Dependable Software by Design

Computers fly our airliners and run most of the world's banking, communications, retail and manufacturing systems. Now powerful analysis tools will at last help software engineers ensure the reliability of their designs

By Daniel Jackson

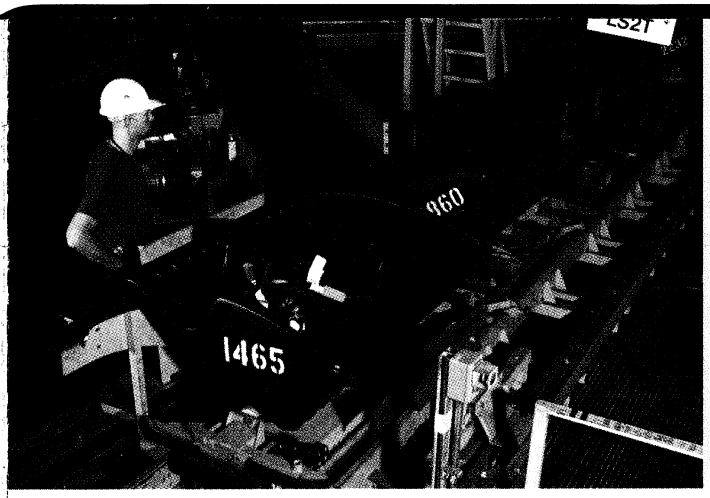

Almost all grave software problems can be traced to conceptual mistakes made before programming started.

Failed automated baggage system at Denver International Airport.

An architectural marvel when it opened 11 years ago, the new Denver International Airport's high-tech jewel was to be its automated baggage handler. It would autonomously route luggage around 26 miles of conveyors for rapid, seamless delivery to planes and passengers. But software problems dogged the system, delaying the airport's opening by 16 months and adding hundreds of millions of dollars in cost overruns. Despite years of tweaking, it never ran reliably. Last summer airport managers finally pulled the plug—reverting to traditional manually loaded baggage carts and tugs with human drivers. The mechanized handler's designer, BAE Automated Systems, was liquidated, and United Airlines, its principal user, slipped into bankruptcy, in part because of the mess.

The high price of poor software design is paid daily by millions of frustrated users. Other notorious cases include costly debacles at the U.S. Internal Revenue Service (a failed $4-billion modernization effort in 1997, followed by an equally troubled $8-billion updating project); the Federal Bureau of Investigation (a $170-million virtual case-file management system was scrapped in 2005); and the Federal Aviation Administration (a lingering and still unsuccessful attempt to renovate its aging air-traffic control system).

Such massive failures occur because crucial design flaws are discovered too late. Only after programmers began building the code—the instructions a computer uses to execute a program—do they discover the inadequacy of their designs. Sometimes a fatal inconsistency or omission is at fault, but more often the overall design is vague and poorly thought out. As the code grows with the addition of piecemeal fixes, a detailed design structure indeed emerges—but it is a design full of special cases and loopholes, without coherent principles. As in a building, when the software's foundation is unsound, the resulting structure is unstable.

Managers involved in high-profile software blowouts could claim in their defense that they followed standard industry practices, and unfortunately they would be right. Developers rarely articulate their designs precisely and analyze them to check that they embody the desired properties. But with computers now flying airplanes, driving trains and cars, and running most of the financial, communications, trading and production machinery of the world, society has an urgent need to improve software dependability.

Now a new generation of software design tools is emerging [*see box on page 74*]. Their analysis engines are similar in principle to tools that engineers increasingly use to check computer hardware designs. A developer models a software design using a high-level (summary) coding notation and then applies a tool that explores billions of possible executions of the system, looking for unusual conditions that would cause it to behave in an unexpected way. This process catches subtle flaws in the design before it is even coded, but more important, it results in a design that is precise, robust and thoroughly exercised. One example of such a tool is Alloy, which my research group and I constructed. Alloy (which is freely available on the Web) has proved useful in applications as varied as avionics software, telephony, cryptographic systems and the design of machines used in cancer therapy [*see box on page 73*].

