

Design Principles for Internal Domain-Specific Languages: A Pattern Catalog Illustrated by Ruby

Sebastian Günther, School of Computer Science, University of Magdeburg
and Thomas Cleenewerck, Software Language Engineering Lab, Vrije Universiteit Brussel

Some programming languages, especially dynamic programming languages, offer suitable mechanisms for the construction of internal domain-specific languages (DSL). Internal DSLs inherit the facilities of their host language such as the availability of libraries, frameworks, tool support, and other DSLs. When developing an internal DSL, there are two challenges. First, to cope with the host language's syntactic and semantic restrictions. Careful and thoughtful extensions and modifications of the host language are crucial to overcome these restrictions without reverting to poor language design. Second, to support several design principles that are genuine for a DSL. Although there is an extensive body of knowledge about DSL design principles and desirable quality properties, it remains difficult to apply them, or to reason about whether a particular DSL exhibits specific desirable principles.

Our objective is to put the two perspectives together. We research the most important design principles of a DSL and show how different patterns can be used to support these principles. This allows us to produce an extensive pattern catalog which is the foundation of principled approach for designing internal DSLs. The patterns can be used to assess the design quality of a DSL and structure its implementation. In particular, we show a complex DSL example illustrating each principle and the corresponding patterns. While we stick to Ruby for the explanation and application of the patterns, we also name known uses in Python, Scala, and Smalltalk. Patterns are explained with their context, problem, forces, solution, and consequences. We also explain the patterns with the classical structure of intent, motivation, forces, implementation and their consequences. Finally we reflect upon the pattern utilization by discussing their strengths and weaknesses.

Categories and Subject Descriptors: D.3.3 [Programming Languages]: Language Constructs and Features—*Patterns*

General Terms: Design

Additional Key Words and Phrases: design principles, domain-specific languages, dynamic programming languages

1. INTRODUCTION

Domain-specific languages are programming languages or executable specification languages that offer, through appropriate notations and abstractions, expressive power tailored for a specific problem domain or application area [Van Deursen et al. 2000; Hudak 1998]. At the implementation level, two different types of DSLs can be identified. Specially crafted external DSLs require their own parsers, interpreter or compilers, while internal DSLs are built on top of an existing programming language (called the host language), allowing them to reuse the host language's infrastructure including IDEs and compilers [Mernik et al. 2005]. DSL are a common practice in software engineering as several examples show [Arnold et al. 1995; Groote et al. 1995; Thibault et al. 1997; Latry et al. 2007; Munnely and Clarke 2008].

The wide spectrum of domains, internal DSLs, and external DSLs, is threatening to deteriorate the original meaning of the term DSL. The term is applied often and without to much consideration what the properties of a

Author's address: Sebastian Günther, Universitätsplatz 2, D-39106 Magdeburg, Germany, email: sebastian.guenther@ovgu.de; Thomas Cleenewerck, Vrije Universiteit Brussel, Faculty of Sciences, DINF - SOFT, Pleinlaan 2, B-1050 Brussels, Belgium, email: tcleenew@vub.ac.be
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission. A preliminary version of this paper was presented in a writers' workshop at the 17th Conference on Pattern Languages of Programs (PLoP). PLoP'10, October 16-18, Reno, Nevada, USA. Copyright 2010 is held by the author(s). ACM 978-1-4503-0107-7 PLoP'10, October 16-18, Reno, Nevada, USA. Copyright 2010 is held by the author(s). ACM 978-1-4503-0107-7

DSL are. Analyzing the relevant literature [Mernik et al. 2005; Tanter 2008; Ladd and Ramming 1994; Weinberg 1971; Coplien 1999], however, reveals that there are several design and implementation principles of DSLs such as *abstraction* to make domain concepts explicit programming language constructs, *absorption* to integrate implicit and common domain knowledge in the DSL, or *compression* to yield a concise domain notation.

While these principles are general enough to provide guidelines for the design and implementation of a DSL, it is rather hard to explain whether an implemented DSL supports or embodies these principles. It is also difficult to plan the implementation upfront without a sophisticated language that allows expressing design decisions with respect to the principles. Hence the question: how to achieve this? We argue that the careful selection and application of patterns help here. Patterns support the DSL implementation by embodying these principles. By providing a deep understanding of the design principles together with a pattern catalog, we facilitate the planning and implementation of DSLs that exhibit the design principles.

This paper's contribution is a rich pattern catalog that helps to approach DSL design based on a set of design principles. Through using the patterns that adhere to the design principles, the design quality of a DSL can be assessed. To the best of our knowledge, this particular combination of principles and patterns has not been elaborated in other literature yet. For feasibility of this first-hand analysis, the paper focuses upon selected principles and illustrates the application of patterns grouped into form (alternative syntactic forms for the expressions), modification (small changes to the language, like the provision of aliases), and support. The patterns are illustrated using the Ruby programming language and two internal DSL designed within this language. However, the patterns are available in other programming languages too, which we show for Python, Scala, and Smalltalk by pointing to known uses.

The paper is intended for novice as well as experienced DSL engineers. For experienced DSL engineers our work consolidates the design qualities of their DSLs enabling them to learn and improve their DSL designs. Novice DSL designers can use the principled pattern approach as a reference work to create internal DSLs.

We structure the paper as follows. Section 2 explains the used pattern structure and grouping. In Section 3, the design principles are explained. Sections 4 to 6 list the patterns with their intent, motivation, forces, structure, consequences, and additional known uses outside the herein presented examples. Section 7 revisits the design principles and, using a larger DSL examples, shows how the patterns are used in combination to design a DSL with respect to the principle. Section 8 discusses our experiences in using combinations of patterns and their influence on all principles. Section 9 explains the related work and finally Section 10 summarizes this paper. We apply following formats: *keywords*, *design principles*, and PATTERNS.

2. PATTERN STRUCTURE AND COMMON EXAMPLE

2.1 Pattern Structure

The pattern explanation has a structure that is close to the initial introduction of patterns [Gamma et al. 1997] and recent publications [Guerra et al. 2009; Correia et al. ; Zdun and Strembeck 2009]. The following structure is used to explain each pattern:

- *Context*: Explains the coarse situation in which the pattern is applied.
- *Problem*: A question expressing the central concern the pattern addresses.
- *Design Impact*: Explains which principles are realized by applying this pattern and explains the particular achieved result.
- *Instructions*: Explanation what language mechanisms are used to provide a solution.
- *Example*: Based on the common examples, a detailed utilization of the pattern using the Ruby programming language is given. Source code examples and sometimes diagrams provide further illustration (see ► Figure 1 for used symbols).

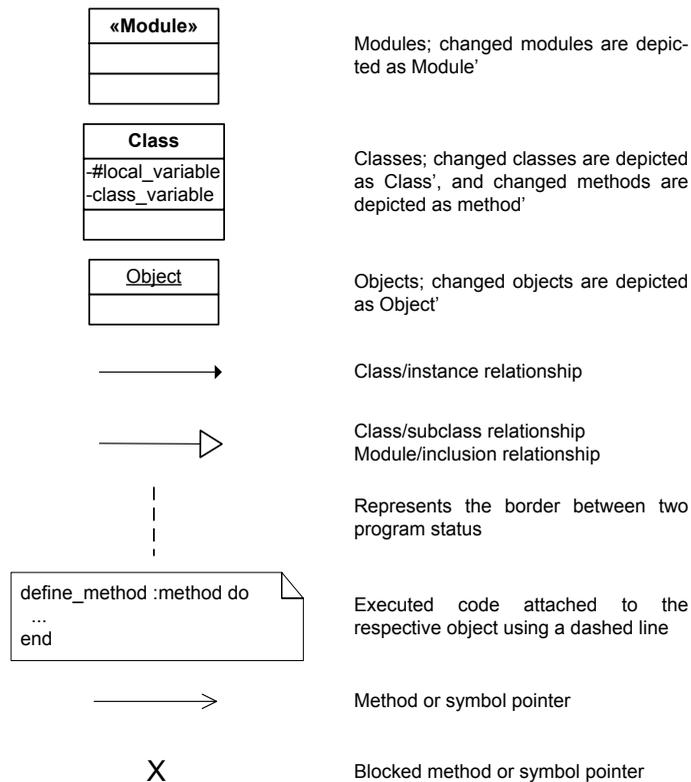


Fig. 1: Symbols used to explain the pattern's structure.

- *Implementation Impact*: Explains how this pattern's utilization impacts the DSL implementation or the DSL itself.
- *Known Uses*: List of applications that use this particular pattern, given for Ruby [Flanagan and Matsumoto 2008; Thomas et al. 2009], Python [Lutz 2009; Summerfield 2010], Scala [Odersky et al. 2008; Wampler and Payne 2009], and Smalltalk [Sharp 1997; Black et al. 2009a; 2009b]. The analyzed DSLs are listed in the next sections.
- *Related Pattern*: Finally, a list of closely related patterns.

2.2 Analyzed DSLs

Following DSLs were analyzed in order to find the uses for our patterns.

- Ruby
 - *Rails*: Web-application framework with a strong Model-View-Controller architecture. Used in several open-source and commercial projects. In addition we analyzed the *Builder* library that is used by Rails to build configuration objects representing XML-like data structures (<http://rubyonrails.com>).
 - *ActiveRecord*: Rails standard database abstraction (<http://ar.rubyonrails.com>)
 - *RSpec*: Framework and DSL for behavior-driven development (<http://rspec.info>).

- *Sinatra*: Lightweight web framework that uses declarative expressions which look like a DSL for web applications (<http://sinatrarb.com>).
- Python
 - *Bottle*: Similar to Sinatra, a lightweight web framework that uses normal Python function declarations and annotations to specify how queries are responded in the application (<http://bottle.paws.de/>).
 - *Should DSL*: Supports simple behavior-driven testing by adding expressions like `should`, `be`, and `be_greater_than` as expressions to Python (<http://github.com/hugobr/should-dsl>).
 - *Cryptlang*: A DSL for defining cryptographic algorithms and random number generators [Agosta and Pelosi 2007].
 - *Constellation Launch Control DSL*: Provides a command like language for “programming test, checkout, and launch processing applications for flight and ground systems” [Bennett et al. 2008].
- Scala
 - *Actors*: Actors is a build-in Scala DSL for asynchronous messaging and frequently deployed for concurrent programming. It uses concepts like a messaging box and special operators used to send and retrieve messages (<http://www.scala-lang.org/api/current/scala/actors/package.html>).
 - *Spec*: DSL for behavior driven development similar to the block-like nature of RSpec (<http://code.google.com/p/specs/>).
 - *Apache Camel*: Helps to express the configuration of enterprise systems that are compatible with Java, allowing to define routes and APIs for the components (<http://camel.apache.org/scala-dsl.html>) Implementation of SQL as a Scala DSL [Spiewak and Zhao 2009].
- Smalltalk
 - *Mondrian* is an information visualization engine that lets the visualization be specified via a script. It is built in VisualWorks Smalltalk and Pharo (<http://www.moosetechnology.org/tools/mondrian>).
 - *Glamour* is an engine for scripting browsers (<http://www.moosetechnology.org/tools/glamour>).

2.3 Pattern Grouping

The patterns are grouped in three classes to distinguish their original concern. *Form* patterns allow to use alternative syntactic forms for expressions, and can also change the basic paradigm expressions stem from (like the change from imperative expressions to declarative expressions using functional programming). *Modification* patterns change the language on a smaller scale such as aliasing existing entities. Opposed to the form pattern, modifications can also require semantic changes such as overriding existing methods. Finally, the *support* pattern’s role is to help the other patterns in achieving their respective goals.

