

# The Spec# Programming System: Challenges and Directions

Mike Barnett, Robert DeLine, Manuel Fähndrich, Bart Jacobs,  
K. Rustan M. Leino, Wolfram Schulte, and Herman Venter

Microsoft Research, Redmond, WA, USA  
{mbarnett,rdeline,maf,leino,schulte,hermanv}@microsoft.com  
bart.jacobs@cs.kuleuven.be

## 1 Introduction

The Spec# programming system [4] is a new attempt to increase the quality of general purpose, industrial software. Using old wisdom, we propose the use of specifications to make programmer assumptions explicit. Using modern technology, we propose the use of tools to enforce the specifications. To increase its chances of having impact, we want to design the system so that it can be widely adopted.

For a programming system to be adopted widely, we think that it must:

- build on a widely used object-oriented programming language; in our case C#;
- build on existing infrastructure and allow interoperability with existing code, here the .NET runtime.
- fully integrate into an development environment; in our case the Microsoft Visual Studio environment.
- build on a teachable and sound methodology; in our case a revised design-by-contract methodology;
- include tools that enforce the methodology; in our case this includes type checking, easily usable dynamic checking, as well as high-assurance automatic static verification;
- support a smooth adoption path whereby programmers can gradually start taking advantage of the benefits of specification;
- be moderate; we added only a few constructs to C#, and soundness is guaranteed only as long as the source comes from constructs that are under our control.

In this extended abstract, we give an overview of the Spec# programming system [as of September 2005], the rationale of its design, and a sketch of some open problems. Spec# is currently under development at Microsoft Research, Redmond.

## 2 The Language

The Spec# language is a superset of C#, an object-oriented language targeted for the .NET Platform. C# features single inheritance whose classes can implement multiple interfaces, object references, dynamically dispatched methods, and exceptions, to mention the features most relevant to this paper.

Spec# adds to C# type support for distinguishing non-null object references from possibly-null object references [14], method specifications like pre- and postconditions, a discipline for managing exceptions [22], and support for constraining the data fields of objects [2]. While conceptually simple, all have some complicated consequences. Next, we give an overview of each feature.

## 2.1 Non-null Types

Many errors in modern programs manifest themselves as null-dereference errors. We have opted to add type support for nullity discrimination to Spec#, because we think types offer the easiest way for programmers to take advantage of nullity distinctions. Here is a list of challenges with some solutions.

**Initialization of fields and arrays.** For type safety, any time a variable of a non-null type is read, the value obtained must be non-null. Without further restrictions, this is not guaranteed for instance fields, static fields, and array elements. Spec# offers a simple syntactic solution for instance fields, which ensures that they are initialized before the object being constructed can be accessed. Handling non-null static fields requires ordering restrictions on the initialization of classes. Spec# supports non-null static fields, but the initialization restrictions are checked only at run time. Supporting non-null element types of arrays is tricky, because there is no language construct that marks the end of the initialization of an array's elements.

**Language interoperability.** The .NET platform supports multiple languages, not all of which may be required to support non-null types. This raises many problems, like calling from such a language methods whose in-parameters have non-null types. One could disallow such calls, but we want to encourage libraries to be updated with explicit non-null information without forcing changes in the clients. A better solution is to add non-null types to the virtual machine and to adapt the just-in-time compiler, but this would greatly impact all virtual-machine implementations. Arrays with non-null element types also require changes in the virtual machine, because of array covariance. For this reason, we do not support them in Spec#.

**Stability of non-nullity.** One must support some form of down-cast of an expression from a maybe-null type to a non-null type. The Spec# compiler infers the non-nullity of local variables by performing a data-flow analysis that takes into account type casts and tests for null. However, this is problematic for fields, whose values may be changed between a test and a use by calls or by other threads about which the compiler has no information.

Non-null types with various degrees of type soundness have also been used in program checking systems like LCLint [13] and are being incorporated into the object-oriented language Eiffel [24]; see Fähndrich and Leino's paper [14] for more information about previous work.