Alloy and related design-checking tools build on a quarter of a century of existing research into ways to prove mathematically whether programs are correct. But rather than requiring proofs to be done by hand, they employ automated reasoning techniques that treat a software design problem as a giant puzzle to be solved. These analyzers operate on designs, not program code, so they

Overview/*Software Design Checkers*

- **Despite the ever increasing importance of computer software in our daily lives, software engineers rarely analyze their designs to ensure reliability. That situation is starting to change with the recent development of software design checking tools such as Alloy.**
- **Alloy combines a language that eases the modeling of complex software designs with an analysis engine that checks extensively for conceptual and structural flaws in an automated fashion, treating designs as huge puzzles to be solved.**
- **In the relatively near future, tools similar to Alloy will greatly improve the dependability of software by basing program development on more robust and constructive design practices.**

CHRIS CARTER *Denver International Airport*

ALLOY IN ACTION

Alloy helps software designers find and fix design flaws by providing both a language that clarifies a program's structure and an automated analyzer that searches the vast numbers of possible executions of a system for a "counterexample" that shows how it could fail to behave as desired. In the simplified example below, an engineer uses Alloy to evaluate the design of a file system—the software that organizes your computer files into folders and stores them on a disk. A crucial task for Alloy is to work out the effects various operations would have on the file structure. Here is how a designer might model and check the operation that moves a folder, or "directory," from one location in the file hierarchy to another.

STEP 1: DEFINE THE OBJECTS

The designer identifies the system's objects—files, directories and the file system as a whole—and their relations with one another. The Alloy model says the file system (FS) has three components: "files" (its set of files), "dirs" (its set of directories) and "contains" (a mapping that gives, for each directory, the set of files and directories it contains).

ALLOY CODE

```
module filesystem

abstract sig Object {}
sig File, Dir extends Object {}

sig FS {
 dirs: set Dir,
 files: set File,
 contains: dirs -> (dirs + files)
 }
```

EFFECT

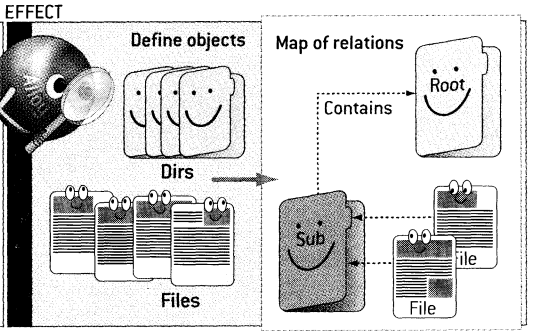

STEP 2: MODEL THE OPERATION

Next, the designer models the move ("move_dir") of the file system before ("fs") to a file system after ("fs'"). The operation involves two directories: "d," the directory being moved, and "to," the place it is being moved to—its new parent. Three constraints follow, which describe the intended effect, on three separate lines: First, both the moved object and its new location are directories of the file system. Second comes the essence of the operation: it says that the new containment mapping is the old one, with every mapping from a directory to "d" removed, and the mapping from "to" to "d" added. The third line says that nothing else changes.

```
pred move_dir (fs, fs': FS, d, to: Dir) {
d + to in fs.dirs
fs'.contains = fs.contains - Dir->d + to->d
fs'.files = fs.files and fs'.dirs = fs.dirs
 }
```

