

Lessons from the JML Project

Gary T. Leavens¹ and Curtis Clifton^{1**}

Department of Computer Science, Iowa State University,
226 Atanasoff Hall, Ames, IA 50011 USA

Abstract. To have impact, a grand challenge should provide a way for diverse research to be integrated in a synergistic fashion. Synergy in the JML project comes from a shared specification language, and thus holds several lessons for the verifying compiler grand challenge. An important lesson is that the project should focus considerable resources on specification language design, which still contains many open research problems. Another important lesson is that, to support such a specification language, the project needs to involve groups doing research on extensible compilers and integrated development environments.

1 Introduction

Hoare's verifying compiler grand challenge is "the construction and application of a verifying compiler that guarantees correctness of a program before running it" [1, p. 63]. This challenge is of such a broad scope that a project meeting the challenge would involve "a significant section of the research community" that would "work together towards a common goal, agreed to be valuable and achievable by a team effort" [1, p. 63].

This position paper focuses on one way that these researchers could work together, drawing lessons from the experience of the Java Modeling Language (JML) project [2,3]. These lessons are relevant because the JML project, although smaller and less ambitious, has many parallels with the verifying compiler grand challenge.

2 What Kind of Specifications?

To verify a program, one must have a specification. The specification can be implicit, such as that the program should not encounter unexpected exceptions due to obvious program errors (such as dereferencing a null pointer or indexing an array beyond its bounds). But finding such bugs is possible with existing tools, such as ESC/Java [4]. While such tools are a subject of current research and engineering, they hardly constitute a grand challenge at this point. Therefore, implicit in the grand challenge are interesting specifications, written in some specification language. Examples of interesting specifications include specifications of functional behavior that involves data values and specifications of safety properties that describe synchronization of concurrent threads.

^{**} Current address for Curtis Clifton: Dept. of Comp. Sci. and Soft. Eng., Rose-Hulman Inst. of Technology, 5500 Wabash Ave., Terre Haute, IN 47803.

The main cost of such interesting specifications is that, in general, they cannot be automatically generated. Instead, they must be written, to some extent, by humans. The reason human input is needed is that only humans can judge the intent of a specification. For example, consider a square root routine. A particular implementation may produce roots with 7 decimal digits of accuracy, but only a human can decide if the intended behavior is 7 digit accuracy, rather than 5 or 10 digit accuracy. Put another way, interesting specifications describe the set of all acceptable implementations abstractly; the intent behind this description allows them to be used as contracts that govern future evolution of both implementations and clients [5]. Ultimately these contracts are a matter for human judgment and negotiation.

3 Why Design a Specification Language?

Strictly speaking, the grand challenge can proceed with only a few interesting specifications. For example, a group of experts might specify the Linux kernel or write a few other interesting specifications as tests of verification technology. If the focus of the project is solely on proving programs correct, then a small set of such specifications would be adequate, and the costs of writing the specifications and designing specification languages could be largely ignored by the project.

However, we believe that including specification language design is necessary to maximize the project's impact, and would also have several other benefits.

The relationship between the project's potential impact and including specification language design in its scope can be seen by considering the alternative. Suppose that the project only works with a small set of test specifications, and does not provide an easily-usable, well-documented specification language. Assuming that the project succeeds, then how can programmers apply its verification technology to code that does not implement the test specifications? In this scenario, such applications might still be very costly, as programmers would have to write new interesting specifications, which we are assuming would be hard. Hence many programming projects would not be able to cost-effectively use the new verification technology.

Another benefit, not to be overlooked, is that a specification language is a good way to coordinate efforts among different verification tools. This is one of the main lessons of the JML project, which has been fairly successful in coordinating the efforts of diverse research groups [2]. Having JML as a common specification language allows these groups to also share users, and thus have a larger pool of users to test their ideas and to obtain feedback. It also facilitates the exchange of ideas among research groups.

Therefore, we believe that one of the project's overall goals should be to make it easy (inexpensive) and valuable (cost-effective) for ordinary programmers to write interesting formal specifications. One milestone would be to replace most informal documentation with formal specifications, while decreasing the overall cost of program development and maintenance, since at that point formal specifications become economically attractive.

Achieving such a goal requires efforts in education, language design, and tools. The educational effort needs to address documentation and training issues. This implies that documentation of the specification language, including examples, should have a high

priority. It also implies that tutorial and teaching materials should have a high priority. This also presents an opportunity include in the project people interested in computer science education, with the goal of integrating more formal methods training into the standard computer science curriculum. One promising approach to doing this would be to promote undergraduate textbooks that use formal specifications.

