

# We should Stop Claiming Generality in our Domain-Specific Language Papers

Daco C. Harkes<sup>a</sup>

<sup>a</sup> Delft University of Technology, The Netherlands

**Abstract** Our community believes that new domain-specific languages should be as general as possible to increase their impact. However, I argue in this essay that we should stop claiming generality for new domain-specific languages. More general domain-specific languages induce more boilerplate code. Moreover, domain-specific languages are co-developed with their applications in practice, and tend to be specific for these applications. Thus, I argue we should stop claiming generality in favor of documenting how domain-specific language based software development is beneficial to the overall software development process. The acceptance criteria for scientific literature should make the same shift: accepting good domain-specific language engineering practice, instead of the next language to rule them all.

ACM CCS 2012

- **Software and its engineering** → *Domain specific languages; Software development methods;*
- **General and reference** → *General literature;*

**Keywords** Domain-specific languages, domain-specific language engineering, scientific rhetoric

## The Art, Science, and Engineering of Programming

---

Submitted April 23, 2018

Published October 1, 2018

doi 10.22152/programming-journal.org/2019/3/5



© Daco C. Harkes

This work is licensed under a “CC BY 4.0” license.

In *The Art, Science, and Engineering of Programming*, vol. 3, Essays, 2019, article 5; 10 pages.

## We should Stop Claiming Generality in our DSL Papers

### 1 Introduction

Domain-specific languages (DSLs) are thriving in academia and in industry. Our scientific community publishes on average 1.1 articles per day with “domain-specific language” in the title or abstract. At the time of writing Google Scholar lists 418 such publications in the last 365 days. In industry DSL based software engineering is growing rapidly as well. This is illustrated both by the number of DSL talks by companies at academic and industry IT events, and by the bold slogans these companies use: “low-code platform”, “auto generated <something>”, “automate production of software”, and “deliver products 5-10 faster with <DSL-tool>”.

Despite advice to not over-generalize [12], scientific arguments for new DSLs in our community usually include an argument for generality. One of the criteria we evaluate DSL research on is its potential impact; and for a DSL to have a large impact, it should be general enough to cover most problems of a domain. In other words, a criterion for accepting DSL research is that a DSL should be widely used potentially.

My goal in this essay is to argue against this criterion, because DSLs can be very useful without generality. This criterion directs research effort in a direction which is not necessarily fruitful. DSLs can become less useful as they become more general, and DSLs which are useful without being general might not be generalizable at all. If despite this, DSL authors still make an argument for generality, it is to be expected that these arguments are weak. They are usually based on anecdotal evidence. Authors list some examples or case studies they could come up with, and which can be expressed in their DSL. Obvious examples which cannot be expressed in the DSL can go to the future work section, but it is best not to list non-obvious limitations because the DSL generality perceived by the reviewers will diminish.

This often leads to claims of generality which are undone by later work. For example a paper stating “we believe that our approach handles many lexically-scoped languages” [13] is followed-up with a paper stating “We also refine and extend our previous theory to address practical concerns including ambiguity checking and support for a wider range of scope relationships” [19]. Similarly, the “optimal-time” [14] incremental evaluation for attribute grammars turned out to be not generally applicable and was improved with “Asymptotic Performance of Symbol Table Modifications” [9]. These weak and false claims for generality leave an impression that work can solve everything. This is harmful for newcomers and dissemination of knowledge to industry as the technical details of a DSL need to be understood before its applicability to a specific problem can be assessed.

We should stop these weak arguments for generality. If authors specifically want to argue for generality, they should have a good methodology for assessing generality. For all other DSL contributions we should stop evaluating generality. Instead, we should report on the whole software engineering process in which DSLs are developed and used, measuring and detailing how DSLs improve the software engineering process and the resulting software.

## 2 More General DSLs versus more Specific DSLs

First, let us look at generality from a theoretical point of view. Suppose a program can be expressed in a more general DSL and in a more specific DSL, then the program will be more concise in the more specific DSL. A more specific DSL makes assumptions which programs cannot work around, while a more general DSL does not make these assumptions. This means that the more specific DSL cannot express as many programs. However, it also means that programs which do require these assumptions do need to make them explicit in the more general DSL, increasing boilerplate code.

