

Special Issue on Domain-Specific Languages Guest Editors' Introduction

Ralf Lämmel¹ and Marjan Mernik²

¹ Free University Amsterdam, The Netherlands

² University of Maribor, Slovenia

15 December, 2001

In need of domain support

Programming languages are a programmer's most basic tools. A language is suitable for a programming problem if it makes the programmer productive, and if it allows the programmer to write highly scalable, generic, readable and maintainable code. There are various ways to classify or to group programming languages. One can, for example, focus on the programming paradigm in such a classification (cf. imperative vs. functional vs. logic vs. constraint vs. object-oriented etc., or multi-paradigm). One can also consider 2GLs, 3GLs, 4GLs, or very high-level languages. We want to emphasize the division of programming languages into general-purpose and domain-specific languages. Here, we admit that there is a gray area hosting languages which neither belong strictly to the general-purpose nor to the domain-specific group. With general-purpose languages, one can address large classes of problems e.g., scientific computing, business processing, symbolic processing while a domain-specific language (DSL) facilitates the solution of problems in a particular domain. To this end, a DSL provides built-in abstractions and notations specific to the domain of concern. In one domain, partial differential equations might be of central importance, and hence they deserve the status of a proper abstraction. In many domains, graphical notations are common, and hence a DSL that is suitable for such a domain would be a visual language. DSLs are usually small, more declarative than imperative. The poor man's approach to DSL implementation is

to implement a library to be used via an API in the general-purpose language of choice. This approach is not suitable when the domain requires:

- dedicated notation, that is, syntax,
- dedicated abstraction mechanisms,
- dedicated scoping or typing rules,
- domain-specific optimizations,
- domain-specific analyses,
- domain-specific error reporting,
- etc.

A historical note

DSLs have been used in various domains such as robot control, animation, music composition, financial product design, description and analysis of abstract syntax trees, web computing, 3D animation, reverse engineering, modelling of reactive systems, just to mention a few. The idea of DSLs is presumably as old as the notion of programming languages. The APT language (Automatically Programmed Tools), developed in the 1955 at MIT, can be regarded as a first DSL. One of the first research groups systematically working on DSLs was N. J. Lehmann's group at the Technical University of Dresden in East-Germany. This group started in the nineteen-seventies to pursue fundamental research on the field, and to develop DSLs commissioned by

industry on a regular basis. At this time, one did not use the term DSL but “specialized language” or “application-oriented language”, and in German “Fachsprache” or “Spezialsprache”.

Trade-offs

Applications of DSLs have clearly illustrated the advantages of DSLs over general-purpose languages in areas such as productivity, reliability, maintainability and flexibility. However, the benefits of DSLs are not for free. Without appropriate methodology and tools, the upfront investment for DSL support can exceed the savings obtained by using a DSL for application development. Since the costs of DSL language development and maintenance have to be taken into account, one of the main questions is “When and how to design and implement a domain-specific language?”. In similarity to general software development, one can identify the following phases for the development of a DSL: analysis, design, implementation, and finally their deployment. In the analysis phase, the problem domain is identified, and domain knowledge has to be gathered. Then, the DSL is designed to concisely describe applications in the domain. The implementation phase can be done using one of the following approaches:

- A proper compiler or interpreter is developed where standard compiler tools can be used, or tools dedicated to the implementation of DSLs.
- The DSL is modelled as an embedded language or as a domain-specific library where some form of abstraction that is available in an existing language (e.g., functions in C or Haskell) is employed to capture domain-specific operations.
- The DSL is supported via preprocessing or macro processing where DSL constructs are translated to statements in a base language.
- The DSL is implemented by means of an extensible compiler or interpreter, that is, the DSL constructs are added to the existing language implementation, e.g., via reflection.

The above steps show that the development of a domain-specific language is itself a significant software engineering effort, requiring

considerable investment of time and resources. In practice, we have to seek for a trade-off between level of DSL support, usability, efficiency, evolvability of the implementation, available resources, and other criteria. The interpretation/compilation approach from above, for example, is very much suited to achieve a full-blown implementation of a DSL. One might employ simple parser generators, or more sophisticated compilers based on executable language definitions—as common for the implementation of general-purpose languages, too. The other approaches listed above (embedding, preprocessing, extensible compiler/interpreter) can be more efficient and attractive in particular cases of DSL development.