The specification language design and tool-building efforts should have as their overall goal making it as easy as possible for programmers to read and write interesting specifications, in order to minimize educational costs. In the rest of this position paper, we will focus on these language design and tool building problems, since they are the ones we have the most experience with in JML.

4 Nature of the Specification Language

Assuming that the project will devote some effort to specification language design, we now consider what kinds of specification languages would be suitable, and whether there should be a single specification language.

A basic decision is whether the specification languages should be tailored to some specific programming language. That is, should they be interface specification languages (like Eiffel, Larch/C++, or JML), or should they be languages that are independent of any particular programming language (as are VDM, Z, or OCL)? Perhaps the specification language could even be that of some theorem prover, such as PVS or Isabelle?

We believe that the specification languages should be interface specification languages. The great advantage of an interface specification language is that it can specify details particular to some programming model, such as exceptions, visibility restrictions, encapsulation, and typing. Furthermore, an interface specification language can be translated into the input of various theorem provers and other tools, serving as a common front end for them. From the perspective of the verification tools, it may be possible to achieve some of the benefits of an interface specification language, while being somewhat language independent, by targeting an execution platform, such as the Java Virtual Machine or Microsoft's Common Language Runtime. However, even if the language targeted such an execution platform, to have impact on programmers the project would still require user-level specification languages that map to such platform-level languages.

One of the lessons of the JML project is that focusing on just a single specification language allows that language to serve as a central coordination mechanism for diverse groups of researchers. The common language:

- gives these groups a *lingua franca* in which to present semantic issues and tool-building issues,
- provides fresh insights as various technologies can be compared and contrasted through how they deal with common language features, and
- relieves groups of some of the tedium of fine-grained language design, which frees them to innovate in specialized areas.

Therefore, for the remainder of this paper, we will assume a single specification language targeting a specific programming language. Most of the arguments we raise will also apply to a lower-level specification language that targets an execution platform.

5 Problems in Specification Language Design

A major lesson of the JML project is that, even for sequential programs, the design of interface specification languages is still an interesting and quite difficult research problem.

The overall problem for specification language design is how the language can give its users sufficient value to justify the cost of specification. This is hard because, in our experience, the costs of writing a fairly complete functional specification of program behavior is usually about the same as that of writing the code to implement it. Therefore, it is necessary to either reduce these costs or to provide a wealth of tool support to compensate.

JML's approach has been to provide a wealth of tool support. This tool support works to both decrease the cost of writing specifications (for example, by improving the reporting of errors in specifications) and increase the usefulness of specifications (by using them to derive testing oracles and to generate documentation, among other things).

In addition to this tool support, JML is also designed to be easy for Java programmers to adopt and use. The main technical ideas here are to use an extended subset of Java expressions for assertions (drawing on the Eiffel experience), provide mathematical models as Java classes (drawing on the Larch and VDM experience), and prohibit side effects in expressions used in assertions [3].

Techniques for avoiding side effects in JML assertions are rather draconian at the moment and need refinement in order to be practical. The basic problem is that in a language like JML, one must be able to call methods in specifications. For example, to specify Java collections, one must be able to call the `equals` method on `Objects`. However, because of the specification inheritance used in JML to guarantee behavioral subtyping [6,7,8,9], saying that `Object`'s `equals` method is free of side effects means that no side effects are allowed in any overriding `equals` method [3]. We are working on allowing side effects that are not observable, which would allow caches in such methods, based on work by Barnett, *et al.* [10,11].

Another area of ongoing research is how to specify frame axioms (modifies clauses) and invariants in a modular fashion. To achieve modularity for frame axioms, there are several current approaches. The Boogie method [12,13] uses new language statements (`unpack` and `pack`), a dynamically unique object owner, and explicit specification of what objects are threatened by a method. While this makes soundness clear and is very flexible, it is (at present) fairly tedious to use in practice. Like the Boogie method, other techniques [14,15] seem to involve some control of aliasing. The same seems true of a modular treatment of invariants. Again, the Boogie method is more explicit, and also uses unique ownership, and other methods [16] also rely on control of aliasing. In any case some integration with alias control is necessary to prevent representation exposure. Thus these problems and current solutions seem to lead the research into the

realm of alias-controlling type systems [17,18]. However, such alias-controlling type systems and their integration with specification languages are still a very active area of research. The work on JML cited above [15,16] builds on one such research project, the Universe type system [19].