An example of this trade-off between scrap-your-boilerplate and generality is visible in the domain of name binding (figure 1). A very specific DSL for name binding is NaBL [22]. It enables very concise specifications of a small set of static analysis concepts in which scopes are implicit. Scope Graphs [13, 19] make scopes explicit, which enables a wider range of name resolution concepts, but also makes trivial name resolution more verbose. Scope Graphs leave implicit how names flow through scopes. On the generic side we have attribute grammar (AG) based [4] and Datalog-based [16] name resolution which both directly encode how names flow through scopes by means of key-value maps or relations. This enables almost arbitrary name resolution patterns, but increases boilerplate.

It is not necessarily better to have a more general DSL. It can be worth the engineering effort to create a specific DSL for a reasonably complicated program. (Note that it is also not necessarily better to have a more specific DSL, for a relatively simple program the extra overhead of creating and maintaining a more specific DSL might be not worth it.)

## 3 Co-development of DSLs and Software

This trade-off between scrap-your-boilerplate and generality tends to favor the former in practice. DSLs are not designed to be general in practice because most DSLs are developed in tandem with the programs written in them. This is apparent in scientific literature and in the projects using DSLs in academia and industry alike.

### 3.1 Every DSL Paper has its own Running Example

Recently, I had the privilege to work on a survey for incremental computing DSLs. Every paper has its own running example, which can either not be expressed in previous work, or cannot be properly incrementally executed in previous work. These running examples are carefully chosen to exercise the key features of the DSL at hand. Moreover, when research branches out, the running examples do not work in other branches. For example, classical attribute grammars [14, 23, 24] are extended in aggregate AGs [9], remote AGs [1], reference AGs (RAGs) [17], and various higher-order AGs [18, 20, 15, 2, 3], but these extensions do not support each others running examples while most authors claim it is a general framework for building compilers. (One such example is that aggregate AGs (1986) [9] optimize edits on map data

## We should Stop Claiming Generality in our DSL Papers

```
Block(_): scopes variable
VarDecl(t, x, _): defines variable x of type t
VarRef(x): refers to variable x
VarRef(x) : t
where definition of x : t
```

```
[[ Block(body) ^ (outer_scope) ]] :=
  new block_scope,
  block_scope -P-> outer_scope,
  [[ body ^ (block_scope) ]].

[[ VarDecl(t, x, _) ^ (scope) ]] :=
  Variable{x} <- scope,
  Variable{x} : t.

[[ VarRef(x) ^ (scope) : x_type ]] :=
  Variable{x} -> scope,
  Variable{x} |-> x_def,
  x_def : x_type.
```

```
inh Variable Name.lookupVariable(String name);
eq Block.getStmt(int index).lookupVariable(String name){
  for(int i = 0; i < index; i++)
    if(getStmt(i).isVarDecl(name))
      return (VarDecl)getStmt(i);
  return lookupVariable(name);
}

syn boolean Stmt.isVarDecl(String name) = false;
eq VarDecl.isVarDecl(String name) =
  name().equals(name);

syn ClassDecl Expr.type() = unknownType();
eq VarRef.type() = lookupVariable(name()) != null ?
  lookupVariable(name()).type() : unknownType();
```

- **Figure 1** Trade-off between specificity and generality in name binding DSLs. All snippets define how variable references get resolved to variable definitions, and how blocks scope these variables. (Top) NaBL: A very specific DSL for name binding [22]. It enables very concise specifications of a small set of static analysis concepts in which scopes are implicit. (Middle) NaBL2: Scope Graphs make scopes explicit [13, 19], which enables a wider range of name resolution concepts, but also makes trivial name resolution more verbose. NaBL2 leaves implicit how names flow through scopes. (Bottom) JastAdd: Attribute grammar based name resolution which directly encode how names flow through scopes by means of inherited and synthesized attributes [4]. This enables almost arbitrary name resolution patterns.

structures, while RAGs (2012) state as future work addressing this exact problem: “Some practical RAGs, such as JastAddJ, make use of large attribute values for collecting local declarations into a map data structure. [...] This use of large values causes the same kind of dependency imprecision that ordinary AGs suffer from” [17].)

It is conceivable that these DSL authors wanted to build a piece of software (a compiler in this case), and that a better DSL was the way to do it. The piece of software they wanted to build ended up as the running example in their publication. The software was the driving factor behind the DSL research, and the DSL was the means to build this software. Thus, DSLs were co-developed with the running examples written in them.

