2006.
- [20] The Generic Modeling Environment Homepage. <http://w3.isis.vanderbilt.edu/Projects/gme/>, 2008.
- [21] P. Graham. *On Lisp – Advanced Techniques for Common Lisp*. Prentice Hall, 1993.
- [22] E. Grant, K. Narayanan, and H. Reza. Rigorously Defined Domain Modeling Languages. In *Proc. of the 4th OOPSLA Workshop on Domain-Specific Modeling*, October 2004.
- [23] J. Greenfield, K. Short, S. Cook, and S. Kent. *Software Factories: Assembling Applications with Patterns, Frameworks, Models & Tools*. J. Wiley and Sons Ltd., 2004.
- [24] R. Herndon and V. Berzins. The Realizable Benefits of a Language Prototyping Language. *IEEE Transactions on Software Engineering (TSE)*, 14(6):803–809, June 1988.
- [25] J. Highsmith and A. Cockburn. Agile Software Development: The Business of Innovation. *IEEE Computer*, 34(9):120–122, September 2001.
- [26] J. Highsmith and A. Cockburn. Agile Software Development: The People Factor. *IEEE Computer*, 34(11):131–133, November 2001.
- [27] A. L. Hors, P. L. Hegaret, L. Wood, G. Nicol, J. Robie, M. Champion, and S. Byrne. Document Object Model (DOM) Level 3 Core Specification, Version 1.0. <http://www.w3.org/TR/DOM-Level-3-Core>, April 2004. W3 Consortium Recommendation.
- [28] P. Hudak. Building Domain-Specific Embedded Languages. *ACM Computing Surveys*, 28, December 1996.
- [29] Information technology – Syntactic metalanguage – Extended BNF – ISO/IEC 14977:1996. http://www.iso.org/iso/iso_catalogue/catalogue_tc/cashcatalogue_detail.htm?csnumber=26153, 1996.
- [30] Information technology – Database languages – SQL – Part 1: Framework (SQL/Framework) – ISO/IEC 9075-1:2003. http://www.iso.org/iso/iso_catalogue/catalogue_tc/cashcatalogue_detail.htm?csnumber=34132, 2003.
- [31] I. Jacobson, G. Booch, and J. Rumbaugh. *The Unified Software Development Process*. Addison-Wesley, 1999.
- [32] N. Jones. An Introduction to Partial Evaluation. *ACM Computing Surveys*, 28(3):480–503, 1996.
- [33] S. Kelly and J. Tolvanen. *Domain-Specific Modeling: Enabling Full Code Generation*. John Wiley & Sons, 2008.
- [34] P. Klint, R. Lämmel, and C. Verhoef. Toward an Engineering Discipline for Grammarware. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 14(3):331–380, July 2005.
- [35] T. Kosar, P. Martinez-Lopez, P. Barrientos, and M. Mernik. A preliminary study on various implementation approaches of domain-specific language. *Information and Software Technology*, 50(5):390–405, April 2008.
- [36] I. Kurtev, J. Bézivin, F. Jouault, and P. Valduriez. Model-based DSL Frameworks. In *Companion to the 21st ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*, pages 602–616, October 2006.
- [37] L. Lamport. *LaTeX: A Document Preparation System (2nd Edition)*. Addison-Wesley, 1994.
- [38] P. Landin. The Next 700 Programming Languages. *Communications of the ACM (CACM)*, 9(3):157–166, March 1966.

- [39] L^AT_EX Project Homepage. <http://www.latex-project.org>, 2008.
- [40] A. Ledeczi, M. Maroti, A. Bakay, and B. Karsai. The Generic Modeling Environment. In *Proc. of the IEEE International Workshop on Intelligent Signal Processing (WISP)*, May 2001.
- [41] Logix: Multi-Language Programming Homepage. <http://www.livelogix.net/logix/>, 2008.
- [42] J. Luoma, S. Kelly, and J. Tolvanen. Defining Domain-Specific Modeling Languages: Collected Experiences. In *Proc. of the 4th OOPSLA Workshop on Domain-Specific Modeling*, October 2004.
- [43] T. Mens and P. V. Gorp. A Taxonomy of Model Transformation. *Electronic Notes in Theoretical Computer Science*, 152:125–142, 2006.
- [44] M. Mernik, J. Heering, and A. Sloane. When and How to Develop Domain-Specific Languages. *ACM Computing Surveys*, 37(4):316–344, December 2005.
- [45] Microsoft. Microsoft Modeling Platform (code named Oslo). <http://msdn.microsoft.com/en-us/library/cc709420.aspx>, 2008.
- [46] G. Neumann and M. Strembeck. Design and Implementation of a Flexible RBAC-Service in an Object-Oriented Scripting Language. In *Proc. of the 8th ACM Conference on Computer and Communications Security (CCS)*, pages 58–67, November 2001.
- [47] G. Neumann and M. Strembeck. A Scenario-driven Role Engineering Process for Functional RBAC Roles. In *Proc. of 7th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 33–42, June 2002.
- [48] OCL 2.0 Specification. <http://www.omg.org/technology/documents/formal/uml.htm>, May 2006. Version 2.0, formal/06-05-01, The Object Management Group.
- [49] OMG Unified Modeling Language (OMG UML): Superstructure. <http://www.omg.org/technology/documents/formal/uml.htm>, November 2007. Version 2.1.2, formal/2007-11-02, The Object Management Group.