Another notable area of research is how to specify and verify callbacks. Higher-order features, and especially callbacks, are well known to cause difficulties both for practical programming [20]. In specification, the main problem is that the specifications become highly parameterized [21,22], and hence difficult to write and read. In JML we are pursuing the “grey-box approach” [23,24,25]: writing specifications in a refinement-calculus style as abstract programs, but interpreting the internal callbacks as observable.

All the above issues in specification language design apply to specification languages for sequential programs. Another host of issues enters when one tries to specify concurrent programs. The JML effort has mostly ignored these problems, because it has focused on a sequential subset of Java. However, one problem that is important for concurrent programs is also a problem for sequential programs: how to specify orderings of events. That is, the ability to specify permitted sequences of operations would be a useful adjunct to traditional Hoare-style specification languages. Some specification languages have used finite-state machine descriptions for such sequencing [26]. Temporal logic is a widely-known formalism for such sequence specifications [27], although other formalisms may be easier for ordinary programmers to write and read [28,29]. There has been some preliminary work in JML on this [30].

Beyond these issues in functional specification languages is the largely unexplored territory of specification languages that combine features for specification of both functional behavior (including data values) and concurrency control (including safety and liveness). Most specification languages and tools address only one of these areas. However, modern software requires a specification language that can help integrate techniques from these two disparate camps. A start towards this problem in JML is found in the work of Rodríguez *et al.* [31]. This work uses atomicity to reduce reasoning about concurrent programs to the sequential case.

6 Keeping up with Evolution

As we argued above, to have maximal impact on practice, a specification language and its tools must be designed for practicing programmers. The problem is that widely-used languages and their programming environments evolve rapidly, and it is difficult for researchers to keep up with industrial evolution of languages and development environments. This is the case for any “standardized” language, such as C (since the standards committees issue a new standard roughly every 5 years), and perhaps more so for languages that are not standardized, such as Java.

The history of the JML project provides several lessons related to this problem. The JML project started in 1998 as an outgrowth of work on the interface specification language Larch/C++.¹ Larch/C++ was an interface specification language tailored to the specification of C++ modules. C++ was originally chosen because it was a popular

¹ The Larch family of interface specification languages [32], was a refinement of Hoare-style specification languages exemplified by VDM.

object-oriented language, and we believed that we could have maximum impact by working with real problems in a broadly-used language.

However, the size and complexity of C++ presented great problems for our tool building efforts. In particular, the grammatical complexity of C++ presented enormous difficulty. In practice, another large difficulty was the unsafe nature of C++, which complicated specifications in many practical ways.

Thus when Java became available, we abandoned work on Larch/C++ and started work on an interface specification language tailored to Java—JML. Java was (at the time) a much simpler language than C++, which initially seemed to solve many of the tool problems.

Unfortunately, Java has since grown much larger than it was originally; furthermore, it has grown rather quickly. Indeed, the latest release of Java (version 1.5 also known as Java 5) introduces several non-trivial features, so it is challenging for an open-source project like JML to keep up. This dilemma is likely to be faced by the verifying compiler challenge also. In short: (a) in order to have impact, it is helpful to target a popular language, (b) but one way that a language stays popular is by evolving rapidly. This puts great pressure on fundamental tool support: parsing and basic compiler infrastructure have to be frequently updated to track the language's evolution.

Another issue for practical adoption is building integrated development environments (IDEs). Few modern programmers use the traditional command line compilers and old-style text editors; instead they demand integrated support for editing, compiling, and debugging (to say nothing of version control, support for refactoring, etc.). Thus, to have an impact—indeed, just to have users—tools must fit into some IDE. Like languages, IDEs are also evolving rapidly, as evidenced by Eclipse². Thus, while it is not a grand challenge to produce a state-of-the-art IDE, the project must keep up with their evolution.

The lesson we draw from this is that the project must have groups that are interested in building and maintaining basic compiler tool and IDE support. Researchers in formal methods are not interested in building basic compiler tools and IDEs; so the project must find people who are interested in these issues, in order to help the formal methods researchers stay focused. That is, formal methods researchers need help from some groups that are dedicated to providing up-to-date and extensible compiler support. These groups could track changes in programming language(s) and relieve some of the pressure in dealing with the evolution of popular languages. Similarly, formal methods researchers need help from groups that are dedicated to providing extensible IDEs.