### 3.2 DSL Projects in Practice

Co-development is also what I see in projects using DSLs, both in academia and industry. In our building, a group of scientific programmers builds web applications with WebDSL [6]. These programmers are also the maintainers of WebDSL. They add new features to WebDSL, or fix bugs, when this is required for one of the applications. In the same way the additions to IceDust [7] in IceDust2 [8] were required to support a specific application.

Also in industry co-development is the pattern. The Dutch tax office builds the Agile Law Execution Factory (Alef) DSL, with which they build the tax declaration application.<sup>1</sup> Similarly, AFAS builds its customer relation management (CRM) system with a DSL, and when they cannot express something they extend their DSL.<sup>2</sup> M-Industries builds their client applications with Alan [11], and whenever they cannot build something they add another sub language to their family of Alan languages. Finally, the use of mbeddr to build a Smart Meter also included creating new language extensions [21]. These are all typical examples of this pattern I observe around me.

We do not believe in the waterfall process anymore for general software development. Neither should we for DSL-based software development. DSL-based development is iterative, just as software development is. And as requirements and application are co-developed in the iterative process, so are DSLs and applications in the DSL-based iterative process.

## 4 Catering for Co-development

If DSLs (specific and general) are co-developed with applications, we as a research community should support and research that process. In our field we currently have two research directions already supporting this process.

<sup>1</sup> Presented at the LangDevCon in 2018 (<http://langdevcon.org/>).

<sup>2</sup> Presented at the Dutch Software Engineering Symposium in 2016 (<http://www.sen-symposium.nl/history/2016/program/>).

## We should Stop Claiming Generality in our DSL Papers

Extensible languages are one way to cater for co-development of language and application. Extensible language developers expect that their DSL is not a perfect fit for their DSL users. These extensible languages can be customized for specific applications. As mentioned, mbeddr is such an extensible language, and during the Smart Meter development new extensions were added [21].

Another way to cater for co-development is to reduce the development effort of DSLs in general. Language Workbenches aim to do this [5]. In order to provide a seamless co-development cycle of language and application changes in the language should directly be visible in the application. This is the research of interactive or live language workbenches [10].

### 5 Scientific Rhetoric Shift

But more importantly, we need a shift in the rhetoric of our scientific DSL community. We should stop claiming generality, and as reviewers we should stop judging based on the criterion of generality. If we drop the criterion of generality, we might be able to do research in directions which can be more fruitful. DSL authors can stop spending effort on creating weak arguments of why their DSL is general, and reviewers do not have to spend effort on reading those arguments anymore.

Instead, we should start claiming co-development is beneficial, and as reviewers we should start judging on benefit in the software engineering process. Thus we need to report on this co-development process with case studies. Real case studies of the software engineering process are costly, they take enormous effort. Maybe we should, as a community, learn from the medical community, who publish medical trials (case studies) without drawing conclusions or claiming anything. Data on the DSL-based software engineering process should be valuable in itself. If the data is valuable in itself, there is no need to claim in every DSL paper that this DSL is better than all previous DSLs.

Let us be honest, and improve our methodology and rhetoric in DSL research. To inspire us to do that, I want to draw attention to a great example. The case study by Völter et al. of mbeddr use in the Smart Meter project details precisely what DSL use brought to the Smart Meter project: “We find that the extensions help significantly with managing the complexity of the software” [21]. Moreover, they detail how they assessed this benefit: “We look at the mbeddr extensions used[...] We qualitatively assess their impact on the complexity of the implementation” [21]. Finally, it is explicit about conclusions which cannot be drawn based on that case study. For example, they cannot compare development effort with a non-DSL Smart Meter, as they did not build a non-DSL version. I would like to see more such research papers.

I envision that with such research papers we, as a community, can improve the DSL-based software engineering process. And, that in 10 years from now, DSLs are as wide-spread as libraries and frameworks are now.

Note that if DSL authors do write papers with a heavy emphasis on language design instead of benefit to the overall software engineering process, then reviewers should reward authors for being explicit about the limits to the applicability of their DSLs.

Because it is hard to establish the applicability of DSLs (authors cannot do multiple industry-size case studies to gather applicability evidence for a single language design paper), I propose to be explicit about the applicability uncertainty by stating lower bound as well as the upper bound: “Our DSL can do X (we did a case study), it cannot do Y because of Z, and we honestly do not know whether it can do A, B, and C”. Being explicit about the applicability, will help newcomers to the field and industry alike.