- [50] Open Architecture Ware. [openArchitectureWare](http://www.openarchitectureware.org/). <http://www.openarchitectureware.org/>, 2008.
- [51] R. Paul. Designing and Implementing a Domain-Specific Language. *Linux Journal*, 2005(135), July 2005.
- [52] S. Pemberton, D. Austin, T. Celik, D. Dominiak, H. Elenbaas, B. Epperson, M. Ishikawa, S. Matsui, S. McCarron, A. Navarro, S. Peruvemba, R. Relyea, S. Schnitzenbaumer, and P. Stark. XHTML 1.0 The Extensible HyperText Markup Language (Second Edition) – A Reformulation of HTML 4 in XML 1.0. <http://www.w3.org/TR/xhtml1>, August 2002.
- [53] E. Raymond. *The Art of Unix Programming*. Addison-Wesley, 2003.
- [54] D. Roberts and R. Johnson. Evolve Frameworks into Domain-Specific Languages. In *Proc. of the Pattern Languages of Programs Conference (PLoP)*, September 1996.
- [55] D. Roberts and R. Johnson. Patterns for Evolving Frameworks. In *Pattern Languages of Program Design 3*, pages 471–486. Addison-Wesley, 1997.
- [56] R. Sandhu, E. Coyne, H. Feinstein, and C. Youman. Role-Based Access Control Models. *IEEE Computer*, 29(2):38–47, February 1996.
- [57] D. C. Schmidt. Model-Driven Engineering – Guest Editor’s Introduction. *Computer*, 39(2):25–31, February 2006.
- [58] B. Selic. The Pragmatics of Model-Driven Development. *IEEE Software*, 20(5):19–25, 2003.
- [59] S. Sendall and W. Kozaczynski. Model Transformation: The Heart and Soul of Model-Driven Software Development. *IEEE Software*, 20(5):42–45, 2003.
- [60] D. Spinellis. Notable design patterns for domain-specific languages. *Journal of Systems and Software*, 56(1):91–99, February 2001.
- [61] T. Stahl and M. Völter. *Model-Driven Software Development*. John Wiley & Sons, 2006.

- [62] M. Strembeck. Conflict Checking of Separation of Duty Constraints in RBAC - Implementation Experiences. In *Proc. of the Conference on Software Engineering (SE 2004)*, pages 224–229, February 2004.
- [63] M. Strembeck. A Role Engineering Tool for Role-Based Access Control,. In *Proc. of the 3rd Symposium on Requirements Engineering for Information Security (SREIS)*, August 2005.
- [64] M. Strembeck and G. Neumann. An Integrated Approach to Engineer and Enforce Context Constraints in RBAC Environments. *ACM Transactions on Information and System Security (TISSEC)*, 7(3):392–427, August 2004.
- [65] M. Strembeck and U. Zdun. Definition of an Aspect-Oriented DSL using a Dynamic Programming Language. In *Proc. Of the Workshop on Open and Dynamic Aspect Languages (ODAL)*, March 2006.
- [66] M. Strembeck and U. Zdun. Modeling Interdependent Concern Behavior Using Extended Activity Models. *Journal of Object Technology (JOT)*, 7(6):143–166, July-August 2008.
- [67] tDOM Homepage. <http://www.tdom.org>, 2008.
- [68] S. Thibault, R. Marlet, and C. Consel. Domain-Specific Languages: From Design to Implementation Application to Video Device Drivers Generation. *IEEE Transactions on Software Engineering (TSE)*, 25(3):363–377, May/June 1999.
- [69] H. Thompson, D. Beech, M. Maloney, and N. Mendelsohn. XML Schema Part 1: Structures Second Edition. <http://www.w3.org/TR/xmlschema-1>, October 2004. W3C Recommendation.
- [70] E. Valentin and A. Verbraeck. Guidelines for Designing Simulation Building Blocks. In *Proc. of the 34th Winter Simulation Conference (WSC)*, pages 563–571, December 2002.
- [71] A. van Deursen and P. Klint. Little Languages: Little Maintenance? *Journal of Software Maintenance*, 10(2):75–92, March/Arpil 1998.
- [72] Q. Wang and G. Gupta. Rapidly Prototyping Implementation Infrastructure of Domain Specific Languages: A Semantics-based Approach. In *Proc. of the ACM Symposium on Applied Computing (SAC)*, pages 1419–1426, March 2005.
- [73] D. Weiss and C. Lai. *Software Product-Line Engineering: A Family-Based Software Development Process*. Addison-Wesley Longman Publishing, 1999.
- [74] D. Wile. Lessons learned from real dsl experiments. *Science of Computer Programming*, 51(3):265–290, June 2004.
- [75] U. Zdun. Concepts for Model-Driven Design and Evolution of Domain-Specific Languages. In *Proc. of the International Workshop on Software Factories at OOPSLA*, pages 1–6, October 2005.
- [76] U. Zdun. Tailorable language for behavioral composition and configuration of software components. *Computer Languages, Systems and Structures: An International Journal*, 32(1):56–82, 2006.
- [77] U. Zdun, M. Strembeck, and G. Neumann. Object-based and class-based composition of transitive mixins. *Information and Software Technology*, 49(8):871–891, August 2007.