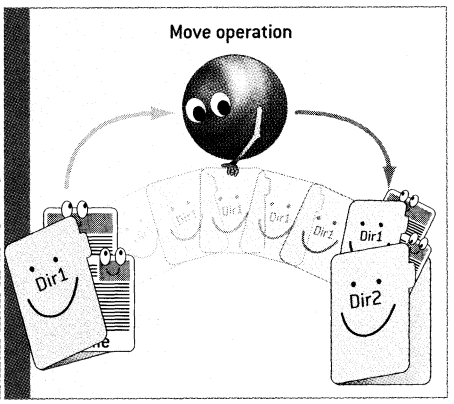

STEP 3: SPECIFY REQUIREMENTS

The designer then formulates a crucial requirement: every file and directory should be "reachable" (have a pathway) from some root. This is recorded in the Alloy model as an "assertion" (called "move_OK"), which says that executing the move operation does not make a file or directory unreachable from a root.

```
pred reachable (fs: FS) {
 some root: fs.dirs | fs.(dirs+files) in root.*(fs.contains)
 }

assert move_OK {
 all fs, fs': FS, d, to: Dir |
   reachable (fs) and move_dir (fs, fs', d, to) implies reachable (fs')
 }
```

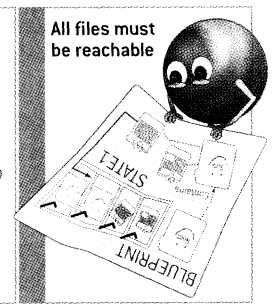

STEP 4: FIND AND FIX THE FLAW

Alloy executes "check move_OK" by generating all possible states of the system (up to a certain size) and checking the assertion for each—thus simulating possible moves as they might occur when the software is run. Alloy finds a counterexample to the assertion—a directory that can be moved to itself. The action would disconnect the directory from a root, making it unreachable. As a remedy, a designer could add a new constraint disallowing a directory to move to itself or any of its descendants.

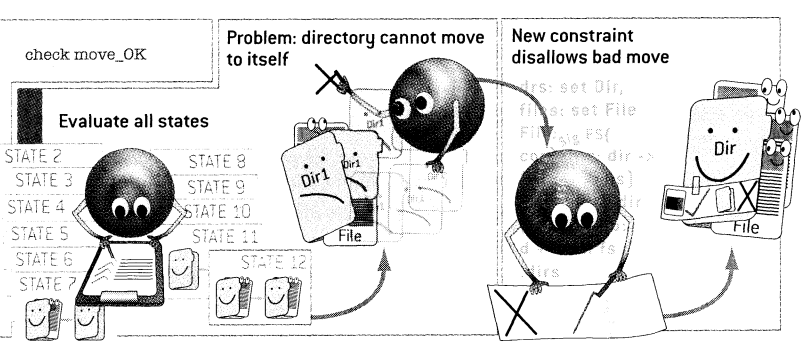

LUCY READING-IKKANDA

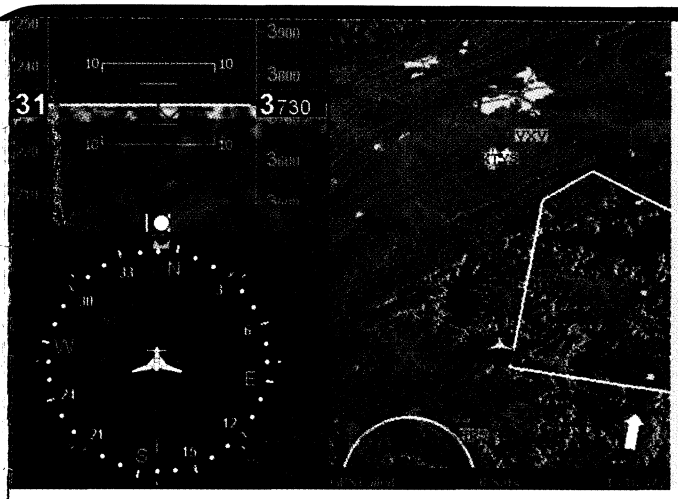

The idea is to simulate every state that the software can take to determine that none leads to a failure.

Alloy helped to make an avionics system hacker-proof.

cannot guarantee that a program will not crash. But they potentially offer software engineers the first practical tools to ensure that designs are robust and free from conceptual flaws and thus provide a firm foundation on which to build reliable software systems.