## 2.2 Method Contracts

To allow programmers to capture more complicated properties, which may involve more than one variable, we advocate using general specifications instead of further

enrichments to the type system. Like for example Eiffel [23], Spec# supports pre- and postconditions of methods; these use ordinary side-effect free boolean expressions.

**Inheritance.** Pre- and postconditions are inherited in method overrides. Spec# allows overrides to declare additional postconditions. However, Spec# does not allow an override to weaken the precondition, despite the fact that this would be sound. Here's our justification: There is a mindset in .NET (and in Java) that the misuse of language primitives (like indexing an array outside its bounds) and methods always be detected and reported immediately when the violation occurs at run time. It would seem to go against this mindset to allow an override to eliminate this detection. Actually, though a fine syntactic convenience, there is never a need for weakening preconditions, because a subclass can always define a new method with a weaker precondition and let the old method call the new one.

Because of interfaces, there is a limited form of multiple inheritance. In these cases, a method implementation can inherit its specification from several sources. However, to avoid weakening preconditions, Spec# allows this multiple inheritance only in certain situations where all the inherited preconditions are the same. Actually, this is not a semantic restriction in Spec#, because a class can provide different implementations for multiply inherited methods.

**Frame conditions.** To reason statically about a call, one needs to know what variables the callee may modify. Specifying this *frame condition* is difficult, because of information hiding: the frame condition must say that the caller's variables are unchanged and must allow the implementation's variables to be changed. But these variables are in general not visible to both the caller and the implementation. Spec# uses syntactic and semantic abstractions to address this problem, but we do not yet have enough evidence to evaluate whether or not this particular solution is the right one.

**Evaluation of contracts.** We take the view that contracts should contribute to the specification of programs, but not to their effects. Therefore, Spec# insists that expressions used in contracts be side-effect free. We do allow procedural abstraction in contracts, as long as any procedure (method, rather) called also is side-effect free, that is, *pure*. A syntactic notion of purity is overly restrictive, so we are exploring notions of *behavioral purity*, but don't have all the answers yet [6,26].

The fact that contracts are given by expressions of the language raises the problem that the evaluation of those expressions can fail. Spec# chooses to hold the specification writer accountable for such errors. That is, contracts are enforced to be totally defined.

## 2.3 Class Contracts

Invariants are key to describing the correctness of programs, and Spec# allows classes to declare invariants that describe the internal consistency of the state of their objects. However, the problem of maintaining invariants that span several objects has been under-studied. Two central problems are delineating when an object invariant holds (which is made difficult because of re-entrancy) and controlling changes to sub-objects (which is made difficult because references to these sub-objects may be leaked).

Spec# uses a sound methodology that addresses both of these central problems (and the problem of writing frame conditions) [2,18,5,19]. Spec# provides a block statement that delineates where object invariants may temporarily be violated, and it uses ownership domains to confine references.

## 2.4 Concurrency

Multi-threading introduces the possibility of race conditions and deadlocks. This makes data consistency even more difficult to achieve in a multi-threaded program. We feel that a programming methodology should permit extensions that cover concurrent programs, too. We have formulated an extension of our object-invariant methodology for multi-threading [15]. It ensures mutual exclusion on ownership domains and maintains data consistency. But our extension still has several shortcomings, including: it does not check for deadlocks [but see our newer work [16]] and its locking is sometimes too coarse-grained.

## 2.5 Data Abstraction

More elaborate specifications of programs require a way to present a view of objects that abstracts away from implementation details. We find that this form of data abstraction already exists in many .NET programs in the form of *properties* (which are essentially parameter-less methods). If a property is a pure function of the state of an object and its immediate ownership domain, then we call it a *model field* [9,17,21,25]. We are hopeful that one can formulate the definition and use of model fields in a way that is amenable to static program verification [see our newer work [20]].

# 3 System Architecture

Architecturally, the Spec# programming system consists of the compiler, a runtime library, and the static program verifier. Spec# has been integrated into the Microsoft Visual Studio environment. For example, violations of the non-null type system are indicated by “red squiggles”, specifications of methods (including any specifications on library methods) are available in tool tips. The static program verifier runs continuously while editing the program and interactively produces red squiggles for semantic errors, but it can also be run as a standalone tool.