## 6 Conclusion

DSLs are co-developed with the applications that are written in them. Co-developed DSLs are usually more useful when they are specific instead of general. Thus, our community should drop the generality criterion when evaluating DSL research contributions. Instead, we should report more on how DSL-based development is beneficial to the overall development process. If we keep evaluating DSLs based on their generality, we will see many more weak and false arguments for specific DSLs claiming generality. This is not good for our research community. I hold that we should embrace co-development of DSLs with their applications. Perhaps you disagree. If you do, I hope to read your counterargument in a published essay soon.

**Acknowledgements** I thank Hendrik van Antwerpen and Sebastian Erdweg for their valuable input to this essay. Also, I thank the anonymous reviewers for their constructive feedback.

## References

- [1] John Boyland. Incremental evaluators for remote attribute grammars. *Electronic Notes in Theoretical Computer Science*, 65(3):9–29, 2002. doi:10.1016/S1571-0661(04)80424-9.
- [2] Jeroen Bransen. *On the Incremental Evaluation of Higher-Order Attribute Grammars*. PhD thesis, Utrecht University, Netherlands, 2015.
- [3] Jeroen Bransen, Atze Dijkstra, and S. Doaitse Swierstra. Incremental evaluation of higher-order attributes. *Science of Computer Programming*, 137:98–124, 2017. doi:10.1016/j.scico.2016.06.001.
- [4] Torbjörn Ekman and Görel Hedin. Modular name analysis for Java using JastAdd. In Ralf Lämmel, João Saraiva, and Joost Visser, editors, *Generative and Transformational Techniques in Software Engineering, International Summer School, GTTSE 2005, Braga, Portugal, July 4-8, 2005. Revised Papers*, volume 4143 of *Lecture Notes in Computer Science*, pages 422–436. Springer, 2006. doi:10.1007/11877028\_18.
- [5] Martin Fowler. Language workbenches: The killer-app for domain specific languages? <http://www.martinfowler.com/articles/languageWorkbench.html>, 2005. Accessed: 2018-08-08.

## We should Stop Claiming Generality in our DSL Papers

- [6] Danny M. Groenewegen, Zef Hemel, Lennart C. L. Kats, and Eelco Visser. Webdsl: a domain-specific language for dynamic web applications. In Gail E. Harris, editor, *Companion to the 23rd Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2008, October 19-13, 2007, Nashville, TN, USA*, pages 779–780. ACM, 2008. doi:10.1145/1449814.1449858.
- [7] Daco Harkes, Danny M. Groenewegen, and Eelco Visser. Icedust: Incremental and eventual computation of derived values in persistent object graphs. In Shriram Krishnamurthi and Benjamin S. Lerner, editors, *30th European Conference on Object-Oriented Programming, ECOOP 2016, July 18-22, 2016, Rome, Italy*, volume 56 of *LIPICs*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016. doi:10.4230/LIPICs.ECOOP.2016.11.
- [8] Daco Harkes and Eelco Visser. Icedust 2: Derived bidirectional relations and calculation strategy composition. In Peter Müller 0001, editor, *31st European Conference on Object-Oriented Programming, ECOOP 2017, June 19-23, 2017, Barcelona, Spain*, volume 74 of *LIPICs*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017. doi:10.4230/LIPICs.ECOOP.2017.14.
- [9] Roger Hoover and Tim Teitelbaum. Efficient incremental evaluation of aggregate values in attribute grammars. In *Proceedings of the 1986 SIGPLAN Symposium on Compiler Construction, Palo Alto, California, USA, June 25-27, 1986*, pages 39–50. ACM, 1986. doi:10.1145/12276.13315.
- [10] Gabriël Konat, Michael J. Steindorfer, Sebastian Erdweg, and Eelco Visser. Pie: A domain-specific language for interactive software development pipelines. *The Art, Science, and Engineering of Programming*, 2(3), 2018. doi:10.22152/programming-journal.org/2018/2/9.
- [11] GJ Kunst. Alanlight: sound, functionally correct, bounded acyclic data flow modeling. Master’s thesis, Delft University of Technology, 2018.
- [12] Marjan Mernik, Jan Heering, and Anthony M. Sloane. When and how to develop domain-specific languages. *ACM Computing Surveys*, 37(4):316–344, 2005. doi:10.1145/1118890.1118892.
- [13] Pierre Néron, Andrew P. Tolmach, Eelco Visser, and Guido Wachsmuth. A theory of name resolution. In Jan Vitek, editor, *Programming Languages and Systems - 24th European Symposium on Programming, ESOP 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, volume 9032 of *Lecture Notes in Computer Science*, pages 205–231. Springer, 2015. doi:10.1007/978-3-662-46669-8\_9.
- [14] Thomas W. Reps. Optimal-time incremental semantic analysis for syntax-directed editors. In *POPL*, pages 169–176, 1982. doi:10.1145/582153.582172.
- [15] João Saraiva, S. Doaitse Swierstra, and Matthijs F. Kuiper. Functional incremental attribute evaluation. In David A. Watt, editor, *Compiler Construction, 9th International Conference, CC 2000, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS 2000, Berlin, Germany, Arch 25 - April 2, 2000, Proceedings*, volume 1781 of *Lecture Notes in Computer Science*, pages 279–294. Springer, 2000. doi:10.1007/3-540-46423-9\_19.