2.4 Common Examples

2.4.1 *Product Line Configuration Language*

Software product lines address the challenge of structuring and reusing software by providing a set of valuable production assets [Czarnecki and Eisenecker 2000]. Such assets have the form of documentation, configuration, source code, libraries and more. Typically, a “product line is a group of products sharing a common, managed set of features” [Withey 1996]. Features describe modularized core functionality of a software [Batory et al. 2003].

Software product lines can be modeled as a feature tree [Czarnecki and Eisenecker 2000] – it structures the features, their relationships, and the constraints they impose on each other (like to choose from a set of alternative features). A sample feature tree, showing the Graph Product Line (GPL) case study explained in [Lopez-Herrejon and Batory 2001], can be seen in ► Figure 2. This product line represents different types of graphs (weighted/unweighted and directed/undirected) and their corresponding property and search algorithms.

In order to model such feature trees, we designed the internal Ruby DSL called Software Product Line Configuration Language (SPLCL). The SPLCL uses a declarative textual notation to specify features with their

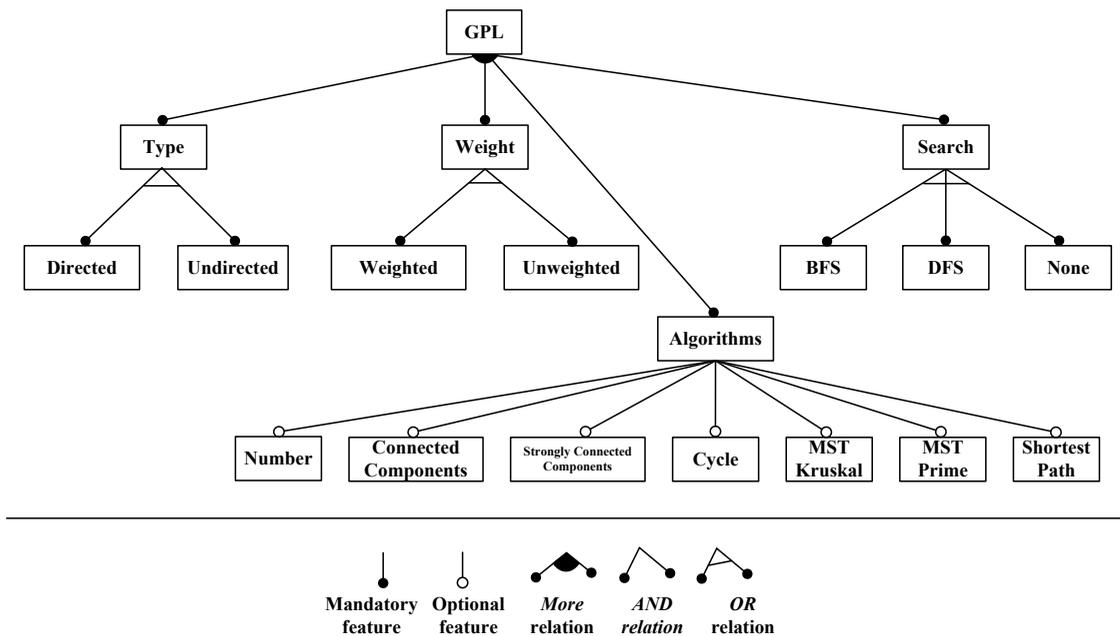


Fig. 2: The Graph Product Line example from [Lopez-Herrejon and Batory 2001].

name, position in the feature tree, subfeatures, and constraints. The DSL was developed as a research project by one of this paper's authors [Günther 2009] and has since then continually evolved and served as vehicle for DSL integration research, amounting to four men months of work effort.

```

1 epl = Feature.configure do
2   name :GPL
3   root
4   subfeatures :Type, :Weight, :Search, :Algorithms
5   requires :GPL => "all :Type, :Weight, :Search, :Algorithms"
6 end

```

Fig. 3: Example of the SPLCL DSL to configure the root feature of the GPL.

► Figure 3 shows an excerpt from this DSL how the root feature is configured. Here is the meaning of this code:

- *Line 1*: Calling a constructor for a feature which receives configuration options in the following lines.
- *Line 2*: Set the name of the feature to GPL.
- *Line 3*: Express that the feature is at the root position.
- *Line 4*: Configure the subfeatures.
- *Line 5*: A constraint in order to activate this feature. It literally expresses that when the GPL feature is activated, the features Type, Weight, Search, and Algorithms have also to be activated.

2.4.2 Infrastructure Management DSLs

Infrastructure management takes care of the initialization, configuration, and maintenance of several servers forming the infrastructure of a stakeholder. Due to continuously increasing requirements with regard to flexibility, high automation support that limits the amount of manual involvement becomes crucial. We designed a set of three DSLs that support various tasks in this domain. The *Boot-DSL* identifies and installs machines, the *Software-Deployment Planning DSL* (SDP-DSL) expresses relationships between packages, and finally the *Configuration Management DSL* (CM-DSL) configures and installs packages. The currently supported operating systems are Linux-based. The supported hypervisors are Amazon EC2¹ and VMware ESX². The DSLs are the collaborative effort of two bachelor student's theses and one of this paper's authors ongoing research project, amounting to more than two men years.

```
1 app_server = Machine "Application Server" do
2   type EC2
3   owner "sebastian.guenther@ovgu.de"
4   os :debian
5   hypervisor do
6     ami "ami-dcf615b5"
7     source "alestic/debian-5.0-lenny-base-2009..."
8     size :m1_small
9     securitygroup "default"
10    private_key "ec2-us-east"
11  end
12  hostname "admantium.com"
13  monitor :cpu, :ram
14  bootstrap!
15 end
16
17 deploy "Redmine" do
18   enroll "Redmine/Database", :on => app_server + aux_server do
19     #...
20   end
21 end
```

Fig. 4: Example of the Boot-DSL and the SDP-DSL to configure a server and installs a package on it.

► Figure 4 shows an example for the Boot-DSL and the SDP-DSL, with the following meaning:

- *Line 1 – 15*: An example of the Boot-DSL that opens a constructor, sets the default name, and passes other configuration options in the form of a block.
 - *Line 2 – 4*: Configures the type, owner, and the operating system of this machine.
 - *Line 5 – 11*: A nested block that specifies the hypervisor-specific options like the Amazon machine image (ami) id and its source, the size, the securitygroup (determines the available services and the open ports of the machine, such as SSH on port 22), and the private key used to access this machine via SSH.
 - *Line 12*: The hostname of the machine which is used inside the network.
 - *Line 13*: Expresses to monitor the CPU and RAM usage of this machine.
 - *Line 14*: Finally an expression to install basic software.

¹<http://aws.amazon.com/ec2/>

²<http://www.vmware.com/products/esx/>

— *Line 17 – 21*: An example of the SDP-DSL to deploy the package Redmine (an open-source project management application, see <http://www.redmine.org/>).

— *Line 18*: Redmine has several dependencies to satisfy, and in this case the database is configured to be enrolled on the application server as well as an auxiliary server for obtaining a master-slave relationship. Other packages would be implemented likewise.

3. DESIGN PRINCIPLES

General purpose languages were designed for expressiveness and scalability to solve a large number of computational problems. They are more concerned with providing efficiency and use language constructs which are independent of a particular application domain. Instead, DSLs absorb concepts of the problem domain directly into the solution domain, and language constructs are used to represent the domain. But this requires a thoughtful approach.

Especially dynamic programming languages, such as Ruby and Python, provide several options to implement DSLs. Through several syntactic alternatives, support for multiparadigm-development, and metaprogramming capabilities, these languages can be customized in terms of their syntax and semantics. However, it is challengingly to know the forms of these modifications and to use them coherently in the process of designing a DSL.

Studying works from the language design community, the software engineering community, and the domain engineering community, we see that a vast amount of principles is offered, but they are rather vague and non-constructive as the following examples show:

— The language design community focuses more on qualities. The important exception is multiparadigm programming by Coplien, where he defined abstraction as a major design principle and gave it a constructive definition in the context of communality and variability [Coplien 1999].

— In software engineering, it is widely acknowledged that “Inside every domain-specific framework, there is a language crying to get out” [Foote B 1998]. But how to construct a language out of packages, classes, and methods is far less obvious.

— The domain engineering community occasionally hints at a principle: classes are said to represent the vocabulary of the problem domain [Booch 1994] and FODA suggests to apply aggregation/decomposition or generalization/specialization to the end of obtaining common assumptions [Kang et al. 1990].

In order to better guide the design of the languages behind DSLs, we analyzed the existing DSL literature [Van Deursen et al. 2000; Bentley 1986; Landin 1966] to discover the following defining characteristics: DSLs are (1) small, (2) concise programming languages or executable specification languages that offer, through (3) appropriate notations and (4) abstractions, (5) expressive power (6) tailored and (7) optimized towards a specific problem domain or application area. Each of those characteristics are embodied by the following design principles:

Generalization

Reduce the amount of concepts by replacing a group of more specific cases with a common case.

In software engineering and especially domain engineering, generalization plays an important role. First, the domain is captured in its full scope consisting of multiple concepts and their relationships. They naturally form a hierarchy in which one concept is more general than other specific concepts. Finding such hierarchies helps the developer to understand the domain. The developer then chooses how to prune, reorganize, and combine the concepts in order to express generic and specific cases. Generalization keeps DSLs small as the reduction of the amount of concepts results in a small amount of needed constructs [Coplien 1999].

Compression

Provide a concise language that is sufficiently descriptive for the domain experts.

The goal of compression is to reduce the amount of expressions or to simplify their appearance while the meaning is not changed. Compression makes DSLs concise programming languages that are still syntactically

clear for the domain experts. This reduces the code footprint of DSL programs [Ladd and Ramming 1994; Weinberg 1971]. Compression is concerned with the syntax of the DSL, so it is not to confuse with the overall expressiveness.

Abstraction

Implement suitable domain-specific abstractions to deemphasize (technical) details.

DSLs abstract twofold: From the domain and from their host language. Abstracting from the domain is to include only those entities and operations in the language that are relevant. Abstractions from the host language are either base level abstractions or meta level abstractions. On the base level, the host's basic operational model – expressed by the existing language constructs – is extended with novel expressions. At the meta-level, the existing constructs are tweaked, to modify or to extend their behavior. Both forms of abstractions are fundamental for DSLs since they have to step away from the language in order to step towards the domain. Abstraction provides DSLs explicit focused expressive power [Coplien 1999].

Representation

Design the DSLs's layout with representative notations for the domain-specific abstractions.

A DSL's representation is its syntax. The syntax consists of all language constructs, keywords, structures, and rules governing their combination. The rules determine the general layout of the language and, being an internal DSL, can also include hooks where other DSLs and host language expressions can be put. Internal DSLs cannot leave the syntactic frontiers of their host language, but they can chose to combine existing syntactic modifications to achieve a form that has little or no resemblance to the host language. Representation equips a DSL with the appropriate notations [Van Deursen et al. 2000] by defining the basic syntax of the language and its grammar to support domain-specific abstractions and notations. Choosing the right representation ensures the DSL's understandability for the domain experts.

Absorption

Implicitly absorb domain commonalities to facilitate DSL usage.

Characteristic properties of domains are commonalities for a broad set of expressions. Users are supported effectively if these commonalities are absorbed by the language – meaning that certain assumptions and concerns are not required to be explicitly expressed. DSLs with high absorption support the user by capturing these assumptions implicitly, relieving him of the burden to repeat himself. This is in stark contrast to a traditional library, where functions are dependent on each other, require a manual setup of the context, and have a restricted call order. The best example is declarative DSLs that just describe a desired state, and the DSL's implementation uses all explicitly and implicitly given information to effectuate this desired state. Absorption provides DSLs implicit focused expressive power [Tanter 2008] by implicitly integrating domain commonalities in the DSL.

