

When and How to Develop Domain-Specific Languages

MARJAN MERNIK

University of Maribor

JAN HEERING

CWI

AND

ANTHONY M. SLOANE

Macquarie University

Domain-specific languages (DSLs) are languages tailored to a specific application domain. They offer substantial gains in expressiveness and ease of use compared with general-purpose programming languages in their domain of application. DSL development is hard, requiring both domain knowledge and language development expertise. Few people have both. Not surprisingly, the decision to develop a DSL is often postponed indefinitely, if considered at all, and most DSLs never get beyond the application library stage.

Although many articles have been written on the development of particular DSLs, there is very limited literature on DSL development methodologies and many questions remain regarding when and how to develop a DSL. To aid the DSL developer, we identify patterns in the decision, analysis, design, and implementation phases of DSL development. Our patterns improve and extend earlier work on DSL design patterns. We also discuss domain analysis tools and language development systems that may help to speed up DSL development. Finally, we present a number of open problems.

Categories and Subject Descriptors: D.3.2 [**Programming Languages**]: Language Classifications—*Specialized Application Languages*

General Terms: Design, Languages, Performance

Additional Key Words and Phrases: Domain-specific language, application language, domain analysis, language development system

Authors' addresses: M. Mernik, Faculty of Electrical Engineering and Computer Science, University of Maribor, Smetanova 17, 2000 Maribor, Slovenia; email: marjan.mernik@uni-mb.si; J. Heering, Department of Software Engineering, CWI, Kruislaan 413, 1098 SJ Amsterdam, The Netherlands; email: Jan.Heering@cw.nl; A.M. Sloane, Department of Computing, Macquarie University, Sydney, NSW 2109, Australia; email: asloane@ics.mq.edu.au.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

©2005 ACM 0360-0300/05/1200-0316 \$5.00

1. INTRODUCTION

1.1. General

Many computer languages are *domain-specific* rather than general purpose. Domain-specific languages (DSLs) are also called application-oriented [Sammet 1969], special purpose [Wexelblat 1981, p. xix], specialized [Bergin and Gibson 1996, p. 17], task-specific [Nardi 1993], or application [Martin 1985] languages. So-called *fourth-generation languages* (4GLs) [Martin 1985] are usually DSLs for database applications. *Little languages* are small DSLs that do not include many features found in general-purpose programming languages (GPLs) [Bentley 1986, p. 715].

DSLs trade generality for expressiveness in a limited domain. By providing notations and constructs tailored toward a particular application domain, they offer substantial gains in expressiveness and ease of use compared with GPLs for the domain in question, with corresponding gains in productivity and reduced maintenance costs. Also, by reducing the amount of domain and programming expertise needed, DSLs open up their application domain to a larger group of software developers compared to GPLs. Some widely used DSLs with their application domains are listed in Table I. The third column gives the language level of each DSL as given in Jones [1996]. Language level is related to productivity as shown in Table II, also from Jones [1996]. Apart from these examples, the benefits of DSLs have often been observed in practice and are supported by quantitative results such as those reported in Herndon and Berzins [1988]; Batory et al. [1994]; Jones [1996]; Kieburtz et al. [1996]; and Gray and Karsai [2003], but their quantitative validation in general as well as in particular cases, is hard and an important open problem. Therefore, the treatment of DSL development in this article will be largely qualitative.

The use of DSLs is by no means new. APT, a DSL for programming numerically-controlled machine tools, was developed in 1957–1958 [Ross 1981]. BNF,

the well-known syntax specification formalism, dates back to 1959 [Backus 1960]. Domain-specific visual languages (DSVLs), such as visual languages for hardware description and protocol specification, are important but beyond the scope of this survey.

We will not give a definition of what constitutes an application domain and what does not. Some consider Cobol to be a DSL for business applications, but others would argue this is pushing the notion of application domain too far. Leaving matters of definition aside, it is natural to think of DSLs in terms of a gradual scale with very specialized DSLs such as BNF on the left and GPLs such as C++ on the right. (The language level measure of Jones [1996] is one attempt to quantify this scale.) On this scale, Cobol would be somewhere between BNF and C++ but much closer to the latter. Similarly, it is hard to tell if command languages like the Unix shell or scripting languages like Tcl are DSLs. Clearly, domain-specificity is a matter of degree.

In combination with an *application library*, any GPL can act as a DSL. The library's Application Programmers Interface (API) constitutes a domain-specific vocabulary of class, method, and function names that becomes available by object creation and method invocation to any GPL program using the library. This being the case, why were DSLs developed in the first place? Simply because they can offer domain-specificity in better ways.

- Appropriate or established *domain-specific notations* are usually beyond the limited user-definable operator notation offered by GPLs. A DSL offers appropriate domain-specific notations from the start. Their importance should not be underestimated as they are directly related to the productivity improvement associated with the use of DSLs.
- Appropriate *domain-specific constructs and abstractions* cannot always be mapped in a straightforward way to functions or objects that can be put in a library. Traversals and error handling are typical examples [Bonachea et al. 1999; Gray and Karsai 2003; Bruntink

Table I. Some Widely Used Domain-Specific Languages

DSL	Application Domain	Level	
BNF	Syntax specification	n.a.	
Excel	Spreadsheets	57	(version 5)
HTML	Hypertext web pages	22	(version 3.0)
L ^A T _E X	Typesetting	n.a.	
Make	Software building	15	
MATLAB	Technical computing	n.a.	
SQL	Database queries	25	
VHDL	Hardware design	17	
Java	General-purpose	6	(comparison only)

Table II. Language Level vs. Productivity as Measured in Function Points (FP)

Level	Productivity Average per Staff Month (FP)
1–3	5–10
4–8	10–20
9–15	16–23
16–23	15–30
24–55	30–50
> 55	40–100

et al. 2005]. A GPL in combination with an application library can only express these constructs indirectly or in an awkward way. Again, a DSL would incorporate domain-specific constructs from the start.

- Use of a DSL offers possibilities for *analysis, verification, optimization, parallelization, and transformation* in terms of DSL constructs that would be much harder or unfeasible if a GPL had been used because the GPL source code patterns involved are too complex or not well defined.
- Unlike GPLs, DSLs *need not be executable*. There is no agreement on this in the DSL literature. For instance, the importance of nonexecutable DSLs is emphasized in Wile [2001], but DSLs are required to be executable in van Deursen et al. [2000]. We discuss DSL executability in Section 1.2.

Despite their shortcomings, application libraries are formidable competitors to DSLs. It is probably fair to say that most DSLs never get beyond the application library stage. These are sometimes called *domain-specific embedded languages* (DSELs) [Hudak 1996]. Even with improved DSL development tools, application libraries will remain the most

cost effective solution in many cases, the more so since the advent of component technologies such as COM and CORBA [Szyperki 2002] has further complicated the relative merits of DSLs and application libraries. For instance, Microsoft Excel’s macro language is a DSL for spreadsheet applications which adds programmability to Excel’s fundamental interactive mode. Using COM, Excel’s implementation has been restructured into an application library of COM components, thereby opening it up to GPLs such as C++, Java, and Basic which can access it through its COM interfaces. This process of componentization is called *automation* [Chappell 1996]. Unlike the Excel macro language, which by its very nature is limited to Excel functionality, GPLs are not. They can be used to write applications transcending Excel’s boundaries by using components from other automated programs and COM libraries in addition to components from Excel itself.

In the remainder of this section, we discuss DSL executability (Section 1.2), DSLs as enablers of reuse (Section 1.3), the scope of this article (Section 1.4), and DSL literature (Section 1.5).

1.2. Executability of DSLs

DSLs are executable in various ways and to various degrees even to the point of being nonexecutable. Accordingly, depending on the character of the DSL in question, the corresponding programs are often more properly called specifications, definitions, or descriptions. We identify some points on the DSL executability scale.

- DSL with well-defined execution semantics (e.g., Excel macro language, HTML).

- Input language of an *application generator* [Cleaveland 1988; Smaragdakis and Batory 2000]. Examples are ATMOL [van Engelen 2001], a DSL for atmospheric modeling, and Hancock [Bonachea et al. 1999], a DSL for customer profiling. Such languages are also executable, but they usually have a more declarative character and less well-defined execution semantics as far as the details of the generated applications are concerned. The application generator is a compiler for the DSL in question.
- DSL not primarily meant to be executable but nevertheless useful for application generation. The syntax specification formalism BNF is an example of a DSL with a purely declarative character that can also act as an input language for a parser generator.
- DSL not meant to be executable. Examples are domain-specific data structure representations [Wile 2001]. Just like their executable relatives, such non-executable DSLs may benefit from various kinds of tool support such as specialized editors, prettyprinters, consistency checkers, analyzers, and visualizers.

1.3. DSLs as Enablers of Reuse

The importance of DSLs can also be appreciated from the wider perspective of the construction of large software systems. In this context, the primary contribution of DSLs is to enable reuse of software artifacts [Biggerstaff 1998]. Among the types of artifacts that can be reused via DSLs are language grammars, source code, software designs, and domain abstractions. Later sections provide many examples of DSLs; here we mention a few from the perspective of reuse.

In his definitive survey of reuse Krueger [1992] categorizes reuse approaches along the following dimensions: abstracting, selecting, specializing, and integrating. In particular, he identifies application generators as an important reuse category. As already noted, application generators often use a DSL as their input language, thereby enabling programmers to reuse

semantic notions embodied in the DSL without having to perform a detailed domain analysis themselves. Examples include BDL [Bertrand and Augeraud 1999] that generates software to control concurrent objects and Teapot [Chandra et al. 1999] that produces implementations of cache coherence protocols. Krueger identifies definition of domain coverage and concepts as a difficult challenge for implementors of application generators. We identify patterns for domain analysis in this article.

DSLs also play a role in other reuse categories identified by Krueger [1992]. For example, software architectures are commonly reused when DSLs are employed because the application generator or compiler follows a standard design when producing code from a DSL input. For example, GAL [Thibault et al. 1999] enables reuse of a standard architecture for video device drivers. DSLs implemented as application libraries clearly enable reuse of source code. Prominent examples are Haskell-based DSLs such as Fran [Elliott 1999]. DSLs can also be used for formal specification of software schemas. For example, Nowra [Sloane 2002] specifies software manufacturing processes and SSC [Buffenbarger and Gruell 2001] deals with subsystem composition.

Reuse may involve exploitation of an existing language grammar. For example, Hancock [Bonachea et al. 1999] piggybacks on C, while SWUL [Bravenboer and Visser 2004] extends Java. Moreover, the success of XML for DSLs is largely based on reuse of its grammar for specific domains. Less formal language grammars may also be reused when notations used by domain experts, but not yet available in a computer language, are realized in a DSL. For example, Hawk [Launchbury et al. 1999] uses a textual form of an existing visual notation.

1.4. Scope of This Article

There are no easy answers to the “when and how” question in the title of this article. The previously mentioned benefits of DSLs do not come free.

- DSL development is hard, requiring both domain and language development expertise. Few people have both.
- DSL development techniques are more varied than those for GPLs, requiring careful consideration of the factors involved.
- Depending on the size of the user community, development of training material, language support, standardization, and maintenance may become serious and time-consuming issues.

These are not the only factors complicating the decision to develop a new DSL. Initially, it is often far from evident that a DSL might be useful or that developing a new one might be worthwhile. This may become clear only after a sizable investment in domain-specific software development using a GPL has been made. The concepts underlying a suitable DSL may emerge only after a lot of GPL programming has been done. In such cases, DSL development may be a key step in software reengineering or software evolution [Bennett and Rajlich 2000].

To aid the DSL developer, we provide a systematic survey of the many factors involved by identifying patterns in the decision, analysis, design, and implementation phases of DSL development (Section 2). Our patterns improve and extend earlier work on DSL design patterns, in particular [Spinellis 2001]. This is discussed in Section 2.6. The DSL development process can be facilitated by using domain analysis tools and language development systems. These are surveyed in Section 3. Finally, conclusions and open problems are presented in Section 4.