Evaluating Designs

BAD SOFTWARE is not a new problem. Warnings of a software crisis go back to the 1960s and have only intensified as computers have been woven into the fabric of society [see "Software's Chronic Crisis," by W. Wayt Gibbs; SCIENTIFIC AMERICAN, September 1994].

Today most software typically is debugged and refined by testing. Human engineers run the program using a wide range of starting conditions (or inputs) to see if it operates as expected. Although the practice catches a raft of small flaws, it often overlooks faults in the basic design of the software. In some sense, these test procedures miss the (diseased) forest for the (rotting) trees.

What is worse, bugs "fixed" during the testing process often exacerbate design problems. As programmers debug the code and insert new features, the software invariably grows barnacles of complexity, creating more opportunities for errors and inefficient operation. This situation is reminiscent of the (incorrect) Ptolemaic theory of planetary motion first developed by the ancient Greeks. In the Middle Ages, as observations showed the predictions to be inaccurate, astronomers adjusted Ptolemy's system, which relied on epicycles. When that proved insufficient, they resorted to adding epicycles to the epicycles. Further fine-tuning over the centuries never solved the problem, because the initial concept was fatally flawed.

Similarly, bad software tends to get more and more complicated and less and less reliable, however much time and money are poured into improving it. It is well known that serious problems with software systems rarely arise from programming errors; almost all grave difficulties can be traced back to conceptual mistakes made before programming even started. In contrast, a small amount of modeling and analysis during the initial determination of requirements, specifications, or program design costs only a tiny fraction of the price tag of checking all the code but provides a large part of the benefit gained from an exhaustive analysis. Focusing on design early saves costly headaches down the road.

Design tools for software have been slow in coming because software does not obey physical laws. Because computer programs are in essence mathematical objects whose values are constructed from bits, software programs are discrete (particlelike) rather than continuous. A mechanical engineer can stress a component with a large force and assume that if it survives it will not fail when subjected to a slightly smaller force. When an object is subject to the (mostly continuous) principles of the physical world, a small change in one quantity generally produces a small change in another. Unfortunately, no such generalities apply to software: one cannot extrapolate between test cases. If one chunk of software works, that fact says nothing about the operations of a similar chunk of code; they are discrete and separate.

In the early days of computer science, researchers hoped that programmers might prove their codings were correct in the same way that mathematicians prove their theorems. With no way to automate the many steps involved, however, a human expert had to do much of the work. These so-called heavy-duty formal methods were impractical except for relatively modest but especially critical pieces of software, such as an algorithm for controlling railroad intersections.

More recently, researchers have adopted a very different approach, one that harnesses the power of today's faster processors to test every possible scenario. This method, known as model checking, is now used extensively to verify integrated-circuit designs. The idea is to simulate every possible sequence of states (the conditions of the system at specific times) that might arise in practice and to determine that none leads to a failure. For a microchip design, the number of states to evaluate is often huge: 10^{100} or more. The challenge is far more stringent for software. But clever encoding techniques (by which large sets

THE AUTHOR

DANIEL JACKSON leads the Computer Science and Artificial Intelligence Laboratory's Software Design Group at the Massachusetts Institute of Technology. His main research interest is software engineering, with a focus on software design, specification and analysis, particularly of critical systems. Jackson received an M.A. from the University of Oxford in physics, and his S.M. and Ph.D. from M.I.T. in computer science. Before his professorship at M.I.T., he taught at Carnegie Mellon University. An avid photographer, Jackson recently exhibited his work at the Newton Free Library outside of Boston.

GREG RICHTER *Blue Mountain Avionics*

of software states can be represented very compactly) make it possible to check every state by considering these large sets simultaneously.

