

Automatic Detection and Diagnosis of Faults in Generated Code for Procedure Calls

Mark W. Bailey, *Member, IEEE*, and Jack W. Davidson, *Member, IEEE Computer Society*

Abstract—In this paper, we present a compiler testing technique that closes the gap between existing compiler implementations and correct compilers. Using formal specifications of procedure-calling conventions, we have built a target-sensitive test suite generator that builds test cases for a specific aspect of compiler code generators: the procedure-calling sequence generator. By exercising compilers with these specification-derived target-specific test suites, our automated testing tool has exposed bugs in every compiler tested on the MIPS and one compiler on the SPARC. These compilers include some that have been in heavy use for many years. Once a fault has been detected, the system can often suggest the nature of the problem. The testing system is an invaluable tool for detecting, isolating, and correcting faults in today's compilers.

Index Terms—Target-sensitive test suite generation, automatic fault isolation, procedure-calling convention, code generation, compiler testing and debugging.

1 INTRODUCTION

BUILDING compilers that generate correct code is difficult. To achieve this goal, compiler writers rely on automated compiler building tools and thorough testing. Automated tools, such as parser generators, take a specification of a task and generate implementations that are more robust than hand-coded implementations. Conversely, testing tries to make hand-coded implementations more robust by detecting errors. One aspect of a compiler that has traditionally been hand-coded is the portion that generates *calling sequences*—implementations of procedure calls. We have developed a language, called CCL, for specifying procedure-calling conventions. CCL specifications are used to automatically generate calling sequences for the *vpcc/vpo* retargetable optimizing compiler [1]. While experimenting with CCL, we realized that the descriptions could be used to make other compilers more robust without requiring that the compiler implementation use CCL. In this paper, we describe how CCL's underlying finite state machine model can be used to generate tests for hand-coded calling sequence generators in other compilers. This technique has exposed a number of calling convention errors in production-quality compilers that have been used heavily for years. Although the convention examples used here were originally specified using CCL, we omit a description of CCL both for brevity and since the generated finite state machine tables used here serve as equivalent specifications of convention behavior. Only the knowledge of how these automata model convention behavior is

necessary to understand our testing technique or to reap its benefits. However, readers interested in the CCL language and the automatic translation of CCL specifications to finite state automata can find these details in a previous paper [1].

In this paper, we describe several contributions. First, we present a method for automatically testing implementations of procedure-calling conventions. Using this technique, we have found bugs in mature C compilers. This approach, which uses a formal model of procedure-calling conventions, methodically generates tests that offer complete coverage of the specified convention. Second, we introduce an algorithm for intelligently selecting important tests from the complete coverage suite. These tests include boundary cases that are more likely to reveal bugs than exhaustive or randomly generated tests. Third, because the tests focus only on the calling convention, they isolate errors more effectively than tests from a general test suite. Fourth, we describe a method for automatically diagnosing the nature of some types of faults. Finally, we describe a method for quickly determining the conformance of multiple compilers at once.

2 PROCEDURE CALLING CONVENTIONS

An important feature of high-level programming languages that compilers must implement is the procedure call. The interface between procedures facilitates separate compilation of program modules and interoperability of programming languages. This is accomplished by defining a *procedure-calling convention* that dictates the way that program values are communicated and how machine resources are shared between a procedure making a call (the *caller*) and the procedure being called (the *callee*). The calling convention is machine-dependent because the rules for passing values from one procedure to another depend on machine-specific features such as memory alignment restrictions and register usage conventions. The code that

• M. Bailey is with the Department of Computer Science, Hamilton College, 198 College Hill Road, Clinton, NY 13323.
E-mail: mbailey@hamilton.edu.

• J. Davidson is with the Department of Computer Science, University of Virginia, 151 Engineer's Way, Charlottesville, VA 22911.
E-mail: jwd@virginia.edu.

Manuscript received 3 July 2002; revised 24 July 2003; accepted 6 Aug. 2003.
Recommended for acceptance by T. Reps.

For information on obtaining reprints of this article, please send e-mail to: tse@computer.org, and reference IEEECS Log Number 116879.