1.5. Literature

We give some general pointers to the DSL literature; more specific references are given at appropriate points throughout this article rather than in this section. Until recently, DSLs received relatively little attention in the computer science research community, and there are few books on the subject. We mention Martin [1985], an exhaustive account of 4GLs;

Biggerstaff and Perlis [1989], a two-volume collection of articles on software reuse including DSL development and program generation; Nardi [1993], focuses on the role of DSLs in end-user programming; Salus [1998], a collection of articles on little languages (not all of them DSLs); and Barron [2000], which treats scripting languages (again, not all of them DSLs). Domain analysis, program generators, generative programming techniques, and intentional programming (IP) are treated in Czarnecki and Eisenecker [2000]. Domain analysis and the use of XML, DOM, XSLT, and related languages and tools to generate programs are discussed in Cleaveland [2001]. Domain-specific language development is an important element of the software factories method [Greenfield et al. 2004].

Proceedings of recent workshops and conferences partly or exclusively devoted to DSLs are Kamin [1997]; USENIX [1997, 1999]; HICSS [2001, 2002, 2003]; Lengauer et al. [2004]. Several journals have published special issues on DSLs [Wile and Ramming 1999; Mernik and Lämmel 2001, 2002]. Many of the DSLs used as examples in this article were taken from these sources. A special issue on end-user development is the subject of Sutcliffe and Mehandjiev [2004]. A special issue on program generation, optimization, and platform adaptation is authored by Moura et al. [2005]. There are many workshops and conferences at least partly devoted to DSLs for a particular domain, for example, description of features of telecommunications and other software systems [Gilmore and Ryan 2001]. The annotated DSL bibliography [van Deursen et al. 2000] (78 items) has limited overlap with the references in this article because of our emphasis on general DSL development issues.

2. DSL PATTERNS

2.1. Pattern classification

The following are DSL development phases: *decision*, *analysis*, *design*, *implementation*, and *deployment*. In practice,

Table III. Decision Patterns

Pattern	Description
Notation	Add new or existing domain notation Important subpatterns: • Transform visual to textual notation • Add user-friendly notation to existing API
AVOPT	Domain-specific Analysis, Verification, Optimization, Parallelization, and Transformation
Task automation	Eliminate repetitive tasks
Product line	Specify member of software product line
Data structure representation	Facilitate data description
Data structure traversal	Facilitate complicated traversals
System front-end Interaction	Facilitate system configuration Make interaction programmable
GUI construction	Facilitate GUI construction

DSL development is not a simple sequential process, however. The decision process may be influenced by preliminary analysis which, in turn, may have to supply answers to unforeseen questions arising during design, and design is often influenced by implementation considerations.

We associate classes of patterns with each of the development phases except deployment which is beyond the scope of this article. The decision phase corresponds to the “when” part of DSL development, the other phases to the “how” part. *Decision patterns* are common situations that potential developers may find themselves in for which successful DSLs have been developed in the past. In such situations, use of an existing DSL or development of a new one is a serious option. Similarly, *analysis patterns*, *design patterns*, and *implementation patterns* are common approaches to, respectively, domain analysis, DSL design, and DSL implementation. Patterns corresponding to different DSL development phases are independent. For a particular decision pattern, virtually any analysis or design pattern can be chosen, and the same is true for design and implementation patterns. Patterns in the same class, on the other hand, need not be independent but may have some overlap.

We discuss each development phase and the associated patterns in a separate section. Inevitably, there may be some patterns we have missed.

2.2. Decision

Deciding in favor of a new DSL is usually not easy. The investment in DSL development (including deployment) has to pay for itself by more economical software development and/or maintenance later on. As mentioned in Section 1.1, a quantitative treatment of the trade-offs involved is difficult. In practice, short-term considerations and lack of expertise may easily cause indefinite postponement of the decision. Obviously, adopting an existing DSL is much less expensive and requires much less expertise than developing a new one. Finding out about available DSLs may be hard, since DSL information is scattered widely and often buried in obscure documents. Adopting DSLs that are not well publicized might be considered too risky, anyway.

To aid in the decision process, we identify the decision patterns, shown in Table III. Underlying them are general, interrelated concerns such as:

- improved software economics,
- enabling of software development by users with less domain and programming expertise, or even by end-users with some domain, but virtually no programming expertise [Nardi 1993; Sutcliffe and Mehandjiev 2004].

The patterns in Table III may be viewed as more concrete and specific subpatterns of these general concerns. We briefly discuss each decision pattern in turn. Examples for each pattern are given in Table IV.

Table IV. Examples for the Decision Patterns in Table III

Pattern	DSL	Application Domain
Notation	MSC [SDL Forum 2000]	Telecom system specification
• Visual-to-textual	Hawk [Launchbury et al. 1999]	Microarchitecture design
	MSF [Gray and Karsai 2003]	Tool integration
	Verischemelog [Jennings and Beuscher 1999]	Hardware design
• API-to-DSL	SPL [Xiong et al. 2001]	Digital signal processing
	SWUL [Bravenboer and Visser 2004]	GUI construction
AVOPT	AL [Guyer and Lin 1999]	Software optimization
	ATMOL [van Engelen 2001]	Atmospheric modeling
	BDL [Bertrand and Augeraud 1999]	Coordination
	ESP [Kumar et al. 2001]	Programmable devices
	OWL-Light [Dean et al. 2003]	Web ontology
	PCSL [Bruntink et al. 2005]	Parameter checking
	PLAN-P [Thibault et al. 1998]	Network programming
Task automation	Teapot [Chandra et al. 1999]	Cache coherence protocols
	Facile [Schnarr et al. 2001]	Computer architecture
	JAMOOS [Gil and Tsoglin 2001]	Language processing
	lava [Sirer and Bershad 1999]	Software testing
	PSL-DA [Fertalj et al. 2002]	Database applications
	RoTL [Mauw et al. 2004]	Traffic control
	SHIFT [Antoniotti and Göllü 1997]	Hybrid system design
Product line	SODL [Mernik et al. 2001]	Network applications
Data structure representation	GAL [Thibault et al. 1999]	Video device drivers
	ACML [Gondow and Kawashima 2002]	CASE tools
	ASDL [Wang et al. 1997]	Language processing
	DiSTiL [Smaragdakis and Batory 1997]	Container data structures
	FIDO [Klarlund and Schwartzbach 1999]	Tree automata
Data structure traversal	ASTLOG [Crew 1997]	Language processing
	Hancock [Bonachea et al. 1999]	Customer profiling
	S-XML [Clements et al. 2004; Felleisen et al. 2004]	XML processing
	TVL [Gray and Karsai 2003]	Tool integration
System front-end	Nowra [Sloane 2002]	Software configuration
	SSC [Buffenbarger and Gruell 2001]	Software composition
Interaction	CHEM [Bentley 1986]	Drawing chemical structures
	FPIC [Kamin and Hyatt 1997]	Picture drawing
	Fran [Elliott 1999]	Computer animation
	Mawl [Atkins et al. 1999]	Web computing
	Service Combinators [Cardelli and Davies 1999]	Web computing
GUI construction	AUI [Schneider and Cordy 2002]	User interface construction
	HyCom [Risi et al. 2001]	Hypermedia applications

Table V. Analysis Patterns

Pattern	Description
Informal	The domain is analyzed in an informal way.
Formal	A domain analysis methodology is used.
Extract from code	Mining of domain knowledge from legacy GPL code by inspection or by using software tools, or a combination of both.

Notation. The availability of appropriate (new or existing) domain-specific notations is the decisive factor in this case. Two important subpatterns are:

- Transform visual to textual notation.* There are many benefits to making an existing visual notation available in textual form such as easier composition of large programs or specifications, and enabling of the AVOPT decision pattern discussed next.
- Add user-friendly notation to an existing API* or turn an API into a DSL.

AVOPT. Domain-specific analysis, verification, optimization, parallelization, and transformation of application programs written in a GPL are usually not feasible because the source code patterns involved are too complex or not well defined. Use of an appropriate DSL makes these operations possible. With continuing developments in chip-level multiprocessing (CMP), domain-specific parallelization will become steadily more important [Kuck 2005]. This pattern overlaps with most of the others.

Task automation. Programmers often spend time on GPL programming tasks that are tedious and follow the same pattern. In such cases, the required code can be generated automatically by an application generator (compiler) for an appropriate DSL.

Product line. Members of a software product line [Weiss and Lay 1999] share a common architecture and are developed from a common set of basic elements. Use of a DSL may often facilitate their specification. This pattern has considerable overlap with both the task automation and system front-end patterns.

Data structure representation. Data-driven code relies on initialized data structures

whose complexity may make them difficult to write and maintain. Such structures are often more easily expressed using a DSL.

Data structure traversal. Traversals over complicated data structures can often be expressed better and more reliably in a suitable DSL.

System front-end. A DSL-based front-end may often be used for handling a system's configuration and adaptation.

Interaction. Text- or menu-based interaction with application software often has to be supplemented with an appropriate DSL for the specification of complicated or repetitive input. For example, Excel's interactive mode is supplemented with the Excel macro language to make Excel programmable.

GUI construction. This is often done using a DSL.

2.3. Analysis

In the analysis phase of DSL development, the problem domain is identified and domain knowledge is gathered. Inputs are various sources of explicit or implicit domain knowledge, such as technical documents, knowledge provided by domain experts, existing GPL code, and customer surveys. The output of domain analysis varies widely but consists basically of domain-specific terminology and semantics in more or less abstract form. There is a close link between domain analysis and knowledge engineering which is only beginning to be explored. Knowledge capture, knowledge representation, and ontology development [Denny 2003] are potentially useful in the analysis phase.

The analysis patterns we have identified are shown in Table V. Examples are given in Table VI. Most of the time, domain analysis is done

Table VI. Examples for the Analysis Patterns in Table V
(References and application domains are given in Table IV. The FODA and FAST domain analysis methodologies are discussed in the text.)

Pattern	DSL	Analysis Methodology
Informal	All DSLs in Table IV except:	
Formal	GAL	FAST commonality analysis
	Hancock	FAST
	RoTL	Variability analysis (close to FODA's)
	Service Combinators	FODA (only in this article—see text)
Extract from code	FPIC	Extracted by inspection from PIC implementation
	Nowra	Extracted by inspection from Odin implementation
	PCSL	Extracted by clone detection from proprietary C code
	Verischemelog	Extracted by inspection from Verilog implementation

informally, but sometimes domain analysis methodologies are used. Examples of such methodologies are DARE (Domain Analysis and Reuse Environment) [Frakes et al. 1998], DSSA (Domain-Specific Software Architectures) [Taylor et al. 1995], FAST (Family-Oriented Abstractions, Specification, and Translation) [Weiss and Lay 1999], FODA (Feature-Oriented Domain Analysis) [Kang et al. 1990], ODE (Ontology-based Domain Engineering) [Falbo et al. 2002], or ODM (Organization Domain Modeling) [Simos and Anthony 1998]. To give an idea of the scope of these methods, we explain the FODA and FAST methodologies in somewhat greater detail. Tool support for formal domain analysis is discussed in Section 3.2.

The output of formal domain analysis is a *domain model* consisting of

- a domain definition defining the scope of the domain,
- domain terminology (vocabulary, ontology),
- descriptions of domain concepts,
- feature models describing the commonalities and variabilities of domain concepts and their interdependencies.

How can a DSL be developed from the information gathered in the analysis phase? No clear guidelines exist, but some are presented in Thibault et al. [1999] and Thibault [1998]. Variabilities indicate precisely what information is required to specify an instance of a system. This in-

formation must be specified directly in, or be derivable from, a DSL program. Terminology and concepts are used to guide the development of the actual DSL constructs corresponding to the variabilities. Commonalities are used to define the execution model (by a set of common operations) and primitives of the language. Note that the execution model of a DSL is usually much richer than that for a GPL. On the basis of a single domain analysis, many different DSLs can be developed, but all share important characteristics found in the feature model.

For the sake of concreteness, we apply the FODA domain analysis methodology [Kang et al. 1990] to the service combinator DSL discussed in Cardelli and Davies [1999]. The latter's goal is to reproduce human behavior, while accessing and manipulating Web resources such as reaction to slow transmission, failures, and many simultaneous links. FODA requires construction of a feature model capturing commonalities (mandatory features) and variabilities (variable features). More specifically, such a model consists of

- a feature diagram representing a hierarchical decomposition of features and their character, that is, whether they are mandatory, alternative, or optional,
- definitions of the semantics of features,
- feature composition rules describing which combinations of features are valid or invalid,
- reasons for choosing a feature.