Model checking alone regrettably cannot handle states with complex structures, which is characteristic of most software designs. My research colleagues and I have developed an approach that shares the same spirit yet employs a different mechanism. Like model checking, it considers all possible scenarios (although in truth, some bounds need to be introduced to keep the problem finite, because software is not restricted by the physical limitations imposed by hardware). Unlike model checking, however, our technique does not examine scenarios in their entirety, one at a time. Instead it searches for a bad scenario—one that results in failure—by filling in each state in an automated fashion, one bit at a time, in no particular order.

The process is in some sense comparable to a robotic arm fitting each piece of a jigsaw puzzle into place one by one until the completed image finally emerges. If that image corresponds to a bad scenario, Alloy would have done its job. Alloy thus treats design analysis as if it were a puzzle to be solved. Some other recently developed software model checkers work this way as well.

The Solution Is a Puzzle

TO UNDERSTAND HOW Alloy solves software design puzzles, it helps to consider an old riddle: A farmer goes to market where he buys a fox, a goose and a bag of corn. On his way home, he has to carry his goods across a river by boat. The skiff will hold only the man and one purchase at a time, however. Herein lies a problem: if left unsupervised, the fox would eat the goose and the goose would eat the corn. So how does the farmer get all of his goods to the far bank intact?

This variety of puzzle involves finding scenarios that satisfy a collection of constraints. Mentally we do this task by imagining a series of steps: The farmer transports the goose first; on the next trip, he takes the fox, whereupon he brings back the goose and then, leaving it behind, crosses with the corn; he then returns to fetch the goose. By checking whether each step satisfies the constraints, we ensure that each item remains safe.

A successful software design imposes a similar, though much more complicated, array of rules. To be useful, a design-checking tool must be able to find counterexamples: solutions to the puzzle that meet all the "good" constraints (and thus could occur when the program is run) and an additional "bad" constraint (and thus yield an unacceptable outcome). If any such counterexamples turn up, they will reveal flaws in the design. So whereas the puzzle solver is happy to find a solution to the "farmer's dilemma," a solution to a software

Debugging Cancer Therapy Machines

Modern medical devices rely on software for almost every aspect of their operation. In a machine used for cancer therapy, even the "emergency stop" button is not an actual electrical switch but a software program: hitting it causes about 15,000 lines of code to execute and shut the system down—unless, of course, there is a bug or design flaw in the software. That is where Alloy comes in—it analyzes programs to find the design problems.

Working with the developers of a cancer-therapy system, for example, we have used Alloy to explore the design of some of its features. In one case, we took a design for a new scheduling system that determines the treatment room to which the beam is sent. We set Alloy to look for scenarios in which interactions between the operator in the main control room and the therapists in the treatment rooms would produce unexpected results. Alloy found various scenarios that had not been anticipated originally.

In another case, we applied Alloy to the design of an elaborate protocol for positioning the patient under the proton beam, which turned out to have a subtle and unexpected consequence: the angle of the gantry crept around over time, even when it was not being intentionally adjusted. With a small Alloy model we showed how, by choosing the right abstractions, this problem could be reduced to the same, rather simple problem as that for designing a car accessory system that remembers driver-seat positions. In fact, the therapy system has many safeguards and the gantry movement was not a dangerous problem. But if the correct abstractions had been used from the start, the design would have been much simpler and operating the software considerably easier. *—D.J.*

CORRECT POSITION of a patient—controlled by software—is critical to control radiation dosage in a cancer therapy machine. Alloy helped to improve the software design for a similar machine.

design puzzle is bad news: it means that an undesirable scenario exists and the design is defective. In practice, the counterexample might not itself lead to any problems. It may instead reveal a discrepancy in how the designer originally characterized the unacceptable outcomes. Either way something needs to be fixed—the design or the designer's expectations.