Standardization

Standardize program writing to a common restrictive set of expressions for assisting users.

Language constructs express language concerns. For example, object-oriented programming languages support declaration of classes, instantiation, provision of instance variables, and more. They focus on the essential concerns of a programming language and the formulation of as few language constructs as possible to express these concerns is called involution in programming languages [65]. DSLs aim for a high involution, meaning that the language will be standardized to a restricted set of operations and expressions based on problem domain. A significant part of their grammatical and semantic specification is therefore dedicated to assist users in how to write a program that will solve their problem. Standardization guides developers in using the tailored and thus restricted domain-specific set of language constructs [Weinberg 1971; McKeeman 1976].

Optimization

Use algorithmic optimizations to improve the DSL's computational performance.

Producing performant code is an essential characteristic of programming languages. In the historic development of computer languages, the successful transition from hand-written assembly code to Fortran and its compiled machine code could only be made by showing that there are no performance decreases [Sebesta 1999]. Likewise, an internal DSL should not be slower than its host language. Also from the domain side, the provision of domain-specific optimization is important. Optimization improves the computational performance by using suitable models and domain-knowledge with an efficient implementation [Bruce 1997].

The relationships between patterns and principles will be explained in the next sections. For feasibility, we focus upon four principles that are more concerned with a DSL's syntax³: *representation*, *absorption*, *compression*, and *generalization*. Because the principles are heavily interwoven with other, we can not give a clear hierarchical decomposition of principles and their supporting pattern. Moreover, one pattern usually affects several other principles. Therefore, the paper continues with a focused explanation of the patterns grouped into the three categories *form*, *modification*, and *support*. Afterwards we will revisit the selected principles and explain how the different patterns support them. And finally, we combine both perspectives and explain our experience in using combinations of patterns and their influence on the principles.

4. FORM PATTERNS

Form patterns are providing the general layout of DSL expressions, they provide the framework in which expressions with a smaller scale are used. Patterns in this category also try to change the expression from the host languages basic paradigm to another one (like the change from imperative expressions to declarative expressions using functional programming).

4.1 Block Scope

Context

Many domains have a natural hierarchy of objects. This hierarchy should be reflected in the DSL's expressions too. For example, when building a structured document like HTML, the DSL should reflect the documents structure with syntactic indentation and at the same time provide a separate lexical scope for its expressions.

Problem

How to provide a vertical code layout for stacking hierarchical information and providing a lexically separated execution scope?

Design Impact

— *Representation*: Group related expressions together and syntactically express the hierarchy of objects as indented blocks of code.

— *Absorption*: Absorb the most relevant execution information, especially if it serves as a closure capturing the surrounding context and use them later to infer contextual information.

— *Compression*: Remove fully-qualified method calls inside a block, and let the execution context be determined by the object or method receiving the block.

³We do not explain *abstraction* because of its inherent semantic character and we also skip *optimization* since this principle is very dependent on the domain.

Instructions

Ruby allows to group expressions inside a block using the `do ... end` notation or curved brackets. The code contained inside this block has two distinguishing properties. First, it encloses its surrounding scope and can serve as a closure for specific execution states. Second, it can be executed in another context, using the variables from the surrounding lexical scope.

BLOCK SCOPE is extensively used in other DSLs for these purposes. It can also be used as a CLEAN ROOM to execute all expressions in the safe scope of a specially crafted object [Perrotta 2010], checking whether all method calls are valid or all values supplied correctly.

In Python, context managers can be used for the same purpose, although fully-qualified method calls still have to be used.

Example

Using the Boot-DSL to configure various virtual machine properties, the BLOCK SCOPE is denoted by the `do ... end` notation in Line 1 and Line 8. The properties are included as statements inside the block.

```
1 machine = Machine "Auxiliary Server" do
2   owner "sebastian.guenther@ovgu.de"
3   os :debian
4   hypervisor :ec2 do
5     ami "ami-dcf615b5"
6     source 'alestic/debian-5.0-lenny-base-2009...'
7     ...
8 end
```

Fig. 5: BLOCK SCOPE: expression example that configures various properties of a virtual machine.

The block is attached to the `Machine` method. Its definition is shown below, starting with the `def` keyword. The method receives two parameters. The first is the machine's hostname, and the second is the block captured in the variable named `block`. Line 2 – 6 checks whether the machine already exists, and if not creates a new one. Then Line 7 evaluates the block in the machine's context.

```
1 def Machine(hostname, &block)
2   machine = Machine.first :hostname => hostname
3   if not machine
4     machine = new()
5     machine.hostname = hostname
6   end
7   machine.instance_eval &block
8   return machine
9 end
```

Fig. 6: BLOCK SCOPE: implementation example. The block is captured as the variable name `block`, and then executed in the context of the created instance.

Implementation Impact

— Possibility of code injection: If the block object is dynamically constructed with input from the users, malicious users could exploit detailed knowledge of the application to read its data or perform file system operations. However, Ruby has a good support for safe levels, as well as tainted and trusted objects [Thomas et al. 2009], which reduces this threat's potential.

— Errors in the block can only be checked by executing the code, which is the task of CLEAN ROOM.

Known Uses

— Ruby

— Using Sinatra, BLOCK SCOPE is applied with the declaration of request handlers, forming a declarative expression how a web application reacts to queried URIs (the request handler is implemented in `sinatra-0.9.4/lib/sinatra/base.rb`, Line 752 – 758).

— Using RSpec, BLOCK SCOPE is applied for the declaration of example groups that begin with the `describe` keyword and group several test cases (implementation of the `describe` method is contained in `rspec-1.3.0/lib/spec/dsl/main.rb`, Line 24 – 29).

— Python

— One of this paper's authors currently develops a SQL-DSL. Thereby, the *with-statement* is used, which is a syntactic extension of Python that allows a local variable to be passed to a set of indented expressions. The expression `with Query("sqlite") as query` starts a block where `query` is a local variable. Using this object, methods representing SQL-statements are called. Because method calls inside the block still need to use fully-qualified names, its syntax is more descriptive than the Ruby version altogether.

— Scala

— In the Specs DSL, BLOCK SCOPE is used to visualize the structure of complex test cases. Test cases, called system under specification in the Specs DSL, start with a `literate` string describing the system, followed by the `should` keyword and a function. Inside the function, examples are specified by using another string description and the `in` keyword. Interestingly, more keywords can be added to these terms, so that expressions like "The web application" `should have response status ...` can be built (the `should` method is defined in `specs-1.6.2.1/src/main/scala/org/specs/specification/SpecificationSystems.scala`, Line 45 – 46, and the `in` method in `specs-1.6.2.1/src/main/scala/org/specs/specification/BaseSpecification.scala`, Line 154).

— Smalltalk

— Although closures are available in Smalltalk, each expression within a block needs a receiver. By combining message cascading with block closures like in expression `doWith: [:result | result message1; message2; ...]` we can achieve similar expressions, but it is required to use the result of one method as the argument to the next at the start of a method cascade. The combination of message cascading and closures is not common, no examples have been found in the analyzed DSLs.

Related Patterns

— KEYWORD ARGUMENTS – Alternatively to the vertical code layout provided with BLOCK SCOPE, code can be arranged horizontally by suitable grouping and naming of arguments.

4.2 Method Chaining

Context

In some domains, expressions involve a high number of concepts and operations. For example in financial transactions, accounts, identification numbers, currency information, stock information and more need to be expressed. This complexity should be reflected in the DSL too.

Problem

How to make sentence-like expressions by chaining methods together?

Design Impact

— *Representation*: Express the domain in a more natural-language like style by chaining method calls and their return values together.

— *Compression*: Use chained methods that transform an input value step-by-step, adopting a more functional-programming oriented style instead of more verbose imperative expressions.

Instructions

Normal method calls are written line by line and usually operate on variables passed between them. If this occurs, then the methods can also be chained in one line by passing the return value of one method as the argument to the next. Because Ruby does not require fully-qualified method calls in the form of “receiver.method”, the “.” can be left out, allowing to chain method calls together that look like language keywords. Scala can use the same layout if the method receives nor or just one argument. Multiple method chains enable provide DSL expressions that do not resemble their host language.

Example

An expression that uses METHOD CHAINING⁴ is shown in the following listing. It is taken from the SPLCL to create a product variant (aka instance) for the GPL.

```
1 activate feature Weighted from productline GPL
```

Fig. 7: Combining METHOD CHAINING and CLEAN METHOD CALLS for a complex yet easy to read expression.

The evaluation of this expression happens from right to left. Implementing METHOD CHAINING demands to define various methods, often empty ones that are just used in the statement but have no semantic meaning. In considering the following implementation (►Figure 8), the `productline` method receives the `GPL` object and sets a local variable (Line 2 – 4). Next, `from` is just an empty method (Line 7 – 11). Subsequently, `Weighted` is a method too, because of Ruby’s lexical scoping. Therefore, the expression `feature Weighted`⁵ is used to access the actual feature (Line 14 – 20). Finally, the `activate` method receives the feature and activates it (Line 23 – 25).

Implementation Impact

— Requires to add several “glue” methods to the surrounding scope, possibly polluting the space with empty methods.

— Limited applicability to modify existing libraries because of the required extensive modifications.

— If used with an existing library, requires to change method internals.

⁴And the later explained CLEAN METHOD CALLS pattern.

⁵Actually an instance of the CUSTOM RETURN OBJECTS pattern.

```

1 # Setting the product line to the specified argument
2 def productline(arg)
3   @productline = arg if arg.is_a? ProductLine
4 end
5
6 # Providing empty methods that are just used inside the expression
7 def from(*args)
8 end
9
10 # Accessor for the 'Weighted' feature
11 def Weighted(*args)
12   :Weighted
13 end
14
15 def feature(arg)
16   Kernel.const_get(arg)
17 end
18
19 # Finally activating the feature
20 def activate(arg)
21   arg.activate
22 end

```

Fig. 8: METHOD CHAINING – implementation of all required methods in order to execute the method chain.

Known Uses

— Ruby

— In RSpec, users chain `should` and `should_not` expression with an assertion statement like `be_nil` or `have(3).items`, combining the statements in a clean and readable way (this works by adding the `should*` method to Ruby's built-in class `Object`, which is done in `rspec-1.3.0/lib/spec/mocks/extensions/object.rb`, Line 1–3).

— Python

— The `should-dsl` is a unique extension because it uses OPERATOR EXPRESSIONS (explained later) together with a normal method call to add a new syntactic element to Python. Basically it allows the same kinds of expressions like RSpec, consisting of an object, a `should` keyword, and a test. Here is an short example: `1 |should_not| be(2)`. Inside this expression, the “|” are actually method calls that receive the left value and the right value, which is then passed to the `should_not` method. This combination allows a nice-to-read METHOD CHAINING (the “|” operators are defined in `should-dsl/src/should_dsl.py`, Line 25–38).

— The Launch Constellation DSL also uses a form of *Method Chaining*: Methods with boolean return values are used to structure conditional blocks that determine how to react to different status of the monitored applications [Bennett et al. 2008].

— Scala

— The Specs DSL uses method chaining in an extreme form because expressions can mirror complete sentences. For example we create a car object with the name “Nissan Primera”. To test whether the carname begins with the string “Niss” and not “Pors”, this expression can be used: `nissan.carname must be matching("Nis*") and not be matching("Pors*")`. Similar matchers for objects other than strings are included in the Specs library or can be added for user-specific types (the `must` method is defined in `specs-1.6.2.1/src/main/scala/org/specs/specification/Expectable.scala`, Line 172, and the string

matcher methods `be` and `matching` are defined in `specs-1.6.2.1/src/main/scala/org/specs/matcher/StringMatchers.scala`, Line 36 and 202).