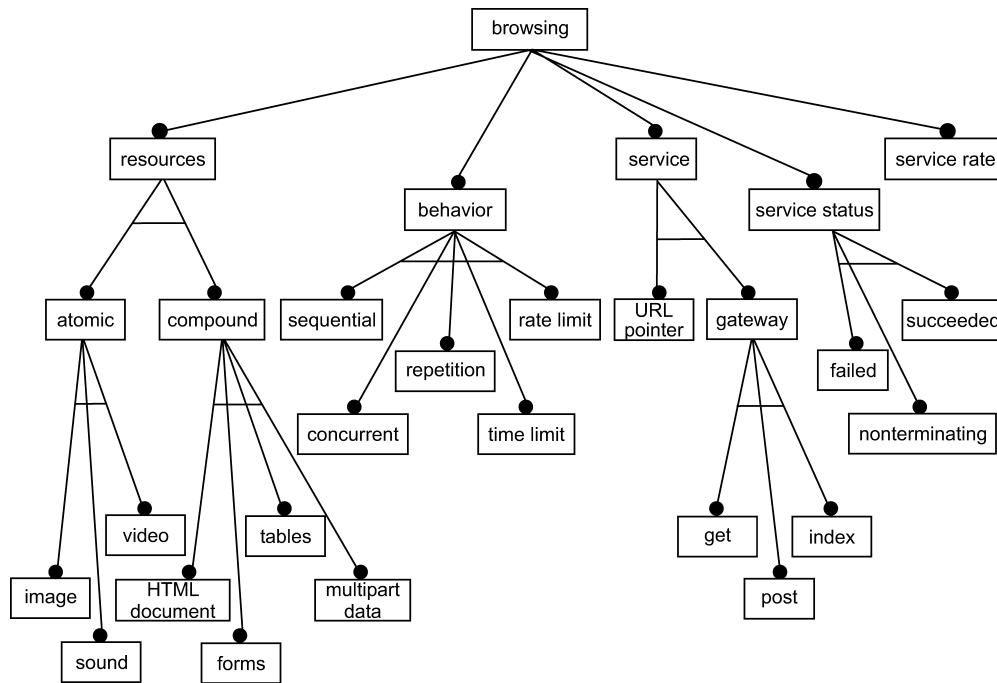


Fig. 1. Feature diagram for Web browsing.

A common feature of a concept is a feature present in all instances of the concept. All mandatory features whose parent is the concept are common features. Also, all mandatory features whose parents are common are themselves common. A variable feature is either optional or alternative (one of, more of). Nodes in the feature diagram to which these features are attached are called variation points.

In the case of our example DSL, the domain consists of resources, browsing behavior, and services (type, status, and rate). Resources can be atomic or compound, access to the resource (service) can be through a URL pointer or a gateway, and browsing behavior can be sequential, concurrent, repetitive, limited by accessing time, or rate. Service has a rate and status (succeeded, failed, or nonterminating). A corresponding feature diagram is shown in Figure 1. The first step in designing the DSL is to look into variabilities and commonalities in the feature diagram. Variable parts must be specified directly in or be derivable from DSL programs. It is

clear that type of service (URL pointer or gateway) and browsing behavior have to be specified in DSL programs. Service status and service rate will be examined and computed while running a DSL program. Therefore, both will be built into the execution model. Type of resource (atomic or compound) are actually types of values that exist during the execution of a DSL program. The basic syntax proposed in Cardelli and Davies [1999]

```

S ::= url(String)           // basic services
| gateway get (String+)
| gateway post (String+)
| index(String, String)
| S1 ? S2                   // sequential execution
| S1 '|' S2                 // concurrent execution
| timeout(Real, S)         // timeout combinator
| limit(Real, Real, S)     // rate limit combinator
| repeat(S)                // repetition
| stall                    // nontermination
| fail                    // failure
  
```

closely resembles our feature diagram. The syntax can later be extended with abstractions and binding.

Table VII. Design Patterns

Pattern	Description
Language exploitation	DSL uses (part of) existing GPL or DSL. Important subpatterns: <ul style="list-style-type: none"> • Piggyback: Existing language is partially used • Specialization: Existing language is restricted • Extension: Existing language is extended
Language invention	A DSL is designed from scratch with no commonality with existing languages
Informal	DSL is described informally
Formal	DSL is described formally using an existing semantics definition method such as attribute grammars, rewrite rules, or abstract state machines

Another domain analysis methodology is FAST (Family-Oriented Abstractions, Specification, and Translation) [Coplien et al. 1998]. FAST is a software development process applying product-line architecture principles, so it relates directly to the product-line decision pattern. A common platform is specified for a family of software products. It is based on the similarities and differences between products. The FAST method consists of the following activities: domain qualification, domain engineering, application engineering, project management, and family change.

During domain engineering, the domain is analyzed and then implemented as a set of domain-specific reusable components. The purpose of domain analysis in FAST is to capture common knowledge about the domain and guide reuse of the implemented components. Domain analysis involves the following steps: decision model definition, commonality analysis, domain design, application modeling language design, creation of standard application engineering process design, and development of the application engineering design environment. An important task of domain analysis is commonality analysis which identifies useful abstractions that are common to all family members. Commonalities are the main source of reuse, thus the emphasis is on finding common parts. Besides the commonalities, variabilities are also discovered during commonality analysis. Variabilities indicate potential sources of change over the lifetime of the family. Commonalities and variabilities in FAST are specified as a structured list. For every variable prop-

erty, the range of variability as well as binding time are specified. Commonality analysis is later used in designing an application modeling language (AML) which is used to generate a family member from specifications.

2.4. Design

Approaches to DSL design can be characterized along two orthogonal dimensions: the relationship between the DSL and existing languages, and the formal nature of the design description. This dichotomy is reflected in the design patterns in Table VII and the corresponding examples in Table VIII.

The easiest way to design a DSL is to base it on an existing language. Possible benefits are easier implementation (see Section 2.5) and familiarity for users, but the latter only applies if users are also programmers in the existing language which may not be the case. We identify three patterns of design based on an existing language. First, we can *piggyback* domain-specific features on part of an existing language. A related approach restricts the existing language to provide a *specialization* targeted at the problem domain. The difference between these two patterns is really a matter of how rigid the barrier is between the DSL and the rest of the existing language. Both of these approaches are often used when a notation is already widely known. For example, many DSLs contain arithmetic expressions which are usually written in the infix-operator style of mathematics.

Another approach is to take an existing language and extend it with new features

Table VIII. Examples for the Design Patterns in Table VII (References and application domains are given in Table IV.)

Pattern	DSL
Language exploitation	
• Piggyback	ACML, ASDL, BDL, ESP, Facile, Hancock, JAMOOS, lava, Mawl, PSL-DA, SPL, SSC, Teapot
• Specialization	OWL-Light
• Extension	AUI, DiSTiL, FPIC, Fran, Hawk, HyCom, Nowra, PLAN-P, SWUL, S-XML, Verischemelog
Language invention	AL, ASTLOG, ATMOL, CHEM, GAL, FIDO, MSF, RoTL, Service Combinators, SHIFT, SODL, TVL
Informal	All DSLs in Table IV except:
Formal	ATMOL, ASTLOG, BDL, FIDO, GAL, OWL-Light, PLAN-P, RoTL, Service Combinators, SHIFT, SODL, SSC

that address domain concepts. In most applications of this pattern, the existing language features remain available. The challenge is to integrate the domain-specific features with the rest of the language in a seamless fashion.

At the other end of the spectrum is a DSL whose design bears no relationship to any existing language. In practice, development of this kind of DSL can be extremely difficult and is hard to characterize. Well-known GPL design criteria such as readability, simplicity, orthogonality, the design principles listed by Brooks [1996], and Tennent's design principles [1977] retain some validity for DSLs. However, the DSL designer has to keep in mind both the special character of DSLs as well as the fact that users need not be programmers. Since ideally the DSL adopts established notations of the domain, the designer should suppress a tendency to improve them. As stated in Wile [2004], one of the lessons learned from real DSL experiments is:

Lesson T2: You are almost never designing a programming language.

Most DSL designers come from language design backgrounds. There the admirable principles of orthogonality and economy of form are not necessarily well-applied to DSL design. Especially in catering to the pre-existing jargon and notations of the domain, one must be careful not to embellish or over-generalize the language.

Lesson T2 Corollary: Design only what is necessary. Learn to recognize your tendency to over-design.

Once the relationship to existing languages has been determined, a DSL

designer must turn to specifying the design before implementation. We distinguish between *informal* and *formal* designs. In an informal design the specification is usually in some form of natural language, probably including a set of illustrative DSL programs. A formal design consists of a specification written using one of the available semantic definition methods [Slonneger and Kurtz 1995]. The most widely used formal notations include regular expressions and grammars for syntax specifications, and attribute grammars, rewrite systems, and abstract state machines for semantic specification.

Clearly, an informal approach is likely to be easiest for most people. A formal approach should not be discounted, however. Development of a formal description of both syntax and semantics can bring problems to light before the DSL is actually implemented. Furthermore, formal designs can be implemented automatically by language development systems and tools, thereby significantly reducing the implementation effort (Section 3).

As mentioned in the beginning of this section, design patterns can be characterized in terms of two orthogonal dimensions: language invention or language exploitation (extension, specialization, or piggyback), and informal or formal description. Figure 2 indicates the position of the DSLs from Table VIII in the design pattern plane. We note that formal description is used more often than informal description when a DSL is designed using the language invention pattern. The opposite is true

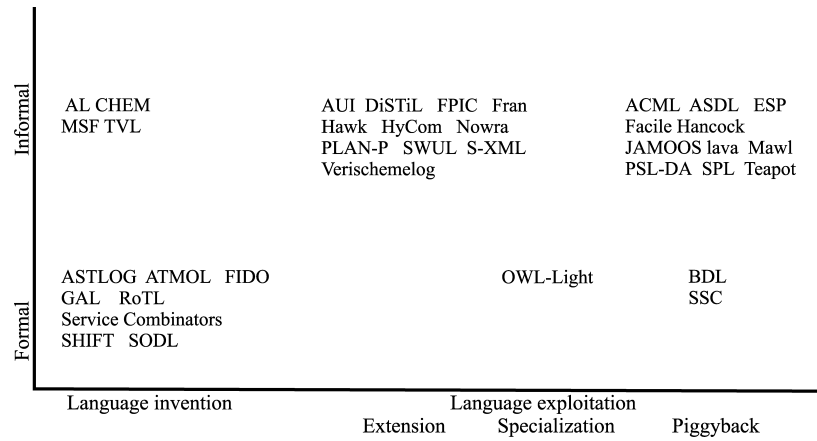


Fig. 2. The DSLs from Table VIII in the design pattern plane.

when a DSL is designed using language exploitation.

2.5. Implementation

2.5.1. Patterns. When an (executable) DSL is designed, the most suitable implementation approach should be chosen. This may be obvious, but in practice it is not, mainly because of the many DSL implementation techniques that have no useful counterpart for GPLs. These DSL-specific techniques are less well known, but can make a big difference in the total effort that has to be invested in DSL development. The implementation patterns we have identified are shown in Table IX. We discuss some of them in more detail. Examples are given in Table X.

Interpretation and compilation are as relevant for DSLs as for GPLs, even though the special character of DSLs often makes them amenable to other, more efficient implementation methods such as preprocessing and embedding. This viewpoint is at variance with Spinellis [2001], where it is argued that DSL development is radically different from GPL development since the former is usually just a small part of a project, and hence DSL development costs have to be modest. This is not always the case, however, and interpreters and compilers or application generators are widely used in practice.

Macros and subroutines are the classic language extension mechanisms used for DSL implementation. Subroutines have given rise to implementation by embedding, while macros are handled by preprocessing. A recent survey of macros is given in Braband and Schwartzbach [2002]. Macro expansion is often independent of the syntax of the base language, and the syntactical correctness of the expanded result is not guaranteed but is checked at a later stage by the interpreter or compiler. This situation is typical for preprocessors.