The need to interact with researchers in extensible compilers and IDEs should be seen as an opportunity to involve more areas of computing in the grand challenge. One way to convince researchers interested in compilers to work with the grand challenge is to let them see additional opportunities for optimization and for exploring language extension techniques. Extensibility is needed not just to track evolution in the programming language, but also to allow experimentation in specification language design. Similarly, IDE researchers might be attracted to the grand challenge if they see opportunities to experiment with extensibility, or to improve the programmer's experience through improved error reporting, specification-enabled visualization, or editing mechanisms.

² Details on Eclipse's rapid evolution are available from <http://www.eclipse.org>.

It is a fortunate fact that, at least in the United States, compiler research and formal methods research are housed in the same division of the National Science Foundation. It will be easier to get support for a grand challenge if we can find ways to involve compilers and other subareas of computing (such as education and human-computer interfaces).

7 Conclusions

To summarize, the grand challenge needs interesting specifications as tests of its verification technology. To have maximum impact, the grand challenge should have as a goal making it easy and cost-effective to write such interesting specifications. This necessarily involves some effort in specification language design. We listed some of the research challenges in specification language design from our experience with the JML project; the grand challenge should not assume that specification languages are completely understood and adequate.

The lessons from the JML project also indicate that the verifying compiler grand challenge should involve researchers in extensible compiler technologies and IDEs. Work in all these areas is necessary for the project to have a practical impact worthy of being a “grand challenge.”

Acknowledgments

Thanks to the program committee for comments that helped clarify our arguments.

The work of Leavens and Clifton was supported in part by the US National Science Foundation through grants CCF-0428078 and CCF-0429567.

References

1. Hoare, T.: The verifying compiler: A grand challenge for computing research. *Journal of the ACM* **50**(1) (January 2003) 63–69
2. Burdy, L., Cheon, Y., Cok, D.R., Ernst, M.D., Kiniry, J.R., Leavens, G.T., Leino, K.R.M., Poll, E.: An overview of JML tools and applications. *International Journal on Software Tools for Technology Transfer* **7**(3) (June 2005) 212–232
3. Leavens, G.T., Cheon, Y., Clifton, C., Ruby, C., Cok, D.R.: How the design of JML accommodates both runtime assertion checking and formal verification. *Science of Computer Programming* **55**(1-3) (March 2005) 185–208
4. Flanagan, C., Leino, K.R.M., Lillibridge, M., Nelson, G., Saxe, J.B., Stata, R.: Extended static checking for Java. In: *Proceedings of the ACM SIGPLAN 2002 Conference on Programming Language Design and Implementation (PLDI’02)*. Volume 37(5) of *SIGPLAN.*, New York, NY, ACM (June 2002) 234–245
5. Meyer, B.: *Object-oriented Software Construction*. second edn. Prentice Hall, New York, NY (1997)
6. Dhara, K.K., Leavens, G.T.: Forcing behavioral subtyping through specification inheritance. In: *Proceedings of the 18th International Conference on Software Engineering, Berlin, Germany, IEEE Computer Society Press (March 1996)* 258–267 A corrected version is ISU CS TR #95-20c, <http://tinyurl.com/s2krg>.