— Smalltalk

— Method chains come very natural in Smalltalk as one of the syntactic design goals of the language was the ability to write nice sentences. In Mondrian, graphic shapes can be used as brushes using sentences such as `view shape rectangle size:`.

— Using Glamour to create a browser for software versions, the expression `a list title: 'versions'`, where “a” is a widget constructor, is used to create a list with a title.

Related Patterns

— **KEYWORD ARGUMENTS:** Both **METHOD CHAINING** and **KEYWORD ARGUMENTS** provide a horizontal code layout. **KEYWORD ARGUMENTS** are especially suitable when a complex expression can be refactored to nested subphrases. The additional benefit to **METHOD CHAINING** is that no “glue” methods are needed.

— **BLOCK SCOPE:** Instead of a horizontal code layout, **BLOCK SCOPE** provides a vertical layout which is especially useful to show hierarchies of objects. This pattern is recommended when very complex expressions need to be structured, where the block-like structure helps to disseminate the information. Of course, individual lines in **BLOCK SCOPE** expressions can be designed with **METHOD CHAINING**.

4.3 Keyword Arguments

Also known as **NAMED ARGUMENTS**

Context

Methods are an important cornerstone in using DSLs. They receive parameters that express the configuration of objects. However, two problems occur. First, always requiring a fixed order of passed arguments can provide problems if three or more arguments are passed. Second, just passing arguments hampers readers of DSL expression to relate the arguments with their intended meaning and relationship to the method.

Problem

How to provide better readability of multiple arguments and gain flexibility in argument passing?

Design Impact

— *Representation:* Avoid complex API with methods receiving more than two parameters and the ambiguity of calling methods with correct arguments in correct order. Instead use keyword-value pairs as relationship declarations, to state mappings, or to declaratively express the desired system state. The chain of keywords and values can furthermore be used to obtain more natural-language like sentences.

— *Absorption:* Provide abundant optional arguments that are passed to other methods. Computations for good defaults for the optional keyword arguments absorb the additional information that can be deduced by the DSL itself.

— *Compression:* Compress several singular method calls to a set of key- value pairs passed to one method.

— *Generalization:* Easily switch to a specific or concrete semantics of a method call by using multiple dispatch on arguments.

Instructions

In Ruby, hashes provide a nice syntax for expressing relationships of objects: `key => value`. Hashes can be created as inline arguments to method calls. That allows to syntactically align the method call and its indexed arguments for better readability. Using these hashes inside the method calls to check for multiple arguments can be done either with the help of an iterator or by explicitly checking for certain keys.

Example

An expression of the SDP-DSL to deploy an application on a machine is used here. The relevant excerpt is the following code.

```
1 deploy "Redmine" do
2   enroll "Redmine/Database", :on => application_server + auxiliary_server do
3     #...
4   end
5 end
```

Fig. 9: KEYWORD ARGUMENTS: the `enroll` statement uses keywords to configure deployment options.

The `enroll` method shown in the following code first receives a name parameter and then a hash. This hash is checked for its contained key-value pairs and appropriate values are set for the installation.

```
1 def enroll(package, hash)
2   @package = package
3   @target = hash[:on]
4   #...
5 end
```

Fig. 10: KEYWORD ARGUMENTS: implementation of the `enroll` method which accesses the passed hash and set values accordingly.

Implementation Impact

- Complex combinations of parsing arguments and error checking can make the method declaration cumbersome.
- More verbose method calls reduce *compression*.

Known Uses

- Ruby
 - Sinatra uses the `set` method for configuring options, it receives a single key-value pair and sets the configuration accordingly (the `set` method is implemented in `sinatra-0.9.4/lib/sinatra/base.rb`, Line 622 – 636).
 - ActiveRecord uses the built-in `autoload` method that receives a module name as a symbol and a filename in which the module is implemented. Modules configured in this way are only imported when they are called in the program (the `autoload` method is defined in Ruby's core class `Module`, which is implemented in C for the MRI version of Ruby).
- Python
 - In Bottle, URL handlers are defined with normal Python functions and annotation. The annotations determine by which URLs the methods are executed. The annotation can contain additional keywords like in this example: `@route('/store/product/:id', method='POST')`. This defines a HTTP POST request to trigger the annotated method (the `@route` annotation method is defined as a decorator in `bottle/lib/bottle.py`, Line 562–566, and calls the method defined in Line 216–232).

- The Launch Constellation DSL uses a variant of keyword passing. Instead of using hash structures directly, first strings are passed to the functions and then the values of the variables addressed by this string. Here is an example: `send_command (discrete ("VALVE1", "ON"))` [Bennett et al. 2008].
- Scala
 - The analyzed Scala DSLs do not use keyword arguments, but it is possible to implement variable length functions and pass any number of tuples inside, like in this contrived example: `configure_route("path" -> "/store/product/", "arg_arity" -> 2, "method" -> "POST")`. The arguments will form a Sequence which can be parsed conveniently with case matchers.
- Smalltalk
 - KEYWORD ARGUMENTS in Smalltalk can be realized by using method cascading, in which the receiver is stated at the beginning of an expression and where subsequent messages are send to the same receiver without having to repeat the receiver. This pattern is used in Mondrian on many occasions, pronounced examples are the sequencing of features of shapes e.g. they can be customized by a message cascade of `width:`, `height:` and `linearFillColor:within:`.
 - In Glamour, hierarchical browsers are created by a sequence of columns, each named after the message that has to be send to descend deeper into the hierarchy.

Related Patterns

- BLOCK SCOPE: The horizontal code layout for key-value pairs can also be represented as a vertical stacking by using BLOCK SCOPE. The tradeoff, however, is to implement several additional methods to handle the argument passing.
- METHOD CHAINING: For an alternative horizontal code layout, represent the key-value pairs as methods inside a chain. However, likewise to block scope, additional methods have to be implemented.

5. MODIFICATION PATTERNS

Modification patterns aim to change host language expressions for better compatibility with the domain, for example by renaming existing methods. Opposed to the form patterns, modification patterns can also demand semantic changes such as overriding existing methods.

5.1 Seamless Constructor

Context

In languages with an object-oriented core, the creation of new objects is typically expressed with keywords like `new`. In a DSL, using this keyword explicitly may not be intended, especially not when the keyword has another, domain-specific meaning. Many languages combine objects on-the-fly by using expressions which uniquely express the type and value of the object to be created. For example, entering any number in Ruby or Python expressions will create a number object with the specified value.

Problem

How to avoid object instantiation when this is not a domain concern?

Design Impact

- *Representation*: Obtain a more fluently readable DSL program by omitting the need for explicit keywords to create instances.
- *Compression*: Eliminate the need to use a specific keyword for creating instances.
- *Absorption*: Avoid the technicalities of creating new instances by absorbing the intention to create a new instance into the DSL.

Instructions

A simple solution is to use a custom method. The method receives parameters which configure several properties of an object, and then this object is returned. A more tricky solution is to use an expression which uses the very same symbol as an existing type. Ruby allows creating a global scope method with the same name as an existing type. The method simply returns a new object. Because of lexical scoping rules, the method can be used in places where method calls are expected (such as in METHOD CHAINING), and thus we can create objects implicitly.

Example

The Boot-DSL uses a SEAMLESS CONSTRUCTOR to instantiate new machines or to update the existing ones. ►Figure 11 shows that normally `Machine` points to the class, but afterwards, it points additionally to a method that returns a new or existing instance.

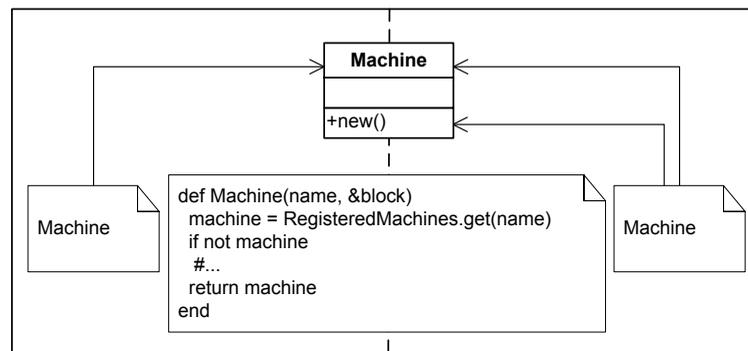


Fig. 11: SEAMLESS CONSTRUCTOR: providing a method with the same name as `Machine` that creates and returns a machine instance.

Implementation Impact

- Lexical scoping determines whether the class or the method is meant when using the symbol.
- Method names written in capital letters violate Ruby conventions and could lead to confusion if not properly explained.

Known Uses

- Ruby
 - No example could be found in the other analyzed DSL, but they are used in the Boot-DSL and SPLCL as shown.
- Python
 - Per default, object instantiation is done by calling the object with attached parentheses. Python allows any variable to point to a method or a method call. This is used in the `should-dsl`: The method `should` actually calls `Should()` to create a new instance (the declaration of the `should` object is contained in `should_dsl/src/should_dsl.py`, Line 224).
- Scala
 - Scala supports SEAMLESS CONSTRUCTOR with two options. Case classes are intended to serve as simple data containers and can be instantiated without the `new` keyword. The other option is to define a method

that returns an instance, because Scala's type system distinguishes between the method call and the object instantiation even if the name is the same for both cases. We could not find this pattern in the analyzed DSLs.

— Smalltalk

— Smalltalk allows defining class level methods that return an instance. This is used in Mondrian to implicitly create nodes that have to be drawn on the screen. Convenience methods exist to automatically create them given any list of elements.

Related Patterns

— ENTITY ALIAS: While SEAMLESS CONSTRUCTOR totally hides the intent to create instances, ENTITY ALIAS makes the intent visible but uses another keyword instead of `new` or similar keywords. The new name can be formed to have an explicit meaning to the domain.

5.2 Entity Alias

Context

A programming language offers extensive functionality through its core-library and several external libraries. However, this functionality is usually expressed in terms which are suitable for the programmer. But these terms may have no or a misleading meaning if used in a DSL. Implementing this functionality again is out of questioning. But still, using appropriate names that fit the domain is crucial for the DSL.

Problem

How to rename built-in classes and external libraries to be consistent with the domain?

Design Impact

— *Representation*: Change the name of existing modules, classes, and methods to better suite the domain by creating aliases to customize the name of an external library or built-in class or method. When working with multiple DSLs, name clashes can be resolved of imported DSLs.

— *Compression*: If existing method or class names are frequently used, compress them with shorter aliases.

— *Generalization*: Provide alias names for already used entities when the internal object hierarchy changes.

Instructions

Ruby is very helpful with this kind of modification: the names of existing modules, classes, and methods can be modified arbitrarily. There are two subpatterns:

— CLASS/MODULE ALIAS: Alias a class or module simply by declaring a variable with the new name.

— METHOD ALIAS: Provide an alias for a method using an interpreter keyword or method.

Example

In the Boot-DSL example, the expression `type EC2` was used to denote the hypervisor for the configured machine. But this is actually METHOD ALIAS for the built-in `include` method. Executing this method includes a module and its methods to the particular machine instance, and it also configures the machine type.

The METHOD ALIAS can be implemented either by using the interpreter keyword `alias` or the method `alias_method`. Both methods first receive the name of the old method and then the name of the new method.

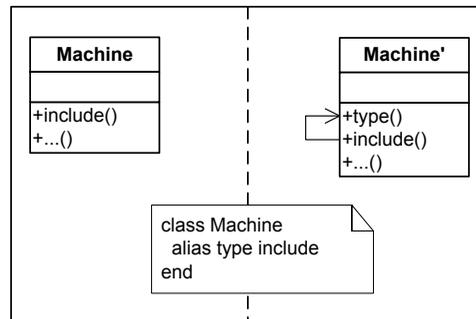


Fig. 12: ENTITY ALIAS/METHOD ALIAS: providing the method name `type` as an alias for `include` inside `Machine`.