C++ supports a language-specific preprocessing approach, *template metaprogramming* [Veldhuizen 1995b; Veldhuizen 1995a]. It uses template expansion to achieve compile-time generation of domain-specific code. Significant mileage has been made out of template metaprogramming to develop mathematical libraries for C++ which have familiar domain notation using C++ user-definable operator notation and overloading but also achieve good performance. An example is Blitz++ [Veldhuizen 2001].

In the embedding approach, a DSL is implemented by extending an existing GPL (the host language) by defining specific abstract data types and operators. A domain-specific problem can then be described with these new constructs. Therefore, the new language has all the power of the host language, but an application

Table IX. Implementation Patterns for Executable DSLs

Pattern	Description
Interpreter	DSL constructs are recognized and interpreted using a standard fetch-decode-execute cycle. This approach is appropriate for languages having a dynamic character or if execution speed is not an issue. The advantages of interpretation over compilation are greater simplicity, greater control over the execution environment, and easier extension.
Compiler/application generator	DSL constructs are translated to base language constructs and library calls. A complete static analysis can be done on the DSL program/specification. DSL compilers are often called application generators.
Preprocessor	DSL constructs are translated to constructs in an existing language (the <i>base language</i>). Static analysis is limited to that done by the base language processor. Important subpatterns: <ul style="list-style-type: none"> • Macro processing: Expansion of macro definitions. • Source-to-source transformation: DSL source code is transformed (translated) into base language source code. • Pipeline: Processors successively handling sublanguages of a DSL and translating them to the input language of the next stage. • Lexical processing: Only simple lexical scanning is required, without complicated tree-based syntax analysis.
Embedding	DSL constructs are embedded in an existing GPL (the <i>host language</i>) by defining new abstract data types and operators. Application libraries are the basic form of embedding.
Extensible compiler/ interpreter	A GPL compiler/interpreter is extended with domain-specific optimization rules and/or domain-specific code generation. While interpreters are usually relatively easy to extend, extending compilers is hard unless they were designed with extension in mind.
Commercial Off-The-Shelf (COTS)	Existing tools and/or notations are applied to a specific domain.
Hybrid	A combination of the above approaches.

Table X. Examples for the Implementation Patterns in Table IX (References and application domains are given in Table IV.)

Pattern	DSL
Interpreter	ASTLOG, Service Combinators
Compiler/application generator	AL, ATMOL, BDL, ESP, Facile, FIDO, Hancock, JAMOOS, lava, Mawl, PSL-DA, RoTL, SHIFT, SODL, SPL, Teapot
Preprocessor	
• Macro processing	S-XML
• Source-to-source transformation	ADSL, AUI, MSF, SWUL, TVL
• Pipeline	CHEM
• Lexical processing	SSC
Embedding	FPIC, Fran, Hawk, HyCom, Nowra, Verischemelog
Extensible compiler/interpreter	DiSTiL
Commercial Off-The-Shelf (COTS)	ACML, OWL-Light
Hybrid	GAL, PLAN-P

engineer can become a programmer without learning too much of it. To approximate domain-specific notations as closely as possible, the embedding approach can use any features for user-definable operator syntax the host language has to offer. For example, it is common to develop C++ class libraries where the existing operators are overloaded with domain-specific semantics. Although this technique is quite powerful, pitfalls exist in overloading familiar operators to have unfamiliar semantics. Although the host language in the embedding approach can be any general-purpose language, functional languages are often appropriate as shown by many researchers [Hudak 1998; Kamin 1998]. This is due to functional language features such as lazy evaluation, higher-order functions, and strong typing with polymorphism and overloading.

Extending an existing language implementation can also be seen as a form of embedding. The difference is usually a matter of degree. In an interpreter or compiler approach, the implementation would usually only be extended with a few features such as new data types and operators for them. For a proper embedding, the extensions might encompass full-blown domain-specific language features. In both settings, however, extending implementations is often very difficult. Techniques for doing so in a safe and modular fashion are still the subject of much research. Since compilers are particularly hard to extend, much of this work is aimed at preprocessors and extensible compilers allowing for the addition of domain-specific optimization rules and/or domain-specific code generation. We mention user-definable optimization rules in the CodeBoost C++ preprocessor [Bagge and Haverlaen 2003] and in the Simplicissimus GCC compiler plug-in [Schupp et al. 2001], the IBM Montana extensible C++ programming environment [Soroker et al. 1997], the user-definable optimization rules in the GHC Haskell compiler [Peyton Jones et al. 2001], and the exploitation of domain-specific semantics of application libraries in the Broadway compiler

[Guyer and Lin 2005]. Some extensible compilers such as OpenC++ [Chiba 1995], support a *metaobject protocol*. This is an object-oriented interface for specifying language extensions and transformations [Kiczales et al. 1991].

The COTS-based approach builds a DSL around existing tools and notations. Typically this approach involves applying existing functionality in a restricted way, according to domain rules. For example, the general-purpose Powerpoint tool has been applied in a domain-specific setting for diagram editing [Wile 2001]. The current prominence of XML-based DSLs is another instance of this approach [Gondow and Kawashima 2002; Parigot 2004]. For an XML-based DSL, grammar is described using a DTD or XML scheme where nonterminals are analogous to elements and terminals to data content. Productions are like element definitions where the element name is the left-hand side and the content model is the right-hand side. The start symbol is analogous to the document element in a DTD. Using a DOM parser or SAX (Simple API for XML) tool, parsing comes for free. Since the parse tree can be encoded in XML as well, XSLT transformations can be used for code generation. Therefore, XML and XML tools can be used to implement a programming language compiler [Germon 2001].

Many DSL endeavors apply a number of these approaches in a hybrid fashion. Thus the advantages of different approaches can be exploited. For instance, embedding can be combined with user-defined domain-specific optimization in an extensible compiler, and the interpreter and compiler approach become indistinguishable in some settings (see next section).

2.5.2. Implementation Trade-Offs. Advantages of the interpreter and compiler or application generator approaches are:

- DSL syntax can be close to the notations used by domain experts,
- good error reporting is possible,

—domain-specific analysis, verification, optimization, parallelization, and transformation (AVOPT) is possible.

Some of its disadvantages are:

- the development effort is large because a complex language processor must be implemented,
- the DSL is more likely to be designed from scratch, often leading to incoherent designs compared with exploitation of an existing language,
- language extension is hard to realize because most language processors are not designed with extension in mind.

However, these disadvantages can be minimized or eliminated altogether when a language development system or toolkit is used so that much of the work of the language processor construction is automated. This presupposes a formal approach to DSL design and implementation. Automation support is discussed further in Section 3.

We now turn to the embedded approach. Its advantages are:

- development effort is modest because an existing implementation can be reused,
- it often produces a more powerful language than other methods since many features come for free,
- host language infrastructure can be reused (development and debugging environments: editors, debuggers, tracers, profilers, etc.),
- user training costs might be lower since many users may already know the host language.

Disadvantages of the embedded approach are:

- syntax is far from optimal because most languages do not allow arbitrary syntax extension,
- overloading existing operators can be confusing if the new semantics does not have the same properties as the old,
- bad error reporting because messages are in terms of host language concepts instead of DSL concepts,

—domain-specific optimizations and transformations are hard to achieve so efficiency may be affected, particularly when embedding in functional languages [Kamin 1998; Sloane 2002].

Advocates of the embedded approach often criticize DSLs implemented by the interpreter or compiler approach in that too much effort is put into syntax design, whereas the language semantics tends to be poorly designed and cannot be easily extended with new features [Kamin 1998]. However, the syntax of a DSL is extremely important and should not be underestimated. It should be as close as possible to the notation used in a domain.

In the functional setting, and in particular if Haskell is used, some of these shortcomings can be reduced by using *monads* to modularize the language implementation [Hudak 1998]. Domain-specific optimizations can be achieved using approaches such as user-defined transformation rules in the GHC compiler [Peyton Jones et al. 2001] or a form of whole-program transformation called *partial evaluation* [Jones et al. 1993; Conzel and Marlet 1998]. In C++, template metaprogramming can be used, and user-defined domain-specific optimization is supported by various preprocessors and compilers. See the references in Section 2.5.1.

The decision diagram on how to proceed with DSL implementation (Figure 3) shows when a particular implementation approach is more appropriate. If the DSL is designed from scratch with no commonality with existing languages (invention pattern), the recommended approach is to implement it by embedding, unless domain-specific analysis, verification, optimization, parallelization, or transformation (AVOPT) is required, a domain-specific notation must be strictly obeyed, or the user community is expected to be large.

If the DSL incorporates (part of) an existing language, one would like to reuse (the corresponding part of) the existing language's implementation as well. Apart from this, various implementation

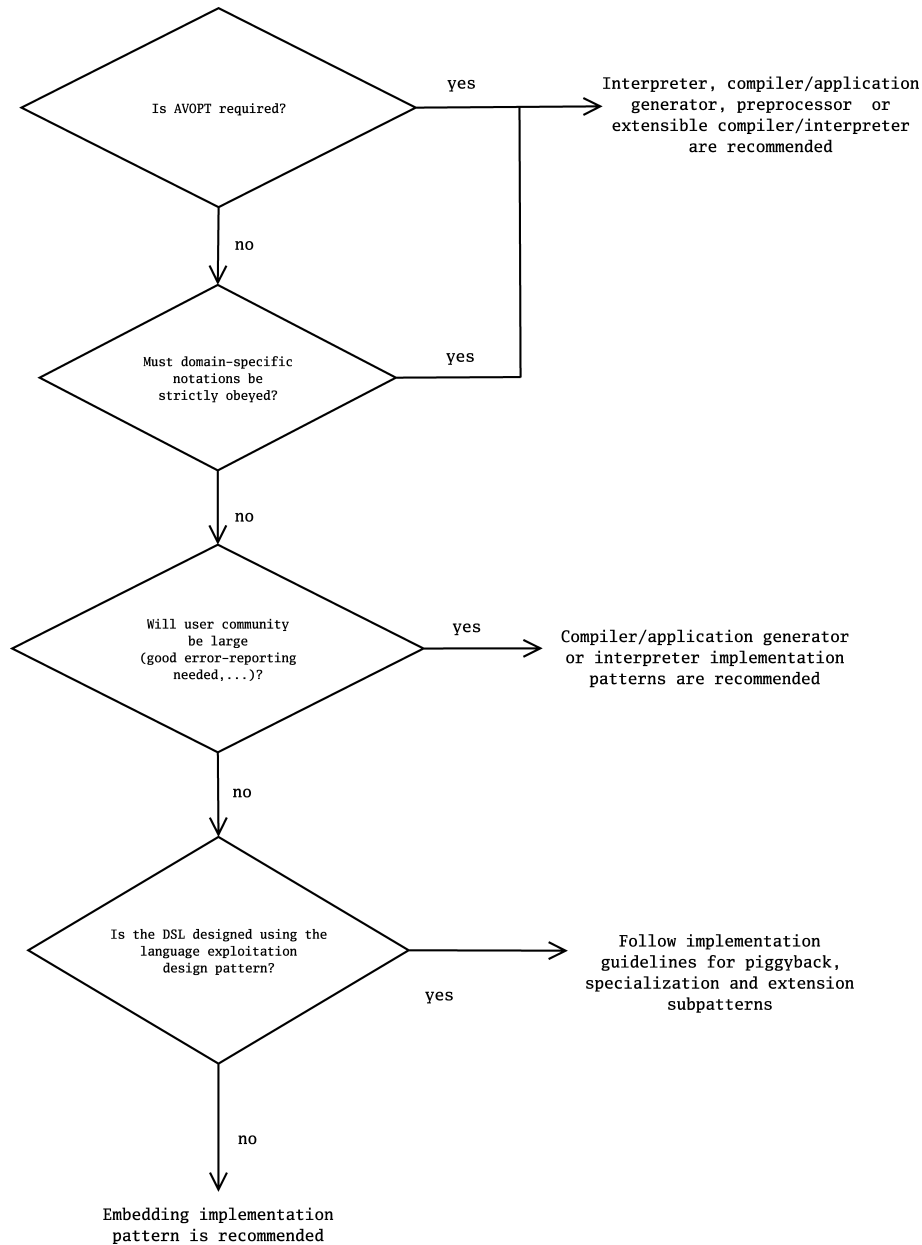


Fig. 3. Implementation guidelines.