The great difficulty in searching for counterexamples is that the number of potential scenarios in a software design of even moderate complexity is typically vast, but only a tiny proportion correspond to counterexamples. Imagine trying to plan who sits next to whom at a wedding reception. If all attendees get along, the solution is trivial. Throw in a few ex-spouses who require separation, and the problem gets trickier. Now consider the seating chart for Romeo and Juliet's reception. If there are 20 seats and any of 10 guests can sit in each, that makes 10^{20} possible combinations. Even checking a billion scenarios per second, a computer would take 3,000 years to explore them all.

In the 1980s, researchers identified problems of this form as a special class of problems that, in the worst case, can be solved only by enumerating all possible scenarios. But in the past decade, with new search strategies and algorithms and by building on ever increasing computational power, researchers have developed tools called SAT (*sat*isfiability) solvers that can handle these problems fairly easily. Many are now freely available and can often solve problems with millions of constraints.

Importance of Abstraction

AS ITS NAME SUGGESTS, Alloy melds two elements that help make software designs more robust. One is a new language that helps to elucidate the structure and behavior of the software design. The other is an automated analyzer (which incorporates a SAT solver) to hash through a multitude of possible scenarios.

The first step in applying Alloy is to create a model of the design: not the rough sketch or flowchart typical in software engineering but a precise model that spells out the "moving parts" and specific behaviors, both desired and undesired, of the system and its components. A software engineer first writes down definitions of the various kinds of objects in the design, then groups those objects into mathematical sets: collections of things that are alike in their structure and behavior (for example, the set of all Capulets) and linked by mathematical relations (such as the relation that associates guests sitting next to one another).

Next come facts that constrain these sets and relations. In a software design, the facts include the mechanism of the software system and assumptions about other components (say, statements about how human users are expected to behave). Some of these facts are simple assumptions—for example, that nobody is both a Capulet and a Montague and that every guest sits next to exactly two other guests. Some of them reflect the design itself: in our seating planner, for instance, the rule that each table, with the exception of the top table, is assigned either to one family or the other.

Finally, there are assertions, which are constraints that are expected to follow from the facts. In our example, with

Tools for Checking Software Designs

Computer scientists have developed a new generation of software design checking tools (in addition to Alloy) that programmers can use to analyze and test their codings for structural and conceptual inconsistencies that could lead to system failure. In general, these commercial and open-source design-evaluation tools are based on specialized high-level languages (notations that summarize blocks of code) that researchers have developed to ease the specification, modeling and simulation of different types of software schemes.

Such tools incorporate automated analysis engines that explore the huge number of potential executions of systems for subtle design flaws that would cause them to behave in undesirable ways (an instance of which is called a counterexample). These software design tools often include facilities that can help designers visualize counterexamples or relations between blocks of code.

LANGUAGE	TOOL	SOURCE	WEB SITE
B	B-Toolkit	B-Core	www.b-core.com
	Atelier-B	Steria	www.atelierb.societe.com
	Pro-B	University of Southampton	www.ecs.soton.ac.uk/~mal/systems/prob.html
CSP	FDR	Formal Systems Europe	www.fsel.com
FSP	LTSA	Imperial College London	www.doc.ic.ac.uk/~jnm/book/ltsa/LTSA.html
Lotos	CADP	INRIA Research Institute	www.inrialpes.fr/vasy/cadp/
OCL	USE	University of Bremen	www.db.informatik.uni-bremen.de/projects/USE/
PROMELA	Spin	Bell Laboratories	spinroot.com/
Statecharts	Statemate	I-Logix	www.ilogix.com
VDM	VDMTools	CSK Corp.	www.csk.com/support_e/vdm/
			www.vdmbook.com/tools.php
Z	Jaza	University of Waikato	www.cs.waikato.ac.nz/~marku/jaza/
Zing	Zing	Microsoft Research	research.microsoft.com/zing/

Alloy has uncovered serious deficiencies in published software designs.

Alloy checked a software program that finds printers on wireless networks.

the exception of Romeo and Juliet, no Capulet should be seated next to a Montague. The assertions say that the system can never get into certain undesirable states and that specific bad sequences of events can never occur.