CLASS ALIAS is simple to implement. Any symbol in Ruby is just a global symbol registered in the symbol table. An assignment of the new name to the old name is sufficient. For example, to provide an CLASS ALIAS for `Machine`, the following declaration must be executed:

```
1 VirtualMachine = Machine
```

Fig. 13: ENTITY ALIAS/CLASSALIAS: providing `VirtualMachine` as an alias for `Machine`.

Implementation Impact

- Aliasing built-in methods can break compatibility with other libraries.
- Possible namespace conflict when executed globally.
- Sometimes, readability suffers if the used names are not obvious for other developers.

Known Uses

- Ruby
 - Rails has a central generator that is used to provide code stubs for models, controllers, and actions. The generators `initialize` method checks the passed parameter. If a controller is generated, the other passed parameters are calling the `action` method instead of `args`, thereby defining controller actions (the ALIAS METHOD for `initialize` is expressed in `rails-2.3.5/lib/rails_generator/base.rb`, Line 204).
 - Rspec allows to use the name `context` to describe an example group. This method is just an alias for `describe` (the ALIAS METHOD is expressed in `rspec-1.3.0/lib/spec/dsl/main.rb`, Line 30).

— Python

— Python makes building aliases easy since each class or method can be selected from a running program (e.g. using the known method name or using reflection capabilities) and assigned to a new name. This is commonly called decorator in Python, and used for several methods in *Bottle* (*bottle/lib/bottle*, line 543–578).

— Scala

— Similar to Python, there is no direct language support for aliasing methods, but manual re-declaration of methods is straightforward to implement. The *Specs DSL* implements a mechanism that works like an alias. Normally the `should` keyword begins a test case. However, executing `def configure = addToSusVerb("configure")` the call can be extended to `should configure`, where `configure` has the same behavior as `should`. However, `configure` can not be used instead of `should`.

— Smalltalk

— For implementing METHOD ALIAS, new methods can be attached to foreign classes (classes belonging to other packages) using extension packages. One example in *Mondrian* is the `Point>>translatedBy:` method, which has been defined for compatibility reasons as it is often used in layouts.

— CLASS ALIASES can be easily added by simply creating another binding in the globals pool of Smalltalk using the expression `Smalltalk at:#MyClass put:Class`. Another method is to create a proxy class that intercepts every message and redirects it to another object by overriding the method `doesNotUnderstand:` `aMessage`. However, we could not find such mechanisms in the analyzed DSLs.

Related Patterns

— SEAMLESS CONSTRUCTOR: If the intent to create instances should be completely hidden, then use SEAMLESS CONSTRUCTOR instead.

5.3 Operator Expressions

Context

Any domain needs to relate its members to each other: compare them, sort them, and select them out of a bigger set. Naturally, symbols for addition, subtraction and so on come to mind. Many Ruby objects have these operations defined. For example, the `Array` class allows performing the set operations difference and joining with the symbols `-` and `&` respectively. These symbols are perfectly suited to express the semantics of similar domain operations.

Problem

How to use symbols like `+`, `*`, `&` in DSLs?

Instructions

In Ruby and Python, most symbols are just method calls. Defining these methods in the classes that are used to call them will greatly enhance the domain specific look of DSLs.

In Ruby, following operators can be redefined [Flanagan and Matsumoto 2008]:

See [Lutz 2009] for Python and [Wampler and Payne 2009] for Scala operators that are redefinable.

Design Impact

— *Representation*: Add domain-specific operations by using appropriate symbols instead of text for operations or relationship declarations.

— *Compression*: Provide shorter alternatives for frequently used method keywords.

— *Generalization*: Common operators can be used for multiple dispatch of arguments, providing different behavior upon the type of their arguments.

Operator	Operation
!	Boolean NOT
+ -	Unary plus and minus (defined with -@ or +@)
+ -	Addition (or concatenation), subtraction
**	Exponentiation
* / %	Multiplication, division, modulo
& ^ ~	Bitwise AND, OR, XOR, and complement
< >	Bitwise shift-left (or append), bitwise shift-right
< <= >= >	Ordering
== === != =~ !~ <=>	Equality, pattern matching, comparison

Table 1. : Redefinable operators in Ruby (from [Flanagan and Matsumoto 2008]).

Structure and Implementation

The SDP-DSL expresses how applications are installed on machines. Applications can be enrolled on multiple machines, and to express this, the symbol “+” is used. Here is the example expression.

```

1 deploy "Redmine" do
2   enroll "Redmine/Database", :on => app_server + aux_server do
3     #...
4   end
5 end

```

Fig. 14: OPERATOR EXPRESSIONS: Using the “+” symbol to combine two or more machines on which applications are installed.

Since this operator is just a method, it needs to be defined in the `Machine` class, and additionally in the built-in `Array` class in order to chain the returned arrays together. This modification is shown in the following figure.

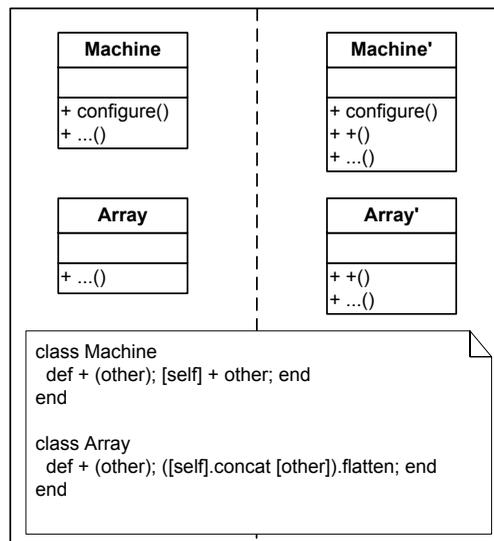


Fig. 15: OPERATOR EXPRESSIONS: adding the operator “+” as a method to the class `Machine` and `Array`.

Implementation Impact

- Operator expressions are just normal method calls, so they need to take care of different argument types and more.
- Some operators are required to be defined on other classes of the application or core library classes too.

Known Uses

- Ruby
 - ActiveRecord uses “<” to add records to associations (defined in activerecord-2.3.5/lib/active_record/associations/association_collection.rb, Line 111–125).
- Python
 - Cryptolang defines 21 operations to express cryptographic algorithms with a very concise syntax. For example, inside the declaration of the blowfish algorithm, the expression `y = y ^ (bs[2] $ S[2])` combines logical-or expressions (“^”) and a substitution (“\$”) by swapping array values [Agosta and Pelosi 2007].
- Scala
 - The Actor DSL uses several methods that are used to fill and query the mailbox object. The method “!” sends a new message, and “!?” sends a message and waits for some milliseconds for a synchronous answers. Messages on the stack can be queried with the “?” method (these methods are defined in scala-lang-2.8.0/src/actors/scala/actors/Channel.scala between Lines 45–108)
- Smalltalk
 - Operators in Smalltalk are binary or unary methods. In Mondrian, “->” is used to construct edges and “@” for points.
 - Glamour uses “->” to select a particular value from a part of the browser. For example, in a browser for software configurations, the selected version is obtained by the expression `#versionsAndBaselines -> #selection`.

Related Patterns

- ENTITY ALIAS/METHOD ALIAS: In some domains, OPERATOR EXPRESSION may seem inappropriate as part of the DSL. To obtain methods instead of the built-in operators, use an alias for defining the particular methods.

6. SUPPORT PATTERNS

Support patterns contribute to the efforts of the other patterns by changing expressions in such a way that they look more separate from the host language.

6.1 Clean Method Calls

Context

In a DSL, it is important to change the syntax for a more natural representation of the domain. Programming languages often require explicit delimiters such as brackets. Except structuring more complex expressions, these symbols usually do not add to a DSL’s readability.

Problem

How to eliminate the clutter of parentheses and other artifacts for your DSL?

Design Impact

- *Representation*: Provide clear expressions that do not resemble their host language but look like keywords of a domain-specific language.

- *Compression*: Remove parentheses, especially for nested method calls, to compress expressions.

Instructions

Ruby allows to drop parentheses in all cases where no lexical ambiguity is present. For example, using any symbol will lead the compiler to look for a constant, but if it is a method (especially with the *Seamless Constructor*), then parentheses must be used. In most cases however, method calls can be dropped from Ruby.

This pattern is not available in Python because parentheses are always required in this language.

Example

Declaring the machines in the Boot DSL uses several methods that set according values. Two expressions are contrasted here. The first one, Line 2 – 6, uses parentheses and the curly brackets for showing blocks. The second form, Line 8 –12, leaves out the parentheses and it uses the `do ... end` block notation.

```
1 # using parentheses
2 Machine("Application Server") {
3   hypervisor(:ec2) {
4     ami("ami-dcf615b5")
5   }
6 }
7 # omitting parentheses
8 Machine "Application Server" do
9   hypervisor :ec2 do
10    ami "ami-dcf615b5"
11  end
12 end
```

Fig. 16: CLEAN METHOD CALLS: removing the parentheses around method calls and using the alternate block syntax simplifies expressions.

Implementation Impact

- Sometimes, explicit knowledge about host language keywords and DSL keywords are more appropriate to understand the DSL completely.

- Lexical scoping determines the meaning of symbols, so using key-value pairs and blocks with curly parentheses can introduce syntactic errors

Known Uses

- Ruby

- In Sinatra, the above mentioned request handlers are typically declared using CLEAN METHOD CALLS (the request handler is implemented in `sinatra-0.9.4/lib/sinatra/base.rb`, Line 752 – 758).

- In RSpec, the example groups are also without using parentheses, which leads to a declarative specification-like language of application behavior (the `describe` method is defined in `rspec-1.3.0/lib/spec/dsl/main.rb`, Line 24 – 29).

- Python

- If a method is called without using parentheses, the Python interpreter treats this as an access to the method object and returns it. Parentheses must be attached in order to call a method – with some rare case implemented by the `should-dsl`. Because built-in operators are defined as methods and cleverly chained, an expression like

`Collection.item_count() |should| be(22)`. can be used. Here, the `should` is actually a method call too, but its parentheses have been skipped (the `should` method is defined in `should-dsl/lib/should_dsl.py`, Line 224).

— Scala

— Scala allows to remove parentheses in some cases. If a method receives no arguments at all, then it can be called without arguments. Especially the Specs DSL uses this capability to write specifications like `should not match("String")`.

— The Apache Camel DSL uses CLEAN METHOD CALLS at several occasions, like in the expression `"app1" throttle (3 per 2 seconds) to ("app2")`⁶ which expresses the intent to only push 3 messages per 2 seconds from the endpoint `app1` to `app2`. Also methods that receive only one argument can be used like this.

— Smalltalk

— This pattern is not supported. Smalltalk parsers are quit rigorous and do not allow for much syntactic freedom. As a result there are parentheses all over the place.

Related Patterns

CLEAN METHOD CALLS is often the foundation of DSL, and therefore combined with all other patterns.

6.2 Custom Return Objects

Context

In a DSL, some methods may require multiple structured objects, such as an array, to be returned to the caller. However, accessing these objects from the outside directly binds the caller to the provided data structures. This is not only difficult to refactor, but also explicitly accessing array values result in an undesired syntax which is tied to language internals.

Problem

How to overcome the limited vocabulary for accessing data in multiple structured objects?

Design Impact

— *Representation*: Utilize keyword access to structured data instead of relying on array positions, providing better readability.

— *Generalization*: Invoke specific semantics in the custom return objects, without changing the caller or the immediate expression environment, thereby providing generalization.