7. Leavens, G.T.: JML's rich, inherited specifications for behavioral subtypes. In Liu, Z., Jifeng, H., eds.: *Formal Methods and Software Engineering: 8th International Conference on Formal Engineering Methods (ICFEM)*. Volume 4260 of *Lecture Notes in Computer Science.*, New York, NY, Springer-Verlag (November 2006) 2–34
8. Leavens, G.T., Naumann, D.A.: Behavioral subtyping, specification inheritance, and modular reasoning. Technical Report 06-20b, Department of Computer Science, Iowa State University, Ames, Iowa, 50011 (September 2006)
9. Leavens, G.T., Poll, E., Clifton, C., Cheon, Y., Ruby, C., Cok, D.R., Müller, P., Kiniry, J., Chalin, P.: JML reference manual. Department of Computer Science, Iowa State University. Available from <http://www.jmlspecs.org> (February 2007)
10. Barnett, M., Naumann, D.A., Schulte, W., Sun, Q.: Allowing state changes in specifications. In Müller, G., ed.: *International Conference on Emerging Trends in Information and Communication Security (ETRICS)*. Volume 3995 of *Lecture Notes in Computer Science.*, Springer (2006) 321–336
11. Naumann, D.A.: Observational purity and encapsulation. In: *Fundamental Aspects of Software Engineering (FASE)*. (2005) 190–204
12. Barnett, M., DeLine, R., Fähndrich, M., Leino, K.R.M., Schulte, W.: Verification of object-oriented programs with invariants. *Journal of Object Technology* **3**(6) (2004) 27–56
13. Leino, K.R.M., Müller, P.: Object invariants in dynamic contexts. In Odersky, M., ed.: *European Conference on Object-Oriented Programming (ECOOP)*. Volume 3086 of *Lecture Notes in Computer Science.*, Springer-Verlag (June 2004) 491–516
14. Leino, K.R.M., Nelson, G.: Data abstraction and information hiding. *ACM Transactions on Programming Languages and Systems* **24**(5) (September 2002) 491–553
15. Müller, P., Poetzsch-Heffter, A., Leavens, G.T.: Modular specification of frame properties in JML. *Concurrency and Computation: Practice and Experience* **15**(2) (February 2003) 117–154
16. Müller, P., Poetzsch-Heffter, A., Leavens, G.T.: Modular invariants for layered object structures. *Science of Computer Programming* **62**(3) (October 2006) 253–286
17. Boyland, J., Noble, J., Retert, W.: Capabilities for sharing. In Knudsen, J.L., ed.: *ECOOP 2001 — Object-Oriented Programming: 15th European Conference, Budapest, Hungary*. Volume 2072 of *Lecture Notes in Computer Science.*, Berlin, Springer-Verlag (June 2001) 1–27
18. Noble, J., Vitek, J., Potter, J.: Flexible alias protection. In Jul, E., ed.: *ECOOP '98 – Object-Oriented Programming, 12th European Conference, Brussels, Belgium*. Volume 1445 of *Lecture Notes in Computer Science.*, Springer-Verlag (July 1998) 158–185
19. Dietl, W., Müller, P.: Universes: Lightweight ownership for JML. *Journal of Object Technology* **4**(8) (October 2005) 5–32
20. Szyperki, C., Gruntz, D., Murer, S.: *Component Software: Beyond Object-Oriented Programming*. Second edition edn. ACM Press and Addison-Wesley, New York, NY (2002)
21. Ernst, G.W., Navlakha, J.K., Ogden, W.F.: Verification of programs with procedure-type parameters. *Acta Informatica* **18**(2) (November 1982) 149–169
22. Goguen, J.A.: Parameterized programming. *IEEE Transactions on Software Engineering* **SE-10**(5) (September 1984) 528–543
23. Büchi, M., Weck, W.: A plea for grey-box components. Technical Report 122, Turku Center for Computer Science, Presented at the Workshop on Foundations of Component-Based Systems, Zürich, September 1997 (1997) <http://tinyurl.com/2833tr>.
24. Büchi, M., Weck, W.: The greybox approach: When blackbox specifications hide too much. Technical Report 297, Turku Center for Computer Science (August 1999) <http://tinyurl.com/ywmuzy>.
25. Büchi, M.: Safe language mechanisms for modularization and concurrency. Technical Report TUCS Dissertations No. 28, Turku Center for Computer Science (May 2000)

26. Ball, T., Rajamani, S.K.: The SLAM project: Debugging system software via static analysis. In: Conference Record of POPL'02: The 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, Oregon (January 16–18, 2002) 1–3
27. Manna, Z., Pnueli, A.: The Temporal Logic of Reactive and Concurrent Systems. Springer-Verlag, New York, NY (1992)
28. Corbett, J.C., Dwyer, M.B., Hatcliff, J., Laubach, S., Pasareanu, C.S., Robby, Zheng, H.: Bandera: Extracting finite-state models from Java source code. In: Proceedings of the 22nd International Conference on Software Engineering, New York, NY, ACM Press (June 2000) 439–448
29. Harel, D.: Statecharts: A visual formalism for complex systems. *Science of Computer Programming* **8**(3) (June 1987) 231–274
30. Cheon, Y., Perumendla, A.: Specifying and checking method call sequences in JML. In Arabnia, H.R., Reza, H., eds.: Proceedings of the 2005 International Conference on Software Engineering Research and Practice (SERP '05), Volume II, Las Vegas, Nevada, June 27-29, 2005, CSREA Press (2005) 511–516
31. Rodríguez, E., Dwyer, M.B., Flanagan, C., Hatcliff, J., Leavens, G.T., Robby: Extending JML for modular specification and verification of multi-threaded programs. In Black, A.P., ed.: ECOOP 2005 — Object-Oriented Programming 19th European Conference, Glasgow, UK. Volume 3586 of Lecture Notes in Computer Science. Springer-Verlag, Berlin (July 2005) 551–576
32. Guttag, J.V., Horning, J.J., Garland, S., Jones, K., Modet, A., Wing, J.: Larch: Languages and Tools for Formal Specification. Springer-Verlag, New York, NY (1993)