Table XI. Pattern Classification Proposed by Spinellis [2001]

Pattern Class	Description
Creational pattern	DSL creation
Structural pattern	Structure of system involving a DSL
Behavioral pattern	DSL interactions

patterns may apply, depending on the language exploitation subpattern used. A piggyback or specialization design can be implemented using an interpreter, compiler or application generator, or preprocessor, but embedding or use of an extensible compiler or interpreter are not suitable,

Table XII. Creational Patterns

Pattern	Description
Language extension	DSL extends existing language with new datatypes, new semantic elements, and/or new syntax.
Language specialization	DSL restricts existing language for purposes of safety, static checking, and/or optimization.
Source-to-source transformation	DSL source code is transformed (translated) into source code of existing language (the base language).
Data structure representation	Data-driven code relies on initialized data structures whose complexity may make them difficult to write and maintain. These structures are often more easily expressed using a DSL.
Lexical processing	Many DSLs may be designed in a form suitable for recognition by simple lexical scanning.

Table XIII. Structural Patterns

Pattern	Description
Piggyback	DSL has elements, for instance, expressions in common with existing language. DSL processor passes those elements to existing language processor.
System front-end	A DSL based front-end may often be used for handling a system's configuration and adaptation.

although specialization can be done using an extensible compiler/interpreter in some languages (Smalltalk, for instance). In the case of piggyback, a preprocessor transforming the DSL to the language it piggybacks on is best from the viewpoint of implementation reuse, but preprocessing has serious shortcomings in other respects. A language extension design can be implemented using all of the previously mentioned implementation patterns. From the viewpoint of implementation reuse, embedding and use of an extensible compiler/interpreter are particularly attractive in this case.

If more than one implementation pattern applies, the one having the highest ratio of benefit (see discussion in this section) to implementation effort is optimal, unless, as in the language invention case, AVOPT is required, a domain-specific notation must be strictly obeyed, or the user community is expected to be large. As already mentioned, a compiler or application generator scores the worst in terms of implementation effort. Less costly are (in descending order) the interpreter, preprocessing, extensible compiler or interpreter, and embedding. On the other hand, a compiler or application generator and interpreter score best as far as benefit to DSL users is concerned. Less benefit is obtained from (in descending order) extensible compiler or interpreter, embedding,

Table XIV. Behavioral Patterns

Pattern	Description
Pipeline	Pipelined processors successively handling sublanguages of a DSL and translating them to input language of next stage.

and preprocessing. In practice, such a cost-benefit analysis is rarely performed, and the decision is driven only by implementor experience. Of course, the latter should be taken into account, but it is not the only relevant factor.

2.6. Comparison With Other Classifications

We start by comparing our patterns with those proposed in Spinellis [2001]. Closely following Gamma et al. [1995], Spinellis distinguishes three classes of DSL patterns as shown in Table XI. The specific patterns for each class are summarized in Tables XII, XIII, and XIV. Most patterns are creational. The piggyback pattern might be classified as creational as well since it is very similar to language extension. This would leave only a single pattern in each of the other two categories.

First, it should be noted that Spinellis's [2001] patterns do not include traditional GPL design and implementation techniques, while ours do, since we consider them to be as relevant for DSLs as for GPLs. Second, Spinellis's

Table XV. Correspondence of Spinellis's [2001] Patterns With Ours
(Since our patterns have a wider scope, many of them have no counterpart in Spinellis's classification. These are not shown in the right-hand column.)

Spinellis's Pattern	Our Pattern
Creational: language extension	Design: language exploitation (extension)
Creational: language specialization	Design: language exploitation (specialization)
Creational: source-to-source transformation	Implementation: preprocessing (source-to-source transformation)
Creational: data structure representation	Decision: data structure representation
Creational: lexical processing	Implementation: preprocessing
Structural: piggyback	Design: language exploitation (piggyback)
Structural: system front-end	Decision: system front-end
Behavioral: pipeline	Implementation: preprocessing (pipeline)

classification does not correspond in an obvious way to our classification in decision, analysis, design, and implementation patterns. The latter are all basically creational, but cover a wider range of creation-related activities than Spinellis's patterns.

The correspondence of Spinellis's [2001] patterns with ours is shown in Table XV. Since our patterns have a wider scope, many of them have no counterpart in Spinellis's classification. These are not shown in the right-hand column. We have retained the terminology used by Spinellis whenever appropriate.

Another classification of DSL development approaches is given in Wile [2001], namely, full language design, language extension, and COTS-based approaches. Since each approach has its own pros and cons, the author discusses them with respect to three kinds of issues, DSL-specific, GPL support, and pragmatic support issues. Finally, the author shows how a hybrid development approach can be used.

3. DSL DEVELOPMENT SUPPORT

3.1. Design and Implementation Support

As we have seen, DSL development is hard, requiring both domain knowledge and language development expertise. The development process can be facilitated by using a *language development system* or *toolkit*. Some systems and toolkits that have actually been used for DSL development are listed in Table XVI. They have widely different capabilities and are in many different stages of development but are based on the same general principle:

they generate tools from language descriptions [Heering and Klint 2000]. The tools generated may vary from a consistency checker and interpreter to an integrated development environment (IDE), consisting of a syntax-directed editor, a prettyprinter, an (incremental) consistency checker, analysis tools, an interpreter or compiler/application generator, and a debugger for the DSL in question (assuming it is executable). As noted in Section 1.2, nonexecutable DSLs may also benefit from various kinds of tool support such as syntax-directed editors, prettyprinters, consistency checkers, and analyzers. These can be generated in the same way.

Some of these systems support a specific DSL design methodology, while others have a largely methodology-independent character. For instance, Sprint [Consel and Marlet 1998] assumes an interpreter for the given DSL and then uses partial evaluation to remove the interpretation overhead by automatically transforming a DSL program into a compiled program. Other systems, such as ASF+SDF [van den Brand et al. 2001], DMS [Baxter et al. 2004], and Stratego [Visser 2003], would not only allow an interpretive definition of the DSL, but would also accept a transformational or translational one. On the other hand, they might not support partial evaluation of a DSL interpreter given a specific program.

The input into these systems is a description of various aspects of the DSL that are developed in terms of specialized metalanguages. Depending on the type of DSL, some important language

Table XVI. Some Language Development Systems and Toolkits That Have Been Used for DSL Development

System	Developer
ASF+SDF [van den Brand et al. 2001]	CWI/University of Amsterdam
AsmL [Glässer et al. 2002]	Microsoft Research, Redmond
DMS [Baxter et al. 2004]	Semantic Designs, Inc.
Draco [Neighbors 1984]	University of California, Irvine
Eli [Gray et al. 1992]	University of Colorado, University of Paderborn, Macquarie University
Gem-Mex [Anlauff et al. 1999]	University of LAquila
InfoWiz [Nakatani and Jones 1997]	Bell Labs/AT&T Labs
JTS [Batory et al. 1998]	University of Texas at Austin
Khepera [Faith et al. 1997]	University of North Carolina
Kodiyak [Herndon and Berzins 1988]	University of Minnesota
LaCon [Kastens and Pfahler 1998]	University of Paderborn (LaCon uses Eli as back-end—see above)
LISA [Mernik et al. 1999]	University of Maribor
metafront [Braband et al. 2003]	University of Aarhus
Metatool [Cleveland 1988]	Bell Labs
POPART [Wile 1993]	USC/Information Sciences Institute
SmartTools [Attali et al. 2001]	INRIA Sophia Antipolis
smgn [Kienle and Moore 2002]	Intel Compiler Lab/University of Victoria
SPARK [Aycok 2002]	University of Calgary
Sprint [Consel and Marlet 1998]	LaBRI/INRIA
Stratego [Visser 2003]	University of Utrecht
TXL [Cordy 2004]	University of Toronto/Queen's University at Kingston

Table XVII. Development Support Provided by Current Language Development Systems and Toolkits for DSL Development Phases/Pattern Classes

Development phase/ Pattern class	Support Provided
Decision	None
Analysis	Not yet integrated—see Section 3.2
Design	Weak
Implementation	Strong

aspects are syntax, prettyprinting, consistency checking, analysis, execution, translation, transformation, and debugging. It so happens that the metalanguages used for describing these aspects are themselves DSLs for the particular aspect in question. For instance, DSL syntax is usually described in something close to BNF, the de facto standard for syntax specification (Table I). The corresponding tool generated by the language development system is a parser.

Although the various specialized metalanguages used for describing language aspects differ from system to system, they are often (but not always) rule based. For instance, depending on the system, the consistency of programs or scripts may have to be checked in terms of *attributed*

syntax rules (an extension of BNF), *conditional rewrite rules*, or *transition rules*. See, for instance, Slonneger and Kurtz [1995] for further details.

The level of support provided by these systems in various phases of DSL development is summarized in Table XVII. Their main strength lies in the implementation phase. Support of DSL design tends to be weak. Their main assets are the metalanguages they support and, in some cases, a meta-environment to aid in constructing and debugging language descriptions but they have little built-in knowledge of language concepts or design rules. Furthermore, to the best of our knowledge, none of them provides any support in the analysis or decision phase. Analysis support tools are discussed in Section 3.2.

Examples of DSL development using the systems in Table XVI are given in Table XVIII. They cover a wide range of application domains and implementation patterns. The Box prettyprinting metalanguage is an example of a DSL developed with a language development system (in this case ASF+SDF) for later use as one of the metalanguages of the system itself. This is common practice. The

Table XVIII. Examples of DSL Development using the Systems in Table XVI

System Used	DSL	Application Domain
ASF+SDF	Box [van den Brand and Visser 1996] Risla [van Deursen and Klint 1998]	Prettyprinting Financial products
AsmL	UPnP [UPnP 2003] XLANG [Thatte 2001]	Networked device protocol Business protocols
DMS	(Various) [Baxter et al. 2004] (Various) [Baxter et al. 2004]	Program transformation Factory control
Eli	Maptool [Kadhim and Waite 1996] (Various) [Pfhler and Kastens 2001]	Grammar mapping Class generation
Gem-Mex	Cubix [Kutter et al. 1998]	Virtual data warehousing
JTS	Jak [Batory et al. 1998]	Syntactic transformation
LaCon	(Various) [Kastens and Pfhler 1998]	Data model translation
LISA	SODL [Mernik et al. 2001]	Network applications
SmartTools	LML [Parigot 2004] BPEL [Courbis and Finkelstein 2004]	GUI programming Business process description
smgn	Hoof [Kienle and Moore 2002] IMDL [Kienle and Moore 2002]	Compiler IR specification Software reengineering
SPARK	Guide [Levy 1998] CML2 [Raymond 2001]	Web programming Linux kernel configuration
Sprint	GAL [Thibault et al. 1999] PLAN-P [Thibault et al. 1998]	Video device drivers Network programming
Stratego	Autobundle [de Jonge 2002] CodeBoost [Bagge and Haverlaen 2003]	Software building Domain-specific C++ optimization

metalanguages for syntax, prettyprinting, attribute evaluation, and program transformation used by DMS were all implemented using DMS, and the Jak transformational metalanguage for specifying the semantics of a DSL or domain-specific language extension in the Jakarta Tool Suite (JTS) was also developed using JTS.

3.2. Analysis Support

The language development toolkits and systems discussed in the previous section do not provide support in the analysis phase of DSL development. Separate frameworks and tools for this have been or are being developed, however. Some of them are listed in Table XIX. We have included a short description of each entry, largely taken from the reference given for it. The fact that a framework or tool is listed does not necessarily mean it is in use or even exists.

As noted in Section 2.3, the output of domain analysis consists basically of domain-specific terminology and semantics in more or less abstract form. It may range from a feature diagram (see FDL entry in Table XIX) to a domain implementation consisting of a set of domain-specific reusable components (see DARE entry in Table XIX), or a theory in the case of sci-

entific domains. An important issue is how to link formal domain analysis with DSL design and implementation. The possibility of linking DARE directly to the Meta-tool metagenerator (i.e., application generator) [Cleveland 1988] is mentioned in Frakes [1998].

4. CONCLUSIONS AND OPEN PROBLEMS

DSLs will never be a solution to all software engineering problems, but their application is currently unduly limited by a lack of reliable knowledge available to (potential) DSL developers. To help remedy this situation, we distinguished five phases of DSL development and identified patterns in each phase, except deployment. These are summarized in Table XX. Furthermore, we discussed language development systems and toolkits that can be used to facilitate the development process especially its later phases.