Instructions

Instead of returning a flat list of values, pack the values inside an object and provide accessor methods to them. Depending on the language, objects can be defined on the fly or even proxies for value-accessing method calls can be provided.

Example

Back in the CLEAN METHOD CALLS pattern, a very verbose expression to activate a single feature was used: `activate feature Weighted from product line GPL`. The expression `feature Weighted` is a custom return object because it returns a specific object for the `activate` method. But as explained, this syntax needs two method calls. Alternatively, the following expression and implementation can be used.

The result of the expression `Feature .Weighted` is returned to the `activate` method which works as before.

⁶Taken from the DSL examples at <http://camel.apache.org/scala-dsl-eip.html>.

```

1 def feature(name)
2   case name
3   when Weighted
4     #... modify the object
5     return modified_weighted
6   end
7 end

```

Fig. 17: CUSTOM RETURN OBJECTS: the `feature` method receives an argument which is checked, customized, and finally returned.

In addition to returning concrete objects, anonymous `Struct` objects are available too. They are simple objects with custom values, define in a concise syntax. For example, if a class should be returned with the methods `to` and `from`, the following expression is sufficient.

```

1 Struct.new(:to, :from).new("Return value of 'to'", "Return value of 'from'")

```

Fig. 18: CUSTOM RETURN OBJECTS – definition of an anonymous `Struct` class with two data fields.

Implementation Impact

- Named classes pollute the global symbol table – use anonymous `Struct` objects instead.
- Border between method calls and objects diminishes, which could make it hard to read for programmers.

Known Uses

- Ruby
 - Using ActiveRecord queries, like `clients.find_by_name("Sebastian")`, returns single values or arrays of records, which are typically included in other statements too (the definition of the `find_by*` methods happens dynamically in `activerecord-2.3.5/lib/active_record/base.rb`, Line 1839 – 1961).
- Python
 - With variable length arguments, an appropriate class, and some metaprogramming for defining accessors to the passed data, objects serving as CUSTOM RETURN OBJECTS can be implemented. However, such approach could not be found in the analyzed DSLs.
- Scala
 - The Apache Camel DSL expression `"app1" throttle (3 per 2 seconds) to ("app2")` uses several CUSTOM RETURN OBJECTS. The expression `3 per 2 seconds` will (a) create a `RichInt` type with the value 3, (b) on which the method `per 2` is called to return a `Period` objects, and (c) the `RichInt` and the `Period` objects instantiate a `Frequency` object on which the `seconds` method is called to return the `Period`'s seconds (the `RichInt` type is defined in `apache-camel/components/camel-scala/src/main/scala/org/apache/camel/scala/RichInt.scala`, and the other types and their methods in `Frequency.scala` and `Period.scala` at the same place).
- Smalltalk
 - CUSTOM RETURN OBJECT in Smalltalk is implemented as prototypical objects that receive a dictionary containing key-value pairs of selectors and closures. The selectors denote the method name and the blocks denote the method body. Upon intercepting a message to that object one can consult the dictionary, look up

the matching closure and execute it. This pattern is used for passing Mondrian nodes along in a sequence of messages where each message can change the node resulting from the previous message. Attributes can be created and added on the fly as the nodes have an arbitrary list of attributes.

Related Patterns

— **METHOD CHAINING**: Along a method chain, passing complex structured data that is explicitly accessed couples the implementation of methods very tightly. Using a **CUSTOM RETURN OBJECT** relieves this tension because the data is used with “mnemonic” methods.

7. PRINCIPLES AND THEIR SUPPORTING PATTERNS

This section details how the selected principles can be supported with the patterns. For each principle, the patterns are listed and a larger DSL example is shown to illustrate how the patterns can be applied in combination to design a DSL with respect to the principle.

7.1 Representation

Provide notations and abstractions both required and suitable for the domain.

Supporting Patterns

— **BLOCK SCOPE**: Group related expressions together and syntactically express the hierarchy of objects be indented blocks of code.

— **METHOD CHAINING**: Use long sentences to verbosely express the domain by chaining method calls and their return values together.

— **KEYWORD ARGUMENTS**: Use keyword-value pairs as relationship declarations, to state mappings, or to declaratively express desired system state.

— **ENTITY ALIAS**: Change the name of existing modules, classes, and methods to suite the domain.

— **SEAMLESS CONSTRUCTOR**: Omit explicit keywords to create instances.

— **OPERATOR EXPRESSIONS**: Use symbols instead of texts for operations or relationship declaration.

— **CLEAN METHOD CALLS**: Use method names that look like keywords of the language.

— **CUSTOM RETURN OBJECTS**: Use the returned object in another **METHOD CHAIN** or start a new **BLOCK SCOPE**.

Example

The infrastructure management DSLs are used in ►Listing 19 to configure a virtual machine and to deploy an application. Using this example we demonstrate that the syntax of the DSL can be carefully designed by the above mentioned patterns to closely match a given domain-specific notation.

— *Line 1 – 15*: Using the **BLOCK SCOPE** and **CLEAN METHOD CALLS** patterns, the configuration options of a machine can be grouped and listed using keywords similar to domain-specific notation for such configurations.

— *Line 1*: A machine with the name “Application Server” is created. Using **SEAMLESS CONSTRUCTOR**, this can be written as a declarative domain notation as there is no need to explicitly call `new` to create a new object.

— *Line 2*: The newly created machine is a **EC2** model. By using the **ALIAS ENTITY** pattern we can adopt the domain concept `type` into Ruby as alias for the built-in method `include`. As such, one expression is used both to include the module and to express which type the machine has.

— *Line 3–4*: The machine’s owner and operating system are configured. Keyword-like methods are used to declaratively express properties.

— *Line 5 – 11*: Another BLOCK SCOPE is used to configure the machine's hypervisor. The contained expressions for the hypervisor are logically grouped in accordance with the valid operation domain of the EC2 hypervisor.

— *Line 12 – 14*: More machine options are configured: The hostname, the monitored resources, and that this machine is bootstrapped (installs some basic software packages). Line 13 exemplifies that a cleaner and more domain-friendly syntax can be achieved by using CLEAN METHOD CALLS to remove the parentheses from method calls (which has been used in the preceding lines as well).

— *Line 17 – 20*: Illustrates the SDP-DSL, expressing the intent to deploy the application Redmine⁷ and its dependencies. Similar to Boot-DSL, the BLOCK SCOPE and CLEAN METHOD CALLS are the DSL's building blocks. By using indented blocks, the DSL expresses the natural hierarchy of the domain objects directly with language expressions.

— *Line 19*: The statement uses KEYWORD ARGUMENTS to better match the domain expressions by more descriptively stating that the Redmine/Database package is enrolled on the application server and an auxiliary server. Note the symbol "" is actually a method call and an example of the OPERATOR EXPRESSIONS pattern.

```
1 app_server = Machine "Application Server" do           SEAMLESS CONSTRUCTOR
2   type EC2                                           ALIASING
3   owner "sebastian.guenther@ovgu.de"
4   os :debian
5   hypervisor do                                     BLOCK SCOPE
6     ami "ami-dcf615b5"
7     source "alestic/debian-5.0-lenny-base-2009..."
8     size :m1_small
9     securitygroup "default"
10    private_key "ec2-us-east"
11  end
12  hostname "admantium.com"
13  monitor :cpu, :ram                                CLEAN METHOD CALLS
14  bootstrap!
15 end
16
17 deploy "Redmine" do
18   enroll "Redmine/Database",
19     :on => app_server & aux_server do             KEYWORD ARGUMENTS, OPERATOR EXPRESSIONS
20     #...
21   end
22 end
```

Fig. 19: Representation: Example of the Boot-DSL (Line 1 – 15) to identify and configure a virtual machine, and the CM-DSL (Line 17 – 20) to deploy an application. Several patterns are used to form the DSLs.

⁷<http://www.redmine.org/>

7.2 Absorption

Implicitly absorb domain commonalities to facilitate DSL usage.

Pattern Support

ABSORPTION depends on the application to have a rich amount of configured properties. The properties are then implicitly contained in language expressions, and the following patterns can be used:

- BLOCK SCOPE: The scope's context absorbs the most relevant execution information, especially if it serves as a closure capturing the surrounding context and using them later to infer contextual information.
- KEYWORD ARGUMENTS: Absorb additional information as optional keyword arguments.
- SEAMLESS CONSTRUCTOR: Absorb the intention to create a new instance.

Examples

In order to explain this design principle, the infrastructure management DSLs are used again. ►Figure 20 shows how the application server's hypervisor is configured.

1	<code>Machine "Application Server" do</code>	SEAMLESS CONSTRUCTOR
2	<code> type EC2</code>	
3	<code> hypervisor do</code>	BLOCK SCOPE
4	<code> ami "ami-dcf615b5"</code>	
5	<code> source 'alestic/debian-5.0-lenny-base-2009...'</code>	
6	<code> size :m1_large, :limit => :costs</code>	KEYWORD ARGUMENTS
7	<code> securitygroup "default"</code>	
8	<code> private_key "ec2-us-east"</code>	
9	<code> end</code>	
10	<code>end</code>	

Fig. 20: Absorption: Boot-DSL and patterns that provide the absorption of context.

- *Line 1*: Expresses the customization of the application server. Whether a new machine is created or whether an existing machine is reconfigured is absorbed with the SEAMLESS CONSTRUCTOR.
- *Line 3 – 9*: All expressions within the hypervisor block are executed against the specific hypervisor EC2. Using the BLOCK SCOPE pattern, this dependency is absorbed into the inner hypervisor block.
- *Line 6*: Expresses to use a large size for the machine, and additionally that the total operating costs as a limit for the (re)sizing. If the costs are calculated as being too high, the small size will be used. The expression `:limit => :costs` is optional absorbed into the DSL by KEYWORD ARGUMENTS.

7.3 Compression

Provide a concise language that is sufficiently descriptive for the domain experts.

Pattern Support

- BLOCK SCOPE: Remove fully-qualified method calls and provide a named context for code execution.
- METHOD CHAINING: Compress several method calls into a chain that operate on each other's return values.
- KEYWORD ARGUMENTS: Compress several singular method calls to a set of key-value pairs passed to one method.
- ENTITY ALIAS: Compress existing method or class names.

- OPERATOR EXPRESSIONS: Provide shorter alternatives for method keywords.
- CLEAN METHOD CALLS: Compress expressions by removing parentheses, especially for nested method calls.

Example

In the following ►Figure 21, the GPL example is used to show three different variant how the root feature can be configured.

```

1 gpl = Feature.new("GPL")
2 gpl.tree_position(:root)
3 gpl.set_subfeatures(:Type, :Weight, :Search, :Algorithms)
4 gpl.add_requirements(self.name, "all :Type, :Weight, :Search, :Algorithms")
5
6 gpl = Feature.configure do                                BLOCK SCOPE
7   name :GPL
8   root
9   subfeatures :Type, :Weight, :Search, :Algorithms
10  requires :GPL => "all :Type, :Weight, :Search, :Algorithms"  KEYWORD ARGUMENTS
11 end
12
13 configure_feature :GPL, as_root do                       CLEAN METHOD CALLS
14   subfeatures :Type, :Weight, :Search, :Algorithms
15   requires self => all(subfeatures)                       METHOD CHAINING
16 end

```

Fig. 21: *Compression*: Showing three different version for modeling features. The tree examples vary in the amount of compression that has been obtained. The first and last example depicts two extremes: the first one uses common object-oriented expressions and requires 181 characters, the last one uses 125 characters. A middle ground is found in the second example where compactness and readability are more balanced.

- *Line 1 – 4*: Normal object-oriented expressions, including method calls with parentheses. These expressions require 181 characters.