Challenges in the DSL field

Without restricting ourselves to DSLs, one can say that language design and implementation are apparently never-ending research activities. If we, for example, look back at the last ten years, we see that major breakthroughs in the field of *modular* language definition and implementation were achieved (cf. modular SOS, action semantics, abstract state machines, use of object-oriented methods and generic programming). Further progress of the DSL field can be expected from these results. The now more and more ubiquitous role of computers, their use in all domains of business and every-day life implies, without any doubt, that DSLs are even more frequently needed in future. End-users need to be enabled to write programs in the notations which are most suitable in a given application domain. Research in the DSL field will hence focus on methods and tools to simplify and to discipline the development of DSLs. Besides these general conceptual challenges, the development of actual DSLs for emerging domains forms a continuous research activity in the field. Finally, we expect an amalgamation of the DSL theme and other research topics, e.g., the more recent field of aspect-oriented programming.

Papers in the special issue

The papers accepted for the CIT special issue on domain-specific languages provide the reader with a contemporary analysis of the spectrum of DSL approaches, with concepts for DSL design, with new shining examples of DSLs,

and with discussions of lightweight vs. heavy-weight tools, or language support for DSL development. In this manner, the selected papers contribute to the answer of the question “When and how to develop a DSL?”. The papers were selected from eleven submissions. There are two invited contributions, and five regular papers. For reasons of page count, the special issue consists of two parts which appear in two consecutive CIT issues. We first list all the papers appearing in the special issue, and then we will shortly introduce the papers included in the present CIT issue.

- D. Wile. Supporting the DSL Spectrum (invited paper)
- A. van Deursen, P. Klint. Domain-Specific Language Design Requires Feature Descriptions (invited paper)
- R. van Engelen. ATMOL: A Domain-Specific Language for Atmospheric Modeling
- A. Berlea, H. Seidl. *fxt* - A Transformation Language for XML Documents
- J. Gil, Y. Tsoglin. *Jamoos* - A Domain-Specific Language for Language Processing
- H. Kienle, D. Moore. *sgmm*: Rapid Prototyping of Small Domain-Specific Languages
- J. Aycock. The Design and Implementation of SPARK, a Toolkit for Implementing Domain-Specific Languages

Part I

Dave S. Wile’s invited contribution provides a splendid overview on DSL approaches using satellite ground controllers as a running example. He covers approaches as different as a graphical approach based on Microsoft PowerPoint vs. a combinator library in Haskell. The author looks back on years of experience in the field of DSL design and tool support. He pursued corresponding research at Tecknowledge Corporation, and before that at the University of Southern California.

Robert A. van Engelen describes the design and implementation of ATMOL: a DSL for the formulation and implementation of atmospheric

models. Such models are based on partial differential equations. A dedicated code generator produces highly efficient FORTRAN code from the DSL programs. The DSL ATMOL, together with the dedicated tool support, forms a prime example of a DSL, because it illustrates best how a DSL can offer very high level specifications which are transformed into efficient code relying on domain-specific optimizations.

Joseph Gil and Yuri Tsoglin outline the language *Jamoos* which can be seen in two ways. One might say that *Jamoos* is a DSL for language processing. One might also say that *Jamoos* is a distinguished object-oriented programming language centred around a tree computing metaphor. In fact, the authors outline the result of a major language design effort, and relate it to general-purpose languages, and other frameworks for the implementation of language processors. Their design involves a number of original features, e.g., object immutability as for the constructed trees, or a generalized call stack regime.

Acknowledgement

We want to thank Arie van Deursen, Paul Klint and Dave S. Wile for their brilliant invited contributions to the special issue. Another important success factor in the project were the referees. Our thanks go to Mikhail Auguston, Dmitry Boulytchev, Mark van den Brand, Kyung-Goo Doh, Gopal Gupta, Gorel Hedin, Jan Heering, Sam Kamin, Günter Kniesel, Tobias Kuipers, Wolfgang Lohmann, Katharina Mehner, Pierre-Etienne Moreau, Günter Robert, Anthony Sloane, Andrey A. Terekhov, Eelco Visser, and Eric Van Wyk. Last but not least, we are grateful for the support by the publisher for this special issue. We appreciate the very smooth cooperation with the CIT editor prof. Sven Lončarić and with Vesna Smoljanović from the CIT Editorial Office.

Ralf Lämmel
Free University Amsterdam
The Netherlands
e-mail: Ralf.Laemmel@cs.vu.nl

Marjan Mernik
University of Maribor
Slovenia
e-mail: marjan.mernik@uni-mb.si