

Preface

As a programmer working for Logica UK in London in the mid-1980's, I became a passionate advocate of formal methods. Extrapolating from small successes with VDM and JSP, I was sure that widespread use of formal methods would bring an end to the software crisis.

One approach especially intrigued me. John Guttag and Jim Horning had developed a language, called Larch, which was amenable to a mechanical analysis. In a paper they'd written a few years earlier [23], and which is still not as widely known as it deserves to be, they showed how questions about a design might be answered automatically. In other words, we would have real software “blueprints”—a way to analyze the essence of the design before committing to code. I went to pursue my PhD with John at MIT, and have been a researcher ever since.

As a researcher though, I soon discovered that formal methods were not the silver bullet I'd hoped they would be. Formal models were hard to construct, and specifying every detail of a system was too hard. Theorem proving, the kind of analysis that Larch relied on, could not be fully automated. Even now, after 20 more years of research, it still requires the careful guidance of a mathematical guru. In my doctoral work, therefore, I took a more conservative route, and worked on automatic detection of bugs in code. But I kept an interest in the more ambitious world of formal methods and design analysis, and hoped one day to return to it.

In 1992, I visited Carnegie Mellon University. By then, I'd become enamored, like many in the formal methods community, with the Z language. The inventors of Z had dispensed with many of the complexities of earlier languages, and based their language on the simplest notions of set theory. And yet Z was even less analyzable than Larch; the only tool in widespread use was a pretty printer and type checker.

On that visit, Ken McMillan showed me his SMV model checker: a tool that could check a state machine of a billion states in seconds, without any aid from the user whatsoever. I was awestruck.

With the invention of model checking, the reputation of formal methods changed almost overnight. The word “verification” became fashionable again, and the adoption of model-checking tools by chip manufac-

turers showed that engineers really could write formal models, and, if the benefit was great enough, would do it of their own accord.

But the languages of model checkers were not suitable for software. They were designed for handling the complexity that arises when a collection of simple state machines interacts concurrently. In software design, complexity arises even in a single machine, from the complex structure of its state. Model checkers can't handle this structure—not even the indirection that is the essence of all software design.

So I began to wonder: could the power of model checking be brought to a language like Z? Here were two cultures, an ocean apart: the gritty automation of SMV, reflecting the steel mills and smokestacks of Pittsburgh, the town of its invention, and the elegance and simplicity of Z, reflecting the beautiful quads of Oxford.

This book is the result of a 10-year effort to bridge this gap, to develop a language that captures the essence of software abstractions simply and succinctly, with an analysis that is fully automatic, and can expose the subtlest of flaws.

The language, Alloy, is deeply rooted in Z. Like Z, it describes all structures (in space and time) with a minimal toolkit of mathematical notions, but its toolkit is even smaller and simpler than Z's. Alloy was also strongly influenced by object modeling notations (such as those of OMT and Syntropy). Like them, it makes it easy to classify objects, and associate properties with objects according to the classification. Alloy supports “navigation expressions,” which are now a mainstay of object modeling, with a syntax that is particularly simple and uniform.

The analysis, embodied in the Alloy Analyzer, actually bears little resemblance to model checking, its original inspiration. Instead, it relies on recent advances in SAT (boolean satisfiability) technology. The Alloy Analyzer translates constraints to be solved from Alloy into boolean constraints, which are fed to an off-the-shelf SAT solver. As solvers get faster, so Alloy's analysis gets faster and scales to larger problems. Using the best solvers of today, the analyzer can examine spaces that are several hundred bits wide (that is, of 10^{60} cases or more). Hardware advances must also get some of the credit. Even had this technology been available 10 years ago, an analysis that takes only seconds on today's machines would have taken an hour back then. (Incidentally, Alloy was by no means the first application of SAT to this kind of problem. SAT had been used for analyzing railway control systems [68], for checking hardware [69], and for planning [45, 17]. Since its adoption in Alloy [33], it has been incorporated into model checkers too [5].)

The experience of exploring a software model with an automatic analyzer is at once thrilling and humiliating. Most modelers have had the benefit of review by colleagues; it's a sure way to find flaws and catch omissions. Few modelers, however, have had the experience of subjecting their models to continual, automatic review. Building a model incrementally with an analyzer, simulating and checking as you go along, is a very different experience from using pencil and paper alone. The first reaction tends to be amazement: modeling is much more fun when you get instant, visual feedback. When you simulate a partial model, you see examples immediately that suggest new constraints to be added.