- *Line 6 – 11*: The current form of our SPLC, with a total amount of 162 characters. The BLOCK SCOPE pattern is used to eliminate the fully qualified method calls as opposed to the first variant. Also, CLEAN METHOD CALLS eliminate the parentheses. Finally Line 10 uses the KEYWORD ARGUMENTS pattern to provide a verbose but readable form of constraints.

- *Line 13 – 16*: A variant with high *compression* using only 125 characters. The name and position properties are specified as arguments to a method using CLEAN METHOD CALLS (Line 13). This argument notation is similar to KEYWORD ARGUMENTS, but without using the keywords and thus allowing a very compressed form. In the BLOCK SCOPE body, the subfeatures are configured as before, but this time combining KEYWORD ARGUMENTS with METHOD CHAINING in Line 16. The constraint rule `all(subfeatures)` uses internal information about the feature to greatly compresses the expression.

It is difficult to argue what amount of *compression* is sufficient or useful for a particular DSL. For example, the first and second alternatives require similar amount of text, but the second version does not use symbols or method calls that are “typical” for programming languages. Although the third version is the most compressed form, the meaning of `self` and the method call `all(subfeatures)` implies knowledge about the host language – people not familiar with Ruby might understand the second alternative better.

7.4 Generalization

Reduce the amount of concepts by replacing a group of more specific cases with a common case.

Pattern Support

Generalization is supported by providing entities and methods that serve as anchors for later specialization refinements. Therefore the following patterns are used:

- KEYWORD ARGUMENTS: Changing semantics by using multiple dispatch on arguments.
- ENTITY ALIAS: Provide alias names for used entities when the hierarchy changes.
- OPERATOR EXPRESSIONS: Common operators can be used for multiple dispatch of arguments, providing different behavior upon the type of their arguments.
- CUSTOM RETURN OBJECTS: The returned object is invisible to the caller. Semantic changes are confined to the objects without having to change the surrounding expressions.

Example

In the following ►Figure 22, example expressions from the SPLCL are shown.

```
1 GPL.configure :root => true,                                KEYWORD ARGUMENTS
2               :subfeatures => [:Type, :Weight, :Search, :Algorithms],
3               #...
4
5 variant.create do
6   name "SimpleVariant"
7   activate feature(Weighted) from productline GPL        METHOD CHAINING, CUSTOM RETURN OBJECTS
8   #...
9 end
```

Fig. 22: *Generalization*: using the SPLCL to configure a feature and then to create a variant. Patterns related to modifying and implementing methods are used to generalize expressions of the DSL.

- *Line 1 – 2*: An alternative form of configuring a product line that uses KEYWORD ARGUMENTS. The used keywords have the same meaning as the method expressions explained before. The expression has a high *generalization* because it focuses on the essential properties of software product lines, yet it is simple to extend with specialized behavior using new keywords.
- *Line 5 – 9*: This expression creates an actual variant of a product line.
 - *Line 6*: Configures the name of the particular product variant that is registered with the product lines.
 - *Line 7*: A complex expression combining METHOD CHAINING and CUSTOM RETURN OBJECTS. It expresses that this variant should use the product line GPL. The expression `feature(Weighted)` is return a custom object that selects the `Weighted` feature from the global namespace. This behavior can be changed transparently by the implementation without changing the language's syntax, thus facilitating *generalization*.

8. DISCUSSION OF PATTERN USAGE

This section details our experience using the patterns and how they collaboratively support the principles.

8.1 Form Patterns

The form patterns are BLOCK SCOPE, METHOD CHAINING, KEYWORD ARGUMENTS.

BLOCK SCOPE is a very dominant pattern. It's hard to find a Ruby DSL that does not use this pattern. The impact on readability achieved with BLOCK SCOPE is fundamental: the `do . . . end` notation reads very natural and fits into several domains. It allows both to syntactically group related expressions and to express the natural hierarchy of objects. And to give the full potential of this pattern: even its semantic capabilities enable many modifications, including deferred evaluation, composition, and execution at any scope of the program. BLOCK SCOPE can be used for the *representation*, *absorption*, and *compression* principle, but is of limited use for *generalization*.

METHOD CHAINING and KEYWORD ARGUMENTS are both an alternative and an addition to the BLOCK SCOPE. These patterns are not concerned with the syntactical grouping of related statements, but they express more complex statements. Since method calls can be used without using the fully qualified name and – using CLEAN METHOD CALLS – can also drop all parentheses, chains of methods that look like language keywords, can be used. This enables more self explanatory statements that can be used for a clear representation of the domain.

The close connection between the form patterns are the following: A block allows to form expressions whose implementation is not available at the place they are specified, but at the place they are executed. Thus, using METHOD CHAINING and KEYWORD ARGUMENTS with CLEAN METHOD CALLS in the context of BLOCK SCOPE allows very free from languages. The overall result are verbose expressions used for the *representation* and also *absorptions*. Depending on the desired verbosity, *compression* can be limited by this combination. *Generalization* is strongly supported since multiple argument dispatch is possible with the KEYWORD ARGUMENTS and the METHOD CHAINING.

8.2 Modification Patterns

The modifications patterns are SEAMLESS CONSTRUCTOR, ENTITY ALIAS, and OPERATOR EXPRESSIONS.

The SEAMLESS CONSTRUCTOR removes the need to use the `new` operator in order to create instances. The solution is to define a method with the same name as the class, and to use the class inside the method statement. The method can receive multiple arguments via KEYWORD ARGUMENTS that are used to customize the implicitly created instance. *Representation*, *absorption*, and *generalization* are supported.

The ALIAS ENTITY pattern can be used in a very similar fashion. Instead of eliminating a specific method call, a viable alternative can be provided. ALIAS ENTITY can be used to customize the whole Ruby programming language, since all classes and methods, even the built-in ones such as `Array` and `String`, can be modified to suit the domain. All design principles are supported, and if the modification of built-in methods is considered, *generalization* is especially well supported.

The final pattern here is OPERATOR EXPRESSIONS. Ruby is a language with a very minimal set of interpreter expressions. The creation of modules, classes, and methods are covered by this, and assignments. But beyond that, a rich set of operators can be defined since they are just method calls on objects. Domains that require such symbolic expressions benefit from the improved *representation* and *compression*.

8.3 Support Patterns

Support patterns are CLEAN METHOD CALLS and CUSTOM RETURN OBJECTS.

CLEAN METHOD CALLS is an extensively used support pattern. Although its intent is simple and its application trivial, it has a strong impact on *representation* and *compression*. Imagine to use METHOD CHAINING while keeping parentheses: only because Ruby allows to eliminate most parentheses, expressions that look like using language keywords can be build. CLEAN METHOD CALLS are primary supporting *representation*, and *compression* to the degree of removing parentheses.

Finally, CUSTOM RETURN OBJECTS is a support pattern for BLOCK SCOPE, METHOD CHAINING, and KEYWORD ARGUMENTS. The primary use case is to simply return an object and access its data fields. Additionally, it can be used inside a method chain too, enabling multiple dispatch arguments. Or it receives a block and can start a BLOCK SCOPE. It is also an alternative for METHOD CHAINING since the returned objects can also define arbitrary data and methods specifically in its scope without polluting the global namespace. Therefore, CUSTOM RETURN OBJECTS supports the *absorption*, *compression*, and *generalization* (it's easy to customize the objects behavior) principle. Its effects on *representation* are usually only achieved in combination with other form patterns.

8.4 Summary

The thoughtful combination of patterns allows DSL developers to plan and discuss a principles-based implementation. Each pattern group has a distinctive role. The form patterns provide the building blocks of a DSL. They characterize the overall *representation* of the domain as either (i) blocks of hierarchically structured expressions, (ii) grouped expression, or (iii) descriptive and complex expressions. Thereby the form patterns use the support patterns to provide method calls without using parentheses or to use objects as an anchor for further block or complex expressions. The modification patterns support adaptation and extension of the available constructs of the domain, which greatly enhances the available forms.

In the pattern's description and the discussion, two facts emerged. First, a pattern supports multiple design principles. It is seldom feasible to use a pattern in isolation to support a single principle. More so, the principles itself are connected and strengthen each other. Second, the patterns have close relationships to each other and are typically used in combinations. For example to use METHOD CHAINING without CLEAN METHOD CALLS would severely impair readability of the expressions. This interleaving means design considerations commonly lead to a closely coupled set of patterns.

9. RELATED WORK

General work on Domain-Specific Languages and their development starts with [Spinellis 2001; Mernik et al. 2005]. These early works are proposing a complete design process including patterns. The patterns however are very coarse grained, like using Unix pipes and filters to chain command or the usage of an interpreter. Our work has more focus on dynamic languages and syntactical language modifications according to design principles.

For graphical modeling languages, criteria how to select different visualization forms for models are presented [Costagliola et al. 2002; Moody 2009]. However, the close focus on graphical notation does not allow a generalization for usage in designing textual languages. There are some empirical studies regarding what concepts to include or not to include in a language [Pane 2004; Høst and Østvold 2008], but these approaches lack an overall conceptual framework while our pattern catalog is the foundation of a principled design method.

In related work in the context of DSLs and dynamic programming languages, few authors try to structure the techniques they use for the host language's modification. They mix language characteristics with concrete mechanisms and design decisions to explain their approach. For example, [Cunningham 2008] explains a DSL for surveys using several Ruby metaprogramming mechanisms. [Dinkelaker and Mezini 2008] explains a Groovy-based DSL for aspect-oriented programming, using an internal interpreter object. For Python, DSLs to secure and restrict the access to resources [Cannon and Wohlstadter 2007] or for cryptography are available [Agosta and Pelosi 2007]. Also, most of these works are focused on semantic extensions of the used host language, considerations of the DSL's syntax are only seldom discussed, like in [Agosta and Pelosi 2007] that uses operator expression for the DSL. Exceptional is a book about the Scala programming language [Wampler and Payne 2009], where several implementation techniques are presented, including different syntaxes for calling methods. Some of the presented techniques look like METHOD CHAINING and CUSTOM RETURN OBJECTS. In yet unpublished work (<http://martinfowler.com/dslwip>), Fowler provides some interesting techniques for implementing internal DSLs that go beyond a particular dynamic language. It is interesting to further check our found patterns with other dynamic languages or even static languages to see how they fare with respect to generalization.

Looking explicitly for pattern catalogs in the context of DSL, some examples can be found. In [Spinellis 2001], eight general design principles are considered grouped into creational, structural, and behavioral. Creational patterns for example elaborate the principle availability of using a given language for specialization or restriction of its host, as well as using the string processing capabilities of a given language for very simple DSLs. Behavioral patterns explain how to link the output of one DSL with another one using Unix pipes. We see these patterns as providing general implementation schemes while our patterns are detailing how to actually change the language. Another paper explains patterns that govern the principal DSL development process (is the DSL created from scratch or is it based on the “language” of an existing system?), the concrete syntax style (is it a textual or a graphical DSL?), and whether to build an external and internal DSL. The patterns explore the general available design space for DSLs and their influence on each other [Zdun and Strembeck 2009]. Another paper explains patterns for different language engineering concerns. Some patterns express the need to start a language with a minimal feature set, growing it over the time with more functionality. Syntax is of great importance and should be put first as design criteria. Once the language is bigger, modules in independent files should be used to facilitate maintaining and continuing development with multiple developers. Also, a standard library should be given to support the language’s core functionality, and explicit extension points can help to focus the language’s evolution [Haase 2007].