- 1) Registers R1, R2, R3, and R4 are 32-bit argument-transmitting registers.
- 2) Arguments are also passed on the stack in increasing memory locations starting at the stack pointer.
- 3) An argument may have type `char` (1 byte), `int` (4 bytes), or `double` (8 bytes).
- 4) An argument is passed in registers (if enough are available to hold the entire argument), and then on the stack.
- 5) Arguments of type `int` are 4-byte aligned on the stack.
- 6) Arguments of type `double` are 8-byte aligned on the stack.
- 7) Stack elements that are skipped over cannot be allocated later.
- 8) Return values are passed in registers R1 and R2.
- 9) Values of registers R6, R7, R8, and R9 must be preserved across a procedure call.

Fig. 1. Rules for a simple calling convention.

implements the calling convention, known as the *calling sequence* [2], must be generated by the code generator. This aspect of the code generator, which we name the *calling sequence generator*, is a source of great difficulty for the compiler writer because it not only suffers from being hand-coded, it also changes each time the compiler is retargeted.

2.1 A Simple Calling Convention

To aid in our discussion of calling conventions, we use a simplified example calling convention. Fig. 1 contains the calling convention rules for a hypothetical machine. Consider the following ANSI C prototype for a function `warp`:

```
int warp(char p1, int p2, int p3, double p4);
```

For the purpose of transmitting procedure arguments for our simple convention, we only consider the *signature* of the procedure. We define a procedure's signature to be the procedure's name, the order and types of its arguments, and its return type. This is analogous to ANSI C's *abstract declarator* [3], which, for the previous function prototype, is:

```
int warp(char, int, int, double);
```

which defines a function that takes four arguments (a `char`, two `int`'s, and a `double`), and returns an `int`.

With `warp`'s signature, we can apply the calling convention in Fig. 1 to determine how to call `warp`. Arguments to `warp` would be placed in the following locations:

- `p1` in register R1,
- `p2` in register R2,
- `p3` in register R3, and
- `p4` on the stack at offsets 0-7.

Notice that, although register R4 is available, `p4` is placed on the stack since it cannot be placed completely in argument-transmitting registers (rule 4). Such restrictions are common in actual calling conventions.

2.2 Convention to Implementation

Once the calling convention has been established, a compiler can be targeted to generate the calling sequence code that implements the procedure calls for the source

language. Traditionally, this code has been handcrafted. In contrast, we use a calling convention specification and an interpreter. The interpreter can generate tables that can be used in the calling-convention-specific portion of *vpcc/vpo* [1] or in a test suite generator. The test suite generator uses information from the table to tailor the test suite to the specific calling convention. The test suite can either be used to confirm that the *vpcc/vpo* implementation properly uses the convention tables or that another, independent compiler conforms to the convention described in the CCL specification. In the next section, we describe the formalism that we use to capture convention details.

3 THE FORMAL MODEL

We use finite state automata to model a calling convention's placement of arguments (and return values) in a machine's memory locations. The use of FSAs for modeling parts of a compiler and as an implementation tool has a long and successful history. For example, FSAs have often been used to implement lexical analyzers [4]. More recently, Proebsting and Fraser [5] and Müller [6] have used finite state automata to model and detect structural hazards in pipelines for instruction scheduling.

The FSA model characterizes only the placement of argument descriptors regardless of type or passing mechanism. For pass-by-value, the placement of the descriptor is the same as the placement of the actual argument value. For pass-by-reference and more complicated passing mechanisms, the model describes the placement of the reference, not the actual argument value. The model does not describe how a descriptor is to be interpreted once it has been transmitted from caller to callee or vice versa. If, however, the caller and callee do not agree on how the descriptor is to be interpreted, this fault will be detected since the argument will not be properly transmitted to the callee.

3.1 P-FSA Representation

An example FSA that we use to model calling convention placement is shown in Fig. 2. This FSA models the placement of procedure arguments for the simple calling convention described in Fig. 1. A placement FSA (P-FSA) takes a procedure's signature as input and produces locations for the procedure's arguments as output. The automaton works by moving from state to state as the location of each value is determined. When a transition is used to move from one state to the next, information about the current parameter is read from the input and the resulting location is written to the output.

The states of the machine represent the state of allocation for the machine's memory resources. For example, the state q_2 (labeled 1100 000) represents the fact that registers R1 and R2 have been allocated (the first two bits: 11), that registers R3 and