The analyzer component of Alloy harnesses a SAT solver to search for counterexamples—possible scenarios of the software system that are permitted by its design but that fail a sanity check (which is accomplished by writing assertions that must be true if the model is correctly designed). In other words, the tool attempts to construct situations that satisfy the facts but violate a stated assertion. In our case, it would generate a seating plan in which a Capulet (other than Juliet) sits next to a Montague (other than Romeo) at the top table. To fix the seating rule, we can add a new fact: that Romeo and Juliet occupy the top table alone. Now Alloy would find no counterexample.

Together the declarations of the sets and relations, the facts, and the assertions make up an abstraction that captures the essence of the software design. Writing all this out makes the limitations of the design explicit and forces engineers to think hard about exactly which abstractions will work best. Bad abstraction choices lie at the root of many unnecessarily complicated or unreliable systems.

Systems that rely on software built on simple and robust abstractions should also be easier to use. Consider how e-ticketing simplified air travel, how universal product codes made shopping easier or how 800-number-based conference calls made teleconferencing more feasible. Each of these innovations stemmed from a transformation in the basic abstractions embodied in the underlying software.

The Road to Reliability

TOOLS AKIN TO ALLOY are currently used primarily in research and in cutting-edge industrial settings. The technology has been employed to explore new architectures for telephone switching systems, to design avionics processors that are secure against hackers and to describe access-control policies for communications networks. We have used it to check widely used and robust software devices, such as protocols for finding printers on networks and tools for synchronizing files across machines.

In addition, Alloy has uncovered serious deficiencies in published software designs—such as a key management protocol that was supposed to enforce special-access rules based on membership in a group but turned out to grant access to former members who should have been rejected. It is noteworthy that many programmers who have used Alloy have been surprised by the number of flaws that the tool turns up in the designs for even their simplest applications.

It is most likely only a matter of time until tools resembling Alloy are adopted more widely in industry. Improvements in the underlying SAT solvers will make analysis tools faster and better able to handle very large systems. Meanwhile a new generation of software designers, educated in these methods, will incorporate them into their work. Modeling is growing in popularity, particularly among managers desperate to see some description of a software system's design beyond the code itself.

At some point, there may come a time when software becomes so essential to our day-to-day infrastructure that society will no longer tolerate bad software. As a result, governments may even establish inspection and licensing regulations that enforce high-quality program construction techniques. One day, perhaps, software systems will be truly robust, predictable and easy to use—by design. SA

MORE TO EXPLORE

Exploring the Design of an Intentional Naming Scheme with an Automatic Constraint Analyzer. Sarfraz Khurshid and Daniel Jackson in *Proceedings of the 15th IEEE International Conference on Automated Software Engineering, Grenoble, France.* IEEE, September 2000. (Describes application of Alloy to the design of a system for finding resources on a network.)

Automating First-Order Relational Logic. Daniel Jackson in *Proceedings of the 8th ACM SIGSOFT International Symposium on Foundations of Software Engineering: Twenty-First Century Applications.* ACM Press, 2000. (Explains Alloy's analysis.)

A Micromodularity Mechanism. Daniel Jackson, Ilya Shlyakhter and Manu Sridharan in *Proceedings of the Joint 8th European Software Engineering Conference (ESEC) and 9th ACM SIGSOFT Symposium on the Foundations of Software Engineering.* ACM Press, 2001. (Explains key concept in the latest version of Alloy language.)

Alloy: A Lightweight Object Modeling Notation. Daniel Jackson in *ACM Transactions on Software Engineering and Methodology,* Vol. 11, Issue 2, pages 256–290; April 2002. (Original description of Alloy.)

Software Abstractions: Logic, Language, and Analysis. Daniel Jackson. MIT Press, 2006.

Daniel Jackson's Web site: **http://people.csail.mit.edu/dnj/**

Alloy Web site: **http://alloy.mit.edu**

EMILY HARRISON