10. SUMMARY

In this paper we presented a rich pattern catalog describing a principled approach for designing internal DSLs. Based on the defining characteristics of a DSL the following four principles have been presented: *representation*, *absorption*, *compression*, and *generalization*. We explained each design principle with motivation, forces, consequences, and showed an example in which several patterns are used. These principles are concerned with the DSL’s syntax, and thus we explained patterns that enable or modify syntactic constructs. We identified three groups of patterns: Form, modification, and support. The form patterns provide the language’s overall layout, a broad choice between closely grouped singular statements that can express the natural hierarchy in the domain, or descriptive and complex statements that look like language keywords. The modification patterns change the existing language by aliasing entities and module. Finally, the support patterns provide additional syntactic changes that can be combined with the other patterns. We explained each patterns with the classical pattern description form, including the implementation and the known uses. The patterns were illustrated with Ruby, but we also provided known uses for Python.

The final discussion contrasted the patterns from each other and explains how they can be connected. Both the given examples and the known uses show that the explained patterns are ideally used in combination with each other, and allow to build DSLs with a distinct form that do not resemble their host-language’s object-oriented nature.

Acknowledgements

We thank our shepherd Uwe Zdun for helpful comments on an earlier draft of this paper. We also wish to thank Andy Kellens for sharing his insights of the Smalltalk language, and Alexandre Bergel for his input on the presence and use of these patterns in Smalltalk DSLs.

Sebastian Günther works with the Very Large Business Applications Lab, School of Computer Science, at the Otto-von-Guericke University of Magdeburg. The Very Large Business Applications Lab is supported by SAP AG. Thomas Cleenewerk works with the Language Engineering Lab at the Vrije Universiteit Brussel.

REFERENCES

AGOSTA, G. AND PELOSI, G. 2007. A Domain Specific Language for Cryptography. In *Proceedings of the Forum on specification and Design Languages (FDL)*. ECSI, 159–164.

- ARNOLD, B. R. T., DEURSEN, A. V., AND RES, M. 1995. Algebraic Specification of a Language for describing Financial Products. In *ICSE-17 Workshop on Formal Methods Application in Software Engineering*. IEEE, 6–13.
- BATORY, D., SARVELA, J. N., AND RAUSCHMAYER, A. 2003. Scaling Step-Wise Refinement. In *Proceedings of the 25th International Conference on Software Engineering (ICSE)*. IEEE Computer Society, Washington, 187–197.
- BENNETT, M., BORGES, R., HAVELUND, K., INGHAM, M., AND WAGNER, D. 2008. Development of a Prototype Domain-Specific Language for Monitor and Control Systems. In *Aerospace Conference*. IEEE Computer Society, 1–18.
- BENTLEY, J. 1986. Programming Pearls: Little Languages. *Communications of the ACM* 29, 8, 711–721.
- BLACK, A. P., DUCASSE, S., NIERSTRASZ, O., AND POLLET, D. 2009a. *Pharo by Example*. Square Bracket Associates.
- BLACK, A. P., DUCASSE, S., NIERSTRASZ, O., AND POLLET, D. 2009b. *Squeak by Example* 1st revised Ed. Square Bracket Associates.
- BOOCH, G. 1994. *Object-Oriented Analysis and Design with Applications* 2nd Ed. Addison Wesley Longman, Redwood City, CA, USA.
- BRUCE, D. 1997. What makes a good Domain-Specific Language? APOSTLE, and its Approach to Parallel Discrete Event Simulation. In *First ACM SIGPLAN Workshop on Domain-Specific Languages (DSL)*. University of Illinois Computer Science, Paris, France, 17–35.
- CANNON, B. AND WOHLSTADTER, E. 2007. Controlling Access to Resources within the Python Interpreter. *Proceedings of the Second EEECE*.
- COPLIEN, J. O. 1999. *Multi-paradigm design for C++*. Addison-Wesley, Boston, San Francisco, et al.
- CORREIA, F. F., FERREIRA, H. S., FLORES, N., AND AGUIAR, A. Patterns for Consistent Software Documentation. In *Proceedings of the 16th Conference for Pattern Languages of Programs (PloP)*.
- COSTAGLIOLA, G., DELUCIA, A., OREFICE, S., AND POLESE, G. 2002. A classification framework to support the design of visual languages. *Journal of Visual Languages & Computing* 13, 6, 573–600.
- CUNNINGHAM, H. C. 2008. A Little Language for Surveys: Constructing an Internal DSL in Ruby. In *Proceedings of the 46th Annual Southeast Regional Conference (ACM-SE)*. ACM, New York, 282–287.
- CZARNECKI, K. AND EISENECKER, U. W. 2000. *Generative Programming: Methods, Tools, and Applications*. Addison-Wesley, Boston, San Francisco et al.
- DINKELAKER, T. AND MEZINI, M. 2008. Dynamically Linked Domain-Specific Extensions for Advice Languages. In *Proceedings of the 2008 AOSD Workshop on Domain-Specific Aspect Languages (DSAL)*. ACM, New York, 1–7.
- FLANAGAN, D. AND MATSUMOTO, Y. 2008. *The Ruby Programming Language*. O'Reilly Media, Sebastopol, USA.
- FOOTE B, J. Y. 1998. Metadata and Active Object Models. In *Proceedings of the 5th Conference on Pattern Languages of Programs (PLOP)*. Washington University Department of Computer Science, Washington, USA.
- GAMMA, E., HELM, R., JOHNSON, R., AND VLISSIDES, J. 1997. *Design Patterns - Elements of Reusable Object-Oriented Software* 10th Ed. Addison-Wesley, Reading, Harlow et al.
- GROOTE, J. F., VAN VLIJMEN, S. F. M., AND KOORN, J. W. C. 1995. The Safety Guaranteeing System at Station Hoorn-Kersenboogerd. In *Proceedings of the Tenth Annual Conference on Computer Assurance Systems Integrity, Software Safety and Process Security (COMPASS 95)*. IEEE, 57–68.
- GUERRA, E., SOUZA, J., AND FERNANDES, C. 2009. A Pattern Language for Metadata-based Frameworks. In *Proceedings of the 16th Conference on Pattern Languages of Programs (PLOP)*. ACM.
- GÜNTHER, S. 2009. Engineering Domain-Specific Languages with Ruby. In 3. *Workshop des Centers for Very Large Business Applications (CVLBA)*, H.-K. Arndt and H. Krömer, Eds. Shaker, Aachen, 11–21.
- HAASE, A. 2007. Patterns for the Definition of Programming Languages. In *Proceedings of 12th European Conference on Pattern Languages of Programs (EuroPLOP)*.
- HØST, E. W. AND ØSTVOLD, B. M. 2008. The java programmer's phrase book. In *Software Language Engineering. First International Conference on Software Language Engineering (SLE), 5452*, 322–341.
- HUDAK, P. 1998. Modular Domain Specific Languages and Tools. In *Proceedings of the Fifth International Conference on Software Reuse (ICSR)*, P. Davenbu and J. Poulin, Eds. IEEE, Washington, 134–142.
- KANG, K., COHEN, S., HESS, J., NOVAK, W., AND PETERSON, A. 1990. Feature-Oriented Domain Analysis (FODA) Feasibility Study. Tech. Rep. CMU/SEI-90-TR-21, Software Engineering Institute, Carnegie Mellon University, USA.
- LADD, D. A. AND RAMMING, J. C. 1994. Two Application Languages in Software Production. In *Proceedings of the USENIX Very High Level Languages Symposium (VHLLS)*. USENIX Association, Berkeley, CA, USA, 10–18.
- LANDIN, P. J. 1966. The Next 700 Programming Languages. *Communications of the ACM* 9, 3, 157–166.
- LATRY, F., MERCADAL, J., AND CONSEL, C. 2007. Staging Telephony Service Creation: A Language Approach. In *Proceedings of the 1st international conference on Principles, systems and applications of IP telecommunications (IPTComm)*. ACM, New York, 99–110.
- LOPEZ-HERREJON, R. E. AND BATORY, D. 2001. A Standard Problem for Evaluating Productline Methodologies. In *Proceedings of the 3rd International Conference on Generative and Component-Based Software Engineering (GCSE)*, J. Bosch, Ed. Lecture Notes in Computer Science Series, vol. 2186. Springer-Verlag, London, 10–24.
- LUTZ, M. 2009. *Learning Python* 4th Ed. O'Reilly Media, Sebastopol.

- MCKEEMAN, W. M. 1976. Programming language design. In *Compiler Construction, An Advanced Course, 2nd ed.* Springer-Verlag, London, UK, 514–524.
- MERNIK, M., HEERING, J., AND SLOANE, A. M. 2005. When and How to Develop Domain-Specific Languages. *ACM Computing Survey* 37, 4, 316–344.
- MOODY, D. L. 2009. The "Physics" of Notations: Toward a Scientific Basis for Constructing Visual Notations in Software Engineering. *IEEE Transactions on Software Engineering* 35, 756–779.
- MUNNELLY, J. AND CLARKE, S. 2008. A Domain-Specific Language for Ubiquitous Healthcare. In *Proceedings of the 3rd International Conference on Pervasive Computing and Applications (ICSPA)*. IEEE, 757–762.
- ODERSKY, M., SPOON, L., AND VENNERS, B. 2008. *Programming in Scala*. Artima Press, Mountain View, California, USA.
- PANE, J. F. 2004. Dissertation. Ph.D. thesis, Carnegie Mellon University, Computer Science Department.
- PERROTTA, P. 2010. *Metaprogramming Ruby*. The Pragmatic Bookshelf, Raleigh.
- SEBESTA, R. W. 1999. *Concepts of Programming Languages*. Addison-Wesley, Reading, Harlow et al.
- SHARP, A. 1997. *Smalltalk by Example*. McGraw-Hill Publishing Co.
- SPIEWAK, D. AND ZHAO, T. 2009. ScalaQL: Language-Integrated Database Queries for Scala. In *Proceedings of the 2nd International Conference on Software Language Engineering (SLE)*, M. van den Brand, D. Gašević, and J. Gray, Eds. Lecture Notes in Computer Science Series, vol. 5969. Springer-Verlag, Berlin, Heidelberg, 154–163.
- SPINELLIS, D. 2001. Notable Design Patterns for Domain-Specific Languages. *Journal of Systems and Software* 56, 1, 91–99.
- SUMMERFIELD, M. 2010. *Programming in Python 3: A Complete Introduction to the Python Programming Language* 2nd Ed. Addison-Wesley, Upper Saddle River, Boston et al.
- TANTER, É. 2008. Contextual Values. In *Proceedings of the 2008 Symposium on Dynamic Languages (DLS)*. ACM.
- THIBAUT, S., MARLET, R., AND CONSEL, C. 1997. A Domain-Specific Language for Video Device Drivers: from Design to Implementation. 11–26.
- THOMAS, D., FOWLER, C., AND HUNT, A. 2009. *Programming Ruby 1.9 - The Pragmatic Programmers' Guide*. The Pragmatic Bookshelf, Raleigh, USA.
- VAN DEURSEN, A., KLINT, P., AND VISSER, J. 2000. Domain-Specific Languages: An Annotated Bibliography. *ACM SIGPLAN Notices* 35, 6, pp. 26–36.
- WAMPLER, D. AND PAYNE, A. 2009. *Programming Scala*. O'Reilly Media, Sebastopol, USA.
- WEINBERG, G. M. 1971. *The Philosophy of Programming Languages*. John Wiley & Sons, New York.
- WITHEY, J. 1996. Investment Analysis of Software Assets for Product Lines. Technical Report CMU/SEI96-TR-010, Software Engineering Institute, Carnegie Mellon University, USA.
- ZDUN, U. AND STREMBECK, M. 2009. Reusable Architectural Decisions for DSL Design: Foundational Decisions in DSL Projects. In *Proceedings of the 14th Annual European Conference on Pattern Languages of Programming (EuroPLoP)*, A. Kelly and M. Weiss, Eds. CEUR, RWTH Aachen, Aachen, Germany.