- [16] Yannis Smaragdakis and Martin Bravenboer. Using datalog for fast and easy program analysis. In Oege de Moor, Georg Gottlob, Tim Furche, and Andrew Jon Sellers, editors, *Datalog Reloaded - First International Workshop, Datalog 2010, Oxford, UK, March 16-19, 2010. Revised Selected Papers*, volume 6702 of *Lecture Notes in Computer Science*, pages 245–251. Springer, 2010. doi:10.1007/978-3-642-24206-9\_14.
- [17] Emma Söderberg and Görel Hedin. Incremental evaluation of reference attribute grammars using dynamic dependency tracking. Technical Report 98, Department of Computer Science, Lund University, 2012.
- [18] Tim Teitelbaum and Richard Chapman. Higher-order attribute grammars and editing environments. In *PLDI*, pages 197–208, 1990. doi:10.1145/93548.93567.
- [19] Hendrik van Antwerpen, Pierre Néron, Andrew P. Tolmach, Eelco Visser, and Guido Wachsmuth. A constraint language for static semantic analysis based on scope graphs. In Martin Erwig and Tiark Rompf, editors, *Proceedings of the 2016 ACM SIGPLAN Workshop on Partial Evaluation and Program Manipulation, PEPM 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 49–60. ACM, 2016. doi:10.1145/2847538.2847543.
- [20] Harald Vogt, S. Doaitse Swierstra, and Matthijs F. Kuiper. Efficient incremental evaluation of higher order attribute grammars. In *PLILP*, pages 231–242, 1991. doi:10.1007/3-540-54444-5\_102.
- [21] Markus Völter, Arie van Deursen, Bernd Kolb, and Stephan Eberle. Using c language extensions for developing embedded software: a case study. In Jonathan Aldrich and Patrick Eugster, editors, *Proceedings of the 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2015, part of SPLASH 2015, Pittsburgh, PA, USA, October 25-30, 2015*, pages 655–674. ACM, 2015. doi:10.1145/2814270.2814276.
- [22] Guido Wachsmuth, Gabriël Konat, Vlad A. Vergu, Danny M. Groenewegen, and Eelco Visser. A language independent task engine for incremental name and type analysis. In Martin Erwig, Richard F. Paige, and Eric Van Wyk, editors, *Software Language Engineering - 6th International Conference, SLE 2013, Indianapolis, IN, USA, October 26-28, 2013. Proceedings*, volume 8225 of *Lecture Notes in Computer Science*, pages 260–280. Springer, 2013. doi:10.1007/978-3-319-02654-1\_15.
- [23] Dashing Yeh. On incremental evaluation of ordered attribute grammars. *BIT*, 23(3):308–320, 1983. doi:10.1007/BF01934460.
- [24] Dashing Yeh and Uwe Kastens. Improvements of an incremental evaluation algorithm for ordered attribute grammars. *SIGPLAN Notices*, 23(12):45–50, 1988. doi:10.1145/57669.57672.

## We should Stop Claiming Generality in our DSL Papers

### About the author

**Daco C. Harkes** is a PhD candidate at Delft University of Technology, where he works on domain-specific languages. His current work includes the domain-specific language IceDust, an incremental computing language for the domain of information systems. You can contact him at [d.c.harkes@tudelft.nl](mailto:d.c.harkes@tudelft.nl).