Our survey also showed many opportunities for further work. As indicated in Table XVII, for instance, there are serious gaps in the DSL development support chain. More specifically, some of the issues needing further attention follow.

Decision. Can useful computer-aided decision support be provided? If so,

Table XIX. Some Domain Analysis Frameworks and Tools

Analysis Framework or Tool	Description
Ariadne [Simos and Anthony 1998]	ODM support framework enabling domain practitioners to collaboratively develop and evolve their own semantic models, and to compose and customize applications incorporating these models as first-class architectural elements.
DARE [Frakes et al. 1998]	Supports the capture of domain information from experts, documents, and code in a domain. Captured domain information is stored in a domain book that will typically contain a generic architecture for the domain and domain-specific reusable components.
DOMAIN [Tracz and Coglianese 1995]	DSSA [Taylor et al. 1995] support framework consisting of a collection of structured editors and a hypertext/media engine that allows the user to capture, represent, and manipulate various types of domain knowledge in a hyper-web. DOMAIN supports a “scenario-based” approach to domain analysis. Users enter scenarios describing the functions performed by applications in the domain of interest. The text in these scenarios can then be used (in a semi-automated manner) to develop a domain dictionary, reference requirements, and domain model, each of which are supported by their own editor.
FDL [van Deursen and Klint 2002]	The Feature Description Language (FDL) is a textual representation of feature diagrams, which are a graphical notation for expressing assertions (propositions, predicates) about systems in a particular application domain. These were introduced in the FODA [Kang et al. 1990] domain analysis methodology. (FDL is an example of the visual-to-textual transformation subpattern in Table III.)
ODE editor [Falbo et al. 2002]	Ontology editor supporting ODE—see also [Denny 2003].

Table XX. Summary of DSL Development Phases and Corresponding Patterns

Development Phase	Pattern
Decision (Section 2.2)	Notation
	AVOPT
	Task automation
	Product line
	Data structure representation
	Data structure traversal
	System front-end
Analysis (Section 2.3)	Interaction
	GUI construction
	Informal
	Formal
Design (Section 2.4)	Extract from code
	Language exploitation
	Language invention
Implementation (Section 2.5)	Informal
	Formal
	Interpreter
	Compiler/application generator
	Preprocessor
	Embedding
	Extensible compiler/interpreter
	COTS
Hybrid	

its integration in existing language development systems or toolkits (Table XVI) might yield additional advantages.

Analysis. Further development and integration of domain analysis support tools. As noted in Section 2.3, there is a close link with knowledge engineering. Existing knowledge engineering tools and frameworks may be useful directly or act as inspiration for further developments in this area. An important issue is how to link formal domain analysis with DSL design and implementation.

Design and Implementation. How can DSL design and implementation be made easier for domain experts not versed in GPL development? Some approaches (not mutually exclusive) are:

—Building DSLs in an incremental, modular, and extensible way from parameterized language building blocks. This

is of particular importance for DSLs since they change more frequently than GPLs [Bosch and Dittrich; Wile 2001]. Progress in this direction is being made [Anlauff et al. 1999; Consel and Marlet 1998; Hudak 1998; Mernik et al. 2000].

- A related issue is how to combine different parts of existing GPLs and DSLs into a new DSL. For instance, in the Microsoft .NET framework, many GPLs are compiled to the Common Language Runtime (CLR) [Gough 2002]. Can this be helpful in including selected parts of GPLs into a new DSL?
- Provide pattern aware development support. The Sprint system [Consel and Marlet 1998], for instance, provides partial evaluation support for the interpreter pattern (see Section 3.1). Other patterns might benefit from specialized support as well. Embedding support is discussed separately in the next paragraph.
- Reduce the need for learning some of the specialized metalanguages of language development systems by supporting description by example (DBE) of selected language aspects like syntax or prettyprinting. The user-friendliness of DBE is due to the fact that examples of intended behavior do not require a specialized metalanguage, or possibly only a small part of it. Grammar inference from example sentences, for instance, may be viable especially since many DSLs are small. This is certainly no new idea [Crespi-Reghezzi et al. 1973; Nardi 1993], but it remains to be realized. Some preliminary results are reported in Črepinšek et al. [2005].
- How can DSL development tools generated by language development systems and toolkits be integrated with other software development tools? Using a COTS-based approach, XML technologies such as DOM and XML-parsers have great potential as a uniform data interchange format for CASE tools. See also Badros [2000] and Cleaveland [2001].

Embedding. GPLs should provide more powerful support for embedding DSLs, both syntactically and semantically. Some issues are:

- Embedding suffers from the very limited user-definable syntax offered by GPLs. Perhaps surprisingly, there has been no trend toward more powerful user-definable syntax in GPLs over the years. In fact, just the opposite has happened. Macros and user-definable operators have become less popular. Java has no user-definable operators at all. On the other hand, some of the language development systems in Table XVI, such as ASF+SDF and to some extent Stratego, support metalanguages featuring fully general user-definable context-free syntax. Although these metalanguages cannot compete directly with GPLs as embedding hosts as far as expressiveness and efficiency are concerned, they can be used to express a source-to-source transformation to translate user-defined DSL syntax embedded in a GPL to appropriate API calls. See Bravenboer and Visser [2004] for an extensive discussion of this approach.
- Improved embedding support is not only a matter of language features, but also of language implementation and, in particular, of preprocessors or extensible compilers allowing the addition of domain-specific optimization rules and/or domain-specific code generation. See the references given in Section 2.5.1 and Granicz and Hickey [2003] and Saraiva and Schneider [2003]. Alternatively, the GPL itself might feature domain-specific optimization rules as a special kind of compiler directive. Such compiler extension makes the embedding process significantly more complex, however, and its cost-benefit ratio needs further scrutiny.

Estimation. Last but not least: In this article, our approach toward DSL development has been qualitative. Can the costs and benefits of DSLs be reliably quantified?

ACKNOWLEDGMENTS

We would like to thank the anonymous reviewers for many useful comments. Arie van Deursen kindly gave us permission to use the source of the annotated DSL bibliography [van Deursen et al. 2000].

REFERENCES

- ANLAUFF, M., KUTTER, P. W., AND PIERANTONIO, A. 1999. Tool support for language design and prototyping with Montages. In *Compiler Construction (CC'99)*, S. Jähnichen, Ed. Lecture Notes in Computer Science, vol. 1575. Springer-Verlag, 296–299.
- ANTONIOTTI, M. AND GÖLLÜ, A. 1997. SHIFT and SMART-AHS: A language for hybrid system engineering modeling and simulation. In *Proceedings of the USENIX Conference on Domain-Specific Languages*, 171–182.
- ATKINS, D., BALL, T., BRUNS, G., AND COX, K. 1999. Mawl: A domain-specific language for form-based services. *IEEE Trans. Softw. Eng.* 25, 3 (May/June), 334–346.
- ATTALI, I., COURBIS, C., DEGENNE, P., FAU, A., PARIGOT, D., AND PASQUIER, C. 2001. SmartTools: A generator of interactive environments tools. In *Compiler Construction: 10th International Conference (CC'01)*, R. Wilhelm, Ed. Lecture Notes in Computer Science, vol. 2027. Springer-Verlag, 355–360.
- AYCOCK, J. 2002. The design and implementation of SPARK, a toolkit for implementing domain-specific languages. *J. Comput. Inform. Tech.* 10, 1, 55–66.
- BACKUS, J. W. 1960. The syntax and semantics of the proposed International Algebraic Language of the Zurich ACM-GAMM conference. In *Proceedings of the International Conference on Information Processing, UNESCO, Paris, 1959*. Oldenbourg, Munich and Butterworth, London, 125–132.
- BADROS, G. 2000. JavaML: A markup language for Java source code. In *Proceedings of the 9th International World Wide Web Conference*. <http://www9.org/w9cdrom/start.html>.
- BAGGE, O. S. AND HAVERAAEN, M. 2003. Domain-specific optimisation with user-defined rules in CodeBoost. In *Proceedings of the 4th International Workshop on Rule-Based Programming (RULE'03)*, J.-L. Giavitto and P.-E. Moreau, Eds. Electronic Notes in Theoretical Computer Science, vol. 86(2). Elsevier. <http://www.sciencedirect.com/>.
- BARRON, D. W. 2000. *The World of Scripting Languages*. John Wiley.
- BATORY, D., LOFASO, B., AND SMARAGDAKIS, Y. 1998. JTS: Tools for implementing domain-specific languages. In *Proceedings of the 5th International Conference on Software Reuse (JCSR'98)*, P. Devanbu and J. Poulin, Eds. IEEE Computer Society, 143–153.
- BATORY, D., THOMAS, J., AND SIRKIN, M. 1994. Reengineering a complex application using a scalable data structure compiler. In *Proceedings of the ACM SIGSOFT International Symposium on the Foundations of Software Engineering*. 111–120.
- BAXTER, I. D., PIDGEON, C., AND MEHLICH, M. 2004. DMS: Program transformation for practical scalable software evolution. In *Proceedings of the 26th International Conference on Software Engineering (ICSE'04)*. IEEE Computer Society, 625–634.
- BENNETT, K. H. AND RAJLICH, V. T. 2000. Software maintenance and evolution: A roadmap. In *The Future of Software Engineering*, A. Finkelstein, Ed. ACM Press, 73–87.
- BENTLEY, J. L. 1986. Programming pearls: Little languages. *Comm. ACM* 29, 8 (August), 711–721.
- BERGIN, T. J. AND GIBSON, R. G., Eds. 1996. *History of Programming Languages II*. ACM Press.
- BERTRAND, F. AND AUGERAUD, M. 1999. BDL: A specialized language for per-object reactive control. *IEEE Trans. Softw. Eng.* 25, 3, 347–362.
- BIGGERSTAFF, T. J. 1998. A perspective of generative reuse. *Annals Softw. Eng.* 5, 169–226.
- BIGGERSTAFF, T. J. AND PERLIS, A. J., Eds. 1989. *Software Reusability*. ACM Press/Addison-Wesley. Vol. I: Concepts and Models, Vol. II: Applications and Experience.
- BONACHEA, D., FISHER, K., ROGERS, A., AND SMITH, F. 1999. Hancock: A language for processing very large-scale data. In *Proceedings of the 2nd USENIX Conference on Domain-Specific Languages*, 163–176.
- BOSCH, J. AND DITTRICH, Y. Domain-specific languages for a changing world. <http://www.cs.rug.nl/bosch/articles.html>.
- BRABAND, C. AND SCHWARTZBACH, M. 2002. Growing languages with metamorphic syntax macros. *ACM SIGPLAN Notices* 37, 3 (March), 31–40.
- BRABAND, C., SCHWARTZBACH, M. I., AND VANGGAARD, M. 2003. The metafront system: Extensible parsing and transformation. In *Proceedings of the 3rd Workshop on Language Descriptions, Tools, and Applications (LDTA'03)*, B. R. Bryant and J. Saraiva, Eds. Electronic Notes in Theoretical Computer Science, vol. 82(3). Elsevier. <http://www.sciencedirect.com/>.
- BRAVENBOER, M. AND VISSER, E. 2004. Concrete syntax for objects: Domain-specific language embedding and assimilation without restrictions. In *Proceedings of the 19th ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'04)*, D. C. Schmidt, Ed. ACM, 365–383.
- BROOKS, JR., F. P. 1996. Language design as design. In *History of Programming Languages II*. T. J. Bergin and R. C. Gibson Eds. ACM Press, 4–15.