Then the sense of humiliation sets in, as you discover that there's almost nothing you can do right. What you write down doesn't mean exactly what you think it means. And when it does, it doesn't have the consequences you expected. Automatic analysis tools are far more ruthless than human reviewers. I now cringe at the thought of all the models I wrote (and even published) that were never analyzed, as I know how error-ridden they must be. Slowly but surely the tool teaches you to make fewer and fewer errors. Your sense of confidence in your modeling ability (and in your models!) grows.

You can use analysis to make models not only more correct but also more succinct and more elegant. When you want to rework a constraint in the model, you can ask the analyzer to check that the new and old constraint have the same meaning. This is like using unit tests to check refactoring in code, except that the analyzer typically checks billions of cases, and there are no test suites to write.

I sometimes call my approach "lightweight formal methods" [39], because it tries to obtain the benefits of traditional formal methods at lower cost, and without requiring a big initial investment. Models are developed incrementally, driven by the modeler's perception of which aspects of the software matter most, and of where the greatest risks lie, and automated tools are exploited to find flaws as early as possible.

But at the same time as I have argued against some of the assumptions of traditional formal methods, my experience in the last decade—teaching software engineering to students at Carnegie Mellon and MIT, building tools with students, and consulting on industrial developments—has convinced me of the validity of their central premise. As Tony Hoare famously put it in his Turing Award lecture [31]:

There are two ways of constructing a software design: One way is to make it so simple there are obviously no deficiencies and

the other way is to make it so complicated that there are no obvious deficiencies.

A commitment to simplicity of design means addressing the essence of design—the abstractions on which software is built—explicitly and up front. Abstractions are articulated, explained, reviewed and examined deeply, in isolation from the details of the implementation. This doesn't imply a waterfall process, in which all design and specification precedes all coding. But developers who have experienced the benefits of this separation of concerns are reluctant to rush to code, because they know that an hour spent on designing abstractions can save days of refactoring.

In this respect, the Alloy language and its analysis are a Trojan horse: an attempt to capture the attention of software developers, who are mired in the tar pit of implementation technologies, and to bring them back to thinking deeply about underlying concepts.

That is why I have chosen the title *Software Abstractions* for this book. The lure of coding, and pressure to deliver elaborate features on short schedules, often draw programmers away from designing abstractions to coping with the intricacies of transient technologies, and to inventing clever tricks to overcome their limitations. If we focused instead on the underlying concepts, and struggled not for small performance gains or ever more complex features, but for simplicity and clarity, our software would be more powerful, more dependable, and more enjoyable to use. Like the best artifacts of civil and mechanical engineering, the best software systems would be a marriage of utility and beauty. And as software designers, we'd have more fun: we'd spend less time working around basic structural flaws in our software, and our ideas would have more lasting impact.

Acknowledgments

I am deeply grateful to the many friends and colleagues who have helped in the writing of this book:

To Ilya Shlyakhter, who invented the modeling idiom that expresses dynamics by adding a column of state atoms to each relation (leading to the design of the signature construct, and making possible Alloy’s precarious balance of expressiveness and tractability), and who designed and built the key algorithms of the Alloy Analyzer.

To Manu Sridharan, who contributed extensively to the language, designed and implemented large parts of the analyzer, was an enthusiast for Alloy before we had credible examples, and has continued to help out despite having left MIT long ago.

To the many undergraduate and master’s students who contributed to the tool implementation: Arturo Arizpe, Emily Chang, Joseph Cohen, Sam Daitch, Greg Dennis, David Kelman, Daniel Kokotov, Edmond Lau, Likuo Lin, Jesse Pavel, Uriel Schafer, Ian Schechter, Ning Song, Emina Torlak, Vincent Yeung, and Andrew Yip; and to those who were guinea pigs in evaluating Alloy in early case studies: Ryan Jazayeri, Sarfraz Khurshid, Edmond Lau, Robert Lee, SeungYong Albert Lee, Kartik Mani, Tina Nolte, Suresh Toby Segaran, Tucker Sylvestro, Mana Taghdiri, Allison Waingold, Hoe Teck Wee, and Jon Whitney; and to MIT’s UROP office for coordinating the undergraduate research program.

To the current members of my research group—Felix Chang, Greg Dennis, Jonathan Edwards, Lucy Mendel, Derek Rayside, Robert Seater, Mana Taghdiri, Emina Torlak, and Vincent Yeung—not only for their intellectual company, but for their many contributions to the Alloy project big and small; especially to Derek who, on his own initiative, took on the task of resolving release problems and platform dependences; to Emina, now Alloy’s lead developer, and Vincent, for their continuing work on the Alloy Analyzer; to Jonathan, who led the design of Alloy’s new type system; to Robert, for his help teaching Alloy; and to Greg, for his work on the Alloy library modules and for answering queries from users. To Viktor Kuncak, for developing the theory behind the “unbounded universal quantifier” problem.