## 3.1 Levels of Checking

Spec# provides three levels of checking.

The first level of checking is provided by the type checker, which runs as part of the compiler and which must accept the program before any code is emitted. The type checker is stronger than some other traditional type checkers in that it is sometimes sensitive to data flow.

The second level of checking is provided by compiler-emitted run-time checks. These checks are always emitted, but we do provide some compiler options that disable them,

because development organizations sometimes feel they have reached a point in the development cycle where they are willing to risk not having the checks and would rather gain performance. The run-time checks enforce many contracts, but do not enforce the entire Spec# programming methodology. For example, frame conditions (modifies clauses) are not checked at run time and, by default, ownership domains are not enforced at run time.

The third level of checking, which is optional, is provided by the Spec# static program verifier. It enforces all contracts and the entire Spec# programming methodology, except assume statements, which are provided for the specific purpose of introducing a run-time check for programmer assumptions that would take great effort to prove statically.

### 3.2 Contract Persistence

The Spec# compiler preserves the specifications as metadata in the same binary assembly as the compiled code. This enables reuse of specifications across tools.

To enable clients of a legacy library to be verified, such a library needs to be retrofitted with specifications. If the library cannot be converted into Spec#, we support the compilation and use of *out-of-band contracts* in shadow assemblies. These give the illusion that the specifications were declared in the library itself.

### 3.3 Static Verification

From MSIL (the .NET bytecode), Spec#'s static program verifier (whose codename is Boogie [1]) constructs a program in its own intermediate language, BoogiePL [12]. BoogiePL is a simple imperative language with procedures. BoogiePL also supports the introduction of uninterpreted function symbols and axioms, which makes it suitable as a stepping stone in program verification. In fact, all of the Spec# to be verified is translated into BoogiePL, including the axiomatization of the Spec# type system, *etc.* From the BoogiePL program, Boogie infers loop invariants using abstract interpretation [10,11,7,8] and generates verification conditions (logical formulas that are valid iff the program is correct) that it passes to an automatic theorem prover [3]. Counterexamples reported by the theorem prover are translated back into error messages about the source code, which are reported to the user.

Technology for abstract interpretation and theorem proving exist. We feel that the work in this area has reached the kind of maturity where what is needed for Spec# consists mainly in adapting and tuning existing technologies for the verification task at hand. One difference from previous work is perhaps the emphasis on the heap in Spec# programs. There is also room for improving the combination of various abstract domains as well as exploring the combination of abstract domains and decision procedures.

In the verification community, various standard or canonical formats have been useful for interoperability, substitutability, and evaluation of tools (for example, the primitive DIMAC format for SAT formulas). Our position is that BoogiePL is a good candidate for playing that same role in the space of verifying programs. Boogie can be invoked not just on MSIL assemblies, but also on BoogiePL programs directly.

This means that other researchers can reuse the abstract interpretation and verification-condition generation in Boogie.

## 4 Conclusion

We are excited about the Verification Grand Challenge working conference. We see the largest remaining challenge to be the formulation of programming methodology that allows modern programs to be specified and verified, as well as the serious engineering effort to build such a programming system.

The Spec# system, including the Boogie program verifier, can be downloaded from <http://research.microsoft.com/SpecSharp>.