- BRUNTINK, M., VAN DEURSEN, A., AND TOURWÉ, T. 2005. Isolating idiomatic crosscutting concerns. In *Proceedings of the International Conference on Software Maintenance (ICSM'05)*. IEEE Computer Society, 37–46.
- BUFFENBARGER, J. AND GRUILL, K. 2001. A language for software subsystem composition. In *IEEE Proceedings of the 34th Hawaii International Conference on System Sciences*.
- CARDELLI, L. AND DAVIES, R. 1999. Service combinators for web computing. *IEEE Trans. Softw. Eng.* 25, 3 (May/June), 309–316.
- CHANDRA, S., RICHARDS, B., AND LARUS, J. R. 1999. Teapot: A domain-specific language for writing cache coherence protocols. *IEEE Trans. Softw. Eng.* 25, 3 (May/June), 317–333.
- CHAPPELL, D. 1996. *Understanding ActiveX and OLE*. Microsoft Press.
- CHIBA, S. 1995. A metaobject protocol for C++. In *Proceedings of the ACM Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA'95)*. ACM, 285–299.
- CLEAVELAND, J. C. 1988. Building application generators. *IEEE Softw.* 5, 4, 25–33.
- CLEAVELAND, J. C. 2001. *Program Generators Using Java and XML*. Prentice-Hall.
- CLEMENTS, J., FELLEISEN, M., FINDLER, R., FLATT, M., AND KRISHNAMURTHI, S. 2004. Fostering little languages. *Dr. Dobb's J.* 29, 3 (March), 16–24.
- CONSEL, C. AND MARLET, R. 1998. Architecturing software using a methodology for language development. In *Principles of Declarative Programming (PLILP'98/ALP'98)*, C. Palamidessi, H. Glaser, and K. Meinke, Eds. Lecture Notes in Computer Science, vol. 1490. Springer-Verlag, 170–194.
- COPLIEN, J., HOFFMAN, D., AND WEISS, D. 1998. Commonality and variability in software engineering. *IEEE Softw.* 15, 6, 37–45.
- CORDY, J. R. 2004. TXL—A language for programming language tools and applications. In *Proceedings of the 4th Workshop on Language Descriptions, Tools, and Applications (LDTA'04)*, G. Hedin and E. van Wyk, Eds. Electronic Notes in Theoretical Computer Science, vol. 110. Elsevier, 3–31. <http://www.sciencedirect.com/>.
- COURBIS, C. AND FINKELSTEIN, A. 2004. Towards an aspect weaving BPEL engine. In *Proceedings of the 3rd AOSD Workshop on Aspects, Components, and Patterns for Infrastructure Software (ACP4IS)*, Y. Coady and D. H. Lorenz, Eds. Tech. rep. NU-CCIS-04-04, College of Computer and Information Science, Northeastern University, Boston, MA.
- ČREPINŠEK, M., MERNIK, M., JAVED, F., BRYANT, B. R., AND SPRAGUE, A. 2005. Extracting grammar from programs: evolutionary approach. *ACM SIGPLAN Notices* 40, 4 (April), 39–46.
- CRESPI-REGHIZZI, S., MELKANOFF, M. A., AND LICHTEN, L. 1973. The use of grammatical inference for designing programming languages. *Comm. ACM* 16, 83–90.
- CREW, R. F. 1997. ASTLOG: A language for examining abstract syntax trees. In *Proceedings of the USENIX Conference on Domain-Specific Languages*, 229–242.
- CZARNECKI, K. AND EISENECKER, U. 2000. *Generative Programming: Methods, Techniques and Applications*. Addison-Wesley.
- DE JONGE, M. 2002. Source tree composition. In *Software Reuse: Methods, Techniques, and Tools: 7th International Conference (ICSR-7)*, C. Gacek, Ed. Lecture Notes in Computer Science, vol. 2319. Springer-Verlag, 17–32.
- DEAN, M., SCHREIBER, G., VAN HARMELEN, F., HENDLER, J., HORROCKS, I., MCGUINNESS, D. L., PATEL-SCHNEIDER, P. F., AND STEIN, L. A. 2003. OWL Web Ontology Language Reference. Working draft, W3C (March). <http://www.w3.org/TR/2003/WD-owl-ref-20030331/>.
- DENNY, M. 2003. Ontology building: A survey of editing tools. Tech. rep., XML.com. <http://www.xml.com/lpt/a/2002/11/06/ontologies.html>.
- ELLIOTT, C. 1999. An embedded modeling language approach to interactive 3D and multimedia animation. *IEEE Trans. Softw. Eng.* 25, 3 (May/June), 291–308.
- FAITH, R. E., NYLAND, L. S., AND PRINS, J. F. 1997. Khepera: A system for rapid implementation of domain specific languages. In *Proceedings of the USENIX Conference on Domain-Specific Languages*, 243–255.
- FALBO, R. A., GUIZZARDI, G., AND DUARTE, K. C. 2002. An ontological approach to domain engineering. In *Proceedings of the 14th International Conference on Software Engineering and Knowledge Engineering (SEKE'02)*. ACM, 351–358.
- FELLEISEN, M., FINDLER, R., FLATT, M., AND KRISHNAMURTHI, S. 2004. Building little languages with macros. *Dr. Dobb's J.* 29, 4 (April), 45–49.
- FERTALJ, K., KALPIČ, D., AND MORNAR, V. 2002. Source code generator based on a proprietary specification language. In *Proceedings of the 35th Hawaii International Conference on System Sciences*.
- FRAKES, W. 1998. Panel: Linking domain analysis with domain implementation. In *Proceedings of the 5th International Conference on Software Reuse*. IEEE Computer Society, 348–349.
- FRAKES, W., PRIETO-DIAZ, R., AND FOX, C. 1998. DARE: Domain analysis and reuse environment. *Annals of Software Engineering* 5, 125–141.
- GAMMA, E., HELM, R., JOHNSON, R., AND VLISSIDES, J. 1995. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley.
- GERMON, R. 2001. Using XML as an intermediate form for compiler development. In *XML Conference Proceedings*. <http://www.idealiance.org/papers/xml2001/index.html>.

- GIL, J. AND TSOGLIN, Y. 2001. JAMOOS—A domain-specific language for language processing. *J. Comput. Inform. Tech.* 9, 4, 305–321.
- GILMORE, S. AND RYAN, M., Eds. 2001. *Language Constructs for Describing Features—Proceedings of the FIREworks Workshop*. Springer-Verlag.
- GLÄSSER, U., GUREVICH, Y., AND VEANES, M. 2002. An abstract communication model. Tech. rep. MSR-TR-2002-55. Microsoft Research, Redmond, WA.
- GONDOW, K. AND KAWASHIMA, H. 2002. Towards ANSI C program slicing using XML. In *Proceedings of the 2nd Workshop on Language Descriptions, Tools, and Applications (LDTA'02)*, M. G. J. van den Brand and R. Lämmel, Eds. Electronic Notes in Theoretical Computer Science, vol. 65(3). Elsevier. <http://www.sciencedirect.com/>.
- GOUGH, J. 2002. *Compiling for the .NET Common Language Runtime (CLR)*. Prentice Hall.
- GRANICZ, A. AND HICKEY, J. 2003. Phobos: Extending compilers with executable language definitions. In *Proceedings of the 36th Hawaii International Conference on System Sciences*.
- GRAY, J. AND KARSAL, G. 2003. An examination of DSLs for concisely representing model traversals and transformations. In *Proceedings of the 36th Hawaii International Conference on System Sciences*.
- GRAY, R. W., LEVI, S. P., HEURING, V. P., SLOANE, A. M., AND WAITE, W. M. 1992. Eli: A complete, flexible compiler construction system. *Comm. ACM* 35, 2 (Feb.), 121–130.
- GREENFIELD, J., SHORT, K., COOK, S., KENT, S., AND CRUPI, J. 2004. *Software Factories: Assembling Applications with Patterns, Models, Frameworks, and Tools*. John Wiley.
- GUYER, S. Z. AND LIN, C. 1999. An annotation language for optimizing software libraries. In *Proceedings of the 2nd USENIX Conference on Domain-Specific Languages*, 39–52.
- GUYER, S. Z. AND LIN, C. 2005. Broadway: A compiler for exploiting the domain-specific semantics of software libraries. In *Proceedings of IEEE*, 93, 2, 342–357.
- HEERING, J. AND KLINT, P. 2000. Semantics of programming languages: A tool-oriented approach. *ACM SIGPLAN Notices* 35, 3 (March) 39–48.
- HERNDON, R. M. AND BERZINS, V. A. 1988. The realizable benefits of a language prototyping language. *IEEE Trans. Softw. Eng.* 14, 803–809.
- HICSS 2001. *Proceedings of the 34th Hawaii International Conference on System Sciences (HICSS'34)*. IEEE.
- HICSS 2002. *Proceedings of the 35th Hawaii International Conference on System Sciences (HICSS'35)*. IEEE.
- HICSS 2003. *Proceedings of the 36th Hawaii International Conference on System Sciences (HICSS'36)*. IEEE.
- HUDAK, P. 1996. Building domain-specific embedded languages. *ACM Comput. Surv.* 28, 4 (Dec).
- HUDAK, P. 1998. Modular domain specific languages and tools. In *Proceedings of the 5th International Conference on Software Reuse (JCSR'98)*, P. Devanbu and J. Poulin, Eds. IEEE Computer Society, 134–142.
- JENNINGS, J. AND BEUSCHER, E. 1999. Verischemelog: Verilog embedded in Scheme. In *Proceedings of the 2nd USENIX Conference on Domain-Specific Languages*. 123–134.
- JONES, C. 1996. SPR Programming Languages Table Release 8.2, <http://www.theadvisors.com/langcomparison.htm>. (Accessed April 2005). Later release not available at publication.
- JONES, N. D., GOMARD, C. K., AND SESTOFT, P. 1993. *Partial Evaluation and Automatic Program Generation*. Prentice Hall.
- KADHIM, B. M. AND WAITE, W. M. 1996. Maptool—Supporting modular syntax development. In *Compiler Construction (CC'96)*, T. Gyimóthy, Ed. Lecture Notes in Computer Science, vol. 1060. Springer-Verlag, 268–280.
- KAMIN, S., Ed. 1997. *DSL'97—1st ACM SIGPLAN Workshop on Domain-Specific Languages in Association with POPL'97*. University of Illinois Computer Science Report.
- KAMIN, S. 1998. Research on domain-specific embedded languages and program generators. *Electro. Notes Theor. Comput. Sci.* 14. <http://www.sciencedirect.com/>.
- KAMIN, S. AND HYATT, D. 1997. A special-purpose language for picture-drawing. In *Proceedings of the USENIX Conference on Domain-Specific Languages*, 297–310.
- KANG, K. C., COHEN, S. G., HESS, J. A., NOVAK, W. E., AND PETERSON, A. S. 1990. Feature-oriented domain analysis (FODA) feasibility study. Tech. rep. CMU/SEI-90-TR-21. Software Engineering Institute, Carnegie Mellon University.
- KASTENS, U. AND PFAHLER, P. 1998. Compositional design and implementation of domain-specific languages. In *IFIP TC2 WG 2.4 Working Conference on System Implementation 2000: Languages, Methods and Tools*, R. N. Horspool, Ed. Chapman and Hall, 152–165.
- KICZALES, G., DES RIVIERES, J., AND BOBROW, D. G. 1991. *The Art of the Metaobject Protocol*. MIT Press.
- KIEBURTZ, R. B., MCKINNEY, L., BELL, J. M., HOOK, J., KOTOV, A., LEWIS, J., OLIVA, D. P., SHEARD, T., SMITH, I., AND WALTON, L. 1996. A software engineering experiment in software component generation. In *Proceedings of the 18th International Conference on Software Engineering (ICSE'96)*. IEEE, 542–552.
- KIENLE, H. M. AND MOORE, D. L. 2002. smgn: Rapid prototyping of small domain-specific languages. *J. Comput. Inform. Tech.* 10, 1, 37–53.
- KLARLUND, N. AND SCHWARTZBACH, M. 1999. A domain-specific language for regular sets of strings and trees. *IEEE Trans. Softw. Eng.* 25, 3 (May/June), 378–386.