To my colleagues who have taught Alloy in their courses, especially Matt Dwyer, John Hatcliff, Cesare Tinelli, and Michael Huth, who developed extensive material when Alloy was much rougher than it is today.

To the readers who gave me comments and suggestions on drafts of the book: Paul Attie, Daniel Le Berre, Paulo Borba, Jin Song Dong, Rohit Gheyi, Tony Hoare, Michael Lutz, Tiago Massoni, Walden Mathews, Joe Moore, Sanjai Narain, David Naumann, Norman Ramsey, Mark Saahtink, Martyn Thomas, and Mandana Vaziri; and especially to Michael Jackson, Jeremy Jacob, Viktor Kuncak, Butler Lampson, Chris Wallace, David Wilczynski, and Pamela Zave, who read the book in its entirety and together found something to fix on almost every page. They have saved me from many embarrassments and the reader from countless frustrations and confusions.

To the National Science Foundation, NASA, IBM, Microsoft, and Doug and Pat Ross, for their support of my research.

To Rod Brooks, Eric Grimson, John Guttag, Rafael Reif, and Victor Zue, for their role in creating the wonderful research and teaching environment that nurtured this work.

To Michael Butler, John Fitzgerald, Martin Gogolla, Peter Gorm Larsen, and Jim Woodcock for contributing solutions in their own languages to the hotel locking problem for appendix E.

To Bob Prior at MIT Press, for his confidence in this book, and his sage advice; to Katherine Almeida, its editor; and to Yasuyo Iguchi, design manager, for her advice on typography.

To my father, Michael Jackson, for his endless encouragement; for the inspiration he has been for me since I joined the family business; and for his tolerance of so many papers, and now a book, where rigor in logic often seems to take precedence over rigor in method. To my mother, Judy Jackson, the most prolific author in the family, whose uplifting emails continued to come even when replies became short and infrequent. To my brother, Adam Jackson, who insisted that my text be optically aligned (and showed me how to do it).

And finally, to my wife Claudia, to whom I dedicate this book, who has taught me so much, especially that analysis isn't everything (and that the New Yorker is much more fun than the Economist). And to my children Rachel, Rebecca and Akiva, who will grow up, I hope, in a world of better and simpler software than we have today.

Acknowledgments

(revised edition)

This new edition brings the book up to date with Alloy 4, the current version of the language. While the changes to the syntax of the language are very minor—consisting primarily of simplifications that make Alloy cleaner and more uniform—the changes to the Alloy Analyzer have been more dramatic. Emina Torlak built a new model finding engine for the Analyzer that improves its performance, often by an order of magnitude. Felix Chang rebuilt the entire front end, and was the one who wisely suggested that the grammar of Alloy might be simplified. Since then, Felix has become a hero to the Alloy community, answering questions online, and posting a comprehensive list of changes to the book—now incorporated in this edition—to bring the examples up to date with Alloy 4. I am very grateful to them both.

Thank you also to my colleagues who found errors in the first edition and made many helpful suggestions: Jasmin Blanchette, Shriram Krishnamurthi, Carroll Morgan, Tobias Nipkow, Carlos Pacheco, Burkhardt Renz, Michael Sperberg-McQueen and Kevin Sullivan.

A Japanese edition of the book has been based on this revised English edition. I'd like to thank the translators—Takeo Imai, Masahiro Sakai, Yusuke Endoh, Josh Kataoka—for their quite extraordinary scrutiny of the English text. They caught typos nobody else had noticed; spotted subtle inconsistencies of typography; and made many good suggestions to improve wording. They also helped me clarify some of the exercises, make the terminology more uniform, and—best of all—drew my attention to some technical issues (for example, the first-order axiomatizability of transitive closure in finite models) that called for revisions to the text. As I worked on this revised edition, they kept me constantly on my toes. While translation into Japanese, like formalization, might be a very productive way to identify flaws, it is by no means a lightweight approach, and I am very grateful to them for their boundless enthusiasm and patience. I'd also like to thank the reviewers of the Japanese translation—especially Akira Tanaka, Masateru Kawaguchi, Keigo Imai and Kazuhiro Inaba—who found additional errors in the technical material, and Shin Nakajima, who supervised the translation.

My current students, Eunsuk Kang, Aleks Millicevic and Joseph Near, have maintained the vitality of Alloy research at MIT, and helped proof-read the new material. Aleks and Joe suggested some nice simplifications to the treatment of integers that have been incorporated into the Alloy language, allowing me to delete several sections from the book that explained complications that are no longer necessary.

Finally, I am grateful to the National Science Foundation for supporting the ongoing development of Alloy, as well as my sabbatical effort in preparing this new edition of the book.