## References

1. Barnett, M., Chang, B.-Y.E., DeLine, R., Jacobs, B., Leino, K.R.M.: Boogie: A Modular Reusable Verifier for Object-Oriented Programs. In: de Boer, F.S., Bonsangue, M.M., Graf, S., de Roever, W.-P. (eds.) FMCO 2005. LNCS, vol. 4111, pp. 364–387. Springer, Heidelberg (2006)
2. 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), [www.jot.fm](http://www.jot.fm)
3. Barnett, M., Leino, K.R.M.: Weakest-precondition of unstructured programs. In: Proceedings of the 2005 ACM SIGPLAN-SIGSOFT Workshop on Program Analysis For Software Tools and Engineering, PASTE 2005, September 2005, ACM, New York (2005)
4. Mike Barnett, K., Rustan, M.: Leino, and Wolfram Schulte. In: Barthe, G., Burdy, L., Huisman, M., Lanet, J.-L., Muntean, T. (eds.) CASSIS 2004. LNCS, vol. 3362, pp. 49–69. Springer, Heidelberg (2005)
5. Barnett, M., Naumann, D.A.: Friends Need a Bit More: Maintaining Invariants Over Shared State. In: Kozen, D. (ed.) MPC 2004. LNCS, vol. 3125, pp. 54–84. Springer, Heidelberg (2004)
6. Barnett, M., Naumann, D.A., Schulte, W., Sun, Q.: 99.44% pure: Useful abstractions in specifications. In: Proceedings, 6th workshop on Formal Techniques for Java-like Programs (June 2004)
7. Chang, B.-Y.E., Leino, K.R.M.: Abstract Interpretation with Alien Expressions and Heap Structures. In: Cousot, R. (ed.) VMCAI 2005. LNCS, vol. 3385, pp. 147–163. Springer, Heidelberg (2005)
8. Chang, B.-Y.E., Leino, K.R.M.: Inferring object invariants. In: Proceedings of First International Workshop on Abstract Interpretation of Object-Oriented Languages (AIOOL 2005) (2005)
9. Cheon, Y., Leavens, G.T., Sitaraman, M., Edwards, S.: Model variables: cleanly supporting abstraction in design by contract. *Software—Practice and Experience* 35(6), 583–599 (2005)
10. Cousot, P., Cousot, R.: Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: Conference Record of the Fourth Annual ACM Symposium on Principles of Programming Languages, January 1977, pp. 238–252. ACM, New York (1977)
11. Cousot, P., Halbwachs, N.: Automatic discovery of linear restraints among variables of a program. In: Conference Record of the Fifth Annual ACM Symposium on Principles of Programming Languages, January 1978, pp. 84–96 (1978)

12. DeLine, R., Rustan, K., Leino, M.: BoogiePL: A typed procedural language for checking object-oriented programs. Technical report, Microsoft Research (2005)
13. Evans, D., Gutttag, J.V., Horning, J.J., Tan, Y.M.: LCLint: A tool for using specifications to check code. In: Wile, D.S. (ed.) SIGSOFT 1994, Proceedings of the Second ACM SIGSOFT Symposium on Foundations of Software Engineering, December 1994. ACM SIGSOFT Software Engineering Notes, vol. 19(5), pp. 87–96 (1994)
14. Fähndrich, M., Leino, K.R.M.: Declaring and checking non-null types in an object-oriented language. In: Crocker, R., Steele Jr., G.L. (eds.) Proceedings of the 2003 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications, OOPSLA 2003, October 2003. SIGPLAN Notices, vol. 38(11), pp. 302–312. ACM, New York (2003)
15. Jacobs, B., Leino, K.R.M., Piessens, F., Schulte, W.: Safe concurrency for aggregate objects with invariants. In: Aichernig, B.K., Beckert, B. (eds.) 3rd International Conference on Software Engineering and Formal Methods, September 2005, pp. 137–146. IEEE, Los Alamitos (2005)
16. Jacobs, B., Smans, J., Piessens, F., Schulte, W.: A statically verifiable programming model for concurrent object-oriented programs. In: Liu, Z., He, J. (eds.) ICFEM 2006. LNCS, vol. 4260, pp. 420–439. Springer, Heidelberg (2006)
17. Leavens, G.T., Baker, A.L., Ruby, C.: Preliminary design of JML: A behavioral interface specification language for Java. Technical Report 98-06-rev28, Iowa State University, Department of Computer Science, See (2003), <http://www.jmlspecs.org>
18. Rustan, K., Leino, M., Müller, P.: Object invariants in dynamic contexts. In: Martin Odersky. In: Odersky, M. (ed.) ECOOP 2004. LNCS, vol. 3086, pp. 491–516. Springer, Heidelberg (2004)
19. Rustan, K., Leino, M., Müller, P.: Modular verification of static class invariants. In: Fitzgerald, J.S., Hayes, I.J., Tarlecki, A. (eds.) FM 2005. LNCS, vol. 3582, pp. 26–42. Springer, Heidelberg (2005)
20. Rustan, K., Leino, M., Müller, P.: A verification methodology for model fields. In: Sestoft, P. (ed.) ESOP 2006 and ETAPS 2006. LNCS, vol. 3924, pp. 115–130. Springer, Heidelberg (2006)
21. Rustan, K., Leino, M., Nelson, G.: Data abstraction and information hiding. ACM Transactions on Programming Languages and Systems 24(5), 491–553 (2002)
22. Rustan, K., Leino, M., Schulte, W.: Exception safety for C#. In: Cuellar, J.R., Liu, Z. (eds.) SEFM 2004—Second International Conference on Software Engineering and Formal Methods, September 2004, pp. 218–227. IEEE, Los Alamitos (2004)
23. Meyer, B.: Object-oriented Software Construction. Series in Computer Science. Prentice-Hall International, New York (1988)
24. Meyer, B.: Attached Types and Their Application to Three Open Problems of Object-Oriented Programming. In: Black, A.P. (ed.) ECOOP 2005. LNCS, vol. 3586, pp. 1–32. Springer, Heidelberg (2005)
25. Müller, P.: Modular Specification and Verification of Object-Oriented Programs. LNCS, vol. 2262. Springer, Heidelberg (2002)
26. Naumann, D.A.: Observational purity and encapsulation. In: Cerioli, M. (ed.) FASE 2005. LNCS, vol. 3442, pp. 190–204. Springer, Heidelberg (2005)