- KRUEGER, C. W. 1992. Software reuse. *ACM Computing Surveys* 24, 2 (June), 131–183.
- KUCK, D. J. 2005. Platform 2015 software: Enabling innovation in parallelism for the next decade. *Technology@Intel Magazine*. <http://www.intel.com/technology/magazine/computing/Parallelism-0405.htm>.
- KUMAR, S., MANDELBAUM, Y., YU, X., AND LI, K. 2001. ESP: A language for programmable devices. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'01)*. ACM, 309–320.
- KUTTER, P. W., SCHWEIZER, D., AND THIELE, L. 1998. Integrating domain specific language design in the software life cycle. In *Applied Formal Methods—FM-Trends 98*, D. Hutter et al., Eds. Lecture Notes in Computer Science, vol. 1641. Springer-Verlag, 196–212.
- LAUNCHBURY, J., LEWIS, J. R., AND COOK, B. 1999. On embedding a microarchitectural design language within Haskell. *ACM SIGPLAN Notices* 34, 9 (Sept.), 60–69.
- LENGAUER, C., BATORY, D., CONSEL, C., AND ODERSKY, M., Eds. 2004. *Domain-Specific Program Generation*. Lecture Notes in Computer Science, vol. 3016. Springer-Verlag.
- LEVY, M. R. 1998. Web programming in Guide. *Softw. Pract. Exper.* 28, 1581–1603.
- MARTIN, J. 1985. *Fourth-Generation Languages*. Vol. I: Principles, Vol II: Representative 4GLs. Prentice-Hall.
- MAUW, S., WIERSMA, W., AND WILLEMSE, T. 2004. Language-driven system design. *Int. J. Softw. Eng. Knowl. Eng.* 14, 1–39.
- MERNIK, M. AND LÄMMEL, R. 2001. Special issue on domain-specific languages, Part I. *J. Comput. Inform. Techn.* 9, 4.
- MERNIK, M. AND LÄMMEL, R. 2002. Special issue on domain-specific languages, Part II. *J. Comput. Inform. Techn.* 10, 1.
- MERNIK, M., LENIĆ, M., AVDIČAUŠEVIĆ, E., AND ŽUMER, V. 2000. Multiple attribute grammar inheritance. *Informatica* 24, 3 (Sept.), 319–328.
- MERNIK, M., NOVAK, U., AVDIČAUŠEVIĆ, E., LENIĆ, M., AND ŽUMER, V. 2001. Design and implementation of simple object description language. In *Proceedings of the 2001 ACM Symposium on Applied Computing (SAC'01)*. ACM, 590–594.
- MERNIK, M., ŽUMER, V., LENIĆ, M., AND AVDIČAUŠEVIĆ, E. 1999. Implementation of multiple attribute grammar inheritance in the tool LISA. *ACM SIGPLAN Notices* 34, 6 (June), 68–75.
- MOURA, J. M. F., PÜSCHEL, M., PADUA, D., AND DONGARRA, J. 2005. Special issue on program generation, optimization, and platform adaptation. *Proceedings of the IEEE* 93, 2.
- NAKATANI, L. AND JONES, M. 1997. Jargons and infocentrism. *1st Acm SIGPLAN Workshop on Domain-Specific Languages*. 59–74. <http://www-sal.cs.uiuc.edu/kamin/dsl/papers/nakatani.ps>.
- NARDI, B. A. 1993. *A Small Matter of Programming: Perspectives on End User Computing*. MIT Press.
- NEIGHBORS, J. M. 1984. The Draco approach to constructing software from reusable components. *IEEE Trans. Softw. Eng.* SE-10, 5 (Sept.), 564–574.
- PARIGOT, D. 2004. Towards domain-driven development: The SmartTools software factory. In *Companion to the 19th Annual ACM SIGPLAN Conference on Object-oriented Programming Systems, Languages, and Applications*. ACM, 37–38.
- PEYTON JONES, S., TOLMACH, A., AND HOARE, T. 2001. Playing by the rules: Rewriting as a practical optimisation technique in GHC. In *Proceedings of the Haskell Workshop*.
- PFÄHLER, P. AND KASTENS, U. 2001. Configuring component-based specifications for domain-specific languages. In *Proceedings of the 34th Hawaii International Conference on System Sciences*.
- RAYMOND, E. S. 2001. The CML2 language: Python implementation of a constraint-based interactive configurator. In *Proceeding of the 9th International Python Conference*. 135–142. <http://www.catb.org/esr/cml2/cml2-paper.html>.
- RISI, W., MARTINEZ-LOPEZ, P., AND MARCOS, D. 2001. Hycom: A domain specific language for hypermedia application development. In *Proceedings of the 34th Hawaii International Conference on System Sciences*.
- ROSS, D. T. 1981. Origins of the APT language for automatically programmed tools. *History of Programming Languages*, R. L. Wexelblat Ed. Academic Press. 279–338.
- SALUS, P. H., Ed. 1998. *Little Languages*. Handbook of Programming Languages, vol. III. MacMillan.
- SAMMET, J. E. 1969. *Programming Languages: History and Fundamentals*. Prentice-Hall.
- SARAIVA, J. AND SCHNEIDER, S. 2003. Embedding domain specific languages in the attribute grammar formalism. In *Proceedings of the 36th Hawaii International Conference on System Sciences*.
- SCHNARR, E., HILL, M. D., AND LARUS, J. R. 2001. Facile: A language and compiler for high-performance processor simulators. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'01)*. ACM, 321–331.
- SCHNEIDER, K. A. AND CORDY, J. R. 2002. AUI: A programming language for developing plastic interactive software. In *Proceedings of the 35th Hawaii International Conference on System Sciences*.
- SCHUPP, S., GREGOR, D. P., MUSSER, D. R., AND LIU, S. 2001. User-extensible simplification—Type-based optimizer generators. In *Compiler Construction (CC'01)*, R. Wilhelm, Ed. Lecture

- Notes in Computer Science, vol. 2027. Springer-Verlag, 86–101.
- SDL FORUM. 2000. MSC-2000: Interaction for the new millenium. <http://www.sdl-forum.org/MS2000present/index.htm>.
- SIMOS, M. AND ANTHONY, J. 1998. Weaving the model web: A multi-modeling approach to concepts and features in domain engineering. In *Proceedings of the 5th International Conference on Software Reuse*. IEEE Computer Society, 94–102.
- SIRER, E. G. AND BERSHAD, B. N. 1999. Using production grammars in software testing. In *Proceedings of the 2nd USENIX Conference on Domain-Specific Languages*. 1–14.
- SLOANE, A. M. 2002. Post-design domain-specific language embedding: A case study in the software engineering domain. In *Proceedings of the 35th Hawaii International Conference on System Sciences*.
- SLONNEGER, K. AND KURTZ, B. L. 1995. *Formal Syntax and Semantics of Programming Languages: A Laboratory Based Approach*. Addison-Wesley.
- SMARAGDAKIS, Y. AND BATORY, D. 1997. DiSTiL: A transformation library for data structures. In *Proceedings of the USENIX Conference on Domain-Specific Languages*. 257–270.
- SMARAGDAKIS, Y. AND BATORY, D. 2000. Application generators. In *Wiley Encyclopedia of Electrical and Electronics Engineering Online*, J. Webster, Ed. John Wiley.
- SOROKER, D., KARASICK, M., BARTON, J., AND STREETER, D. 1997. Extension mechanisms in Montana. In *Proceedings of the 8th Israeli Conference on Computer-Based Systems and Software Engineering (ICCSSE'97)*. IEEE Computer Society, 119–128.
- SPINELLIS, D. 2001. Notable design patterns for domain-specific languages. *J. Syst. Softw.* 56, 91–99.
- SUTCLIFFE, A. AND MEHANDJIEV, N. 2004. Special issue on End-User Development. *Comm. ACM* 47, 9.
- SZYPERSKI, C. 2002. *Component Software—Beyond Object-Oriented Programming*, 2nd Ed. Addison-Wesley/ACM Press.
- TAYLOR, R. N., TRACZ, W., AND COGLIANESE, L. 1995. Software development using domain-specific software architectures. *ACM SIGSOFT Software Engineering Notes* 20, 5, 27–37.
- TENNENT, R. D. 1977. Language design methods based on semantic principles. *Acta Inf.* 8, 97–112.
- THATTE, S. 2001. XLANG: Web services for business process design. Tech. rep. Microsoft. <http://www.getdotnet.com/team/xml.wsspecs/xlang-c/>.
- THIBAUT, S. A. 1998. Domain-specific languages: Conception, implementation and application. Ph.D. thesis, University of Rennes.
- THIBAUT, S. A., CONSEL, C., AND MULLER, G. 1998. Safe and efficient active network programming. In *Proceedings of the 17th IEEE Symposium on Reliable Distributed Systems*. IEEE Computer Society, 135–143.
- THIBAUT, S. A., MARLET, R., AND CONSEL, C. 1999. Domain-specific languages: From design to implementation—Application to video device drivers generation. *IEEE Trans. Softw. Eng.* 25, 3, (May/June), 363–377.
- TRACZ, W. AND COGLIANESE, L. 1995. DOMAIN (Domain Model All INtegrated)—a DSSA domain analysis tool. Tech. rep. ADAGE-LOR-94-11. Loral Federal Systems.
- UPnP 2003. Universal Plug and Play Forum. <http://www.upnp.org/>.
- USENIX 1997. *Proceedings of the USENIX Conference on Domain-Specific Languages*.
- USENIX 1999. *Proceedings of the 2nd USENIX Conference on Domain-Specific Languages (DSL99)*.
- VAN DEN BRAND, M. G. J., VAN DEURSEN, A., HEERING, J., DE JONG, H. A., DE JONGE, M., KUIPERS, T., KLINT, P., MOONEN, L., OLIVER, P. A., SCHEERDER, J., VINJU, J. J., VISSER, E., AND VISSER, J. 2001. The ASF+SDF Meta-Environment: A component-based language development environment. In *Compiler Construction (CC'01)*, R. Wilhelm, Ed. Lecture Notes in Computer Science, vol. 2027. Springer-Verlag, 365–370. <http://www.cwi.nl/projects/MetaEnv>.
- VAN DEN BRAND, M. G. J. AND VISSER, E. 1996. Generation of formatters for context-free languages. *ACM Trans. Softw. Eng. Method.* 5, 1–41.
- VAN DEURSEN, A. AND KLINT, P. 1998. Little languages: Little maintenance? *J. Softw. Maintenance* 10, 75–92.
- VAN DEURSEN, A. AND KLINT, P. 2002. Domain-specific language design requires feature descriptions. *J. Comput. Inform. Tech.* 10, 1, 1–17.
- VAN DEURSEN, A., KLINT, P., AND VISSER, J. 2000. Domain-specific languages: An annotated bibliography. *ACM SIGPLAN Notices* 35, 6 (June), 26–36.
- VAN ENGELEN, R. 2001. ATMOL: A domain-specific language for atmospheric modeling. *J. Comput. Inform. Techn.* 9, 4, 289–303.
- VELDHUIZEN, T. L. 1995a. Expression templates. *C++ Report* 7, 5 (June) 26–31.
- VELDHUIZEN, T. L. 1995b. Using C++ template metaprograms. *C++ Report* 7, 4 (May) 36–43.
- VELDHUIZEN, T. L. 2001. Blitz++ User's Guide. Version 1.2 <http://www.oonumerics.org/blitz/manual/blitz.ps>.
- VISSER, E. 2003. Stratego—Strategies for program transformation. <http://www.stratego-language.org>.
- WANG, D. C., APPEL, A. W., KORN, J. L., AND SERRA, C. S. 1997. The Zephyr abstract syntax description language. In *Proceedings of the USENIX Conference on Domain-Specific Languages*, 213–28.

- WEISS, D. AND LAY, C. T. R. 1999. *Software Product Line Engineering*. Addison-Wesley.
- WEXELBLAT, R. L., Ed. 1981. *History of Programming Languages*. Academic Press.
- WILE, D. S. 1993. *POPART: Producer of Parsers and Related Tools*. USC/Information Sciences Institute. <http://mr.tekknowledge.com/wile/popart.html>.
- WILE, D. S. 2001. Supporting the DSL spectrum. *J. Comput. Inform. Techn.* 9, 4, 263–287.
- WILE, D. S. 2004. Lessons learned from real DSL experiments. *Sci. Comput. Program.* 51, 265–290.
- WILE, D. S. AND RAMMING, J. C. 1999. Special issue on Domain-Specific Languages. *IEEE Trans. Softw. Eng.* SE-25, 3 (May/June).
- XIONG, J., JOHNSON, J., JOHNSON, R. W., AND PADUA, D. A. 2001. SPL: A language and compiler for DSP algorithms. In *Proceedings of the 2001 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI01)*. ACM, 298–308.

Received September 2003; revised May 2005; accepted December 2005

Copyright of *ACM Computing Surveys* is the property of *Association for Computing Machinery* and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.