## A Discussion on Rustan Leino’s Presentation

### Roderick Chapman

I noticed, you spent a lot of time on declaring objects to be of type `int`, and then saying: “Oh no they’re not, they’re positive.” You say it’s an `int`, then you say “oh no it’s not,



it's greater than or equal to nought, or it's greater than nought, it's positive, blah, blah, blah..." and you seem to have to do this quite a lot.

**Rustan Leino:** Yes.

### **Roderick Chapman**

Am I completely missing something about C# or does the language not have a facility to declare integer subrange types?

**Rustan Leino:** No integer subrange types.

### **Roderick Chapman**

Is this not an abominable decision? It shatters me to find that that is not in the language!

### **Rustan Leino**

It would be great to have. Of course, the type system would not do everything for you. You would get runtime checks or the static verifier would pick up where the type checker leaves off, like for so many other things. But sure, that would be great. Actually, there are a number of such things that one can imagine wanting. We've added a few things to C# to make Spec#, but not everything.

### **Bertrand Meyer**

On that last point, I am with you, Rustan. Next, you'll be asked about even integers: Are you going to have a special type for that? But actually, I had a different question. In my experience, it is extremely common to have a postcondition that refers to a secret attribute, and there is absolutely nothing wrong with this conception. Preconditions, of course, are a different business. And the way you addressed the supposed error—which, I think, is not an error-scared me, because then you made the attribute public, and I suspect that in the C++/Java-tradition making an attribute public means making it public for write as well as read. It is not the path that I would like to follow.

### **Rustan Leino**

Yes, thank you for that comment. I made the chunk size read-only here, if that perhaps would make you feel a little bit better. But yes, in general, we would not have to report that error message for postconditions, you are quite right. You have to be a little bit careful that it is at least protected, because if you inherit the method, then the subclass also has to live up to that postcondition. But indeed, preconditions—that's more a requirement-and invariants also help you there, because invariants can talk about the private state. And of course, you need some ways to abstract over the state in other ways, which we will support, but actually, for the moment, we don't, but we will.



### **Jianjun Zhao**

Would you tell us some lessons learned from the design of Spec# for programming language design?

### **Rustan Leino**

To the programming language community, for example. I think we learned many lessons and it would be nice, if the programming language community would pick up some. Let me just say something simple, since we are short on time, maybe we can take it more in discussion later. You saw me move the base class constructor call I move that from the beginning of the method to the end. Well, that is required for soundness, if you have something like non-null types. And the interesting thing is that if you go through your C# programs to convert them into Spec# and you add some non-null features and things, it turns out that almost always, you call the base class constructor last. So, if we should have any default at all, it should be to call it last. But neither C# nor Java supports that today. So, that would be one lesson, but I think, there are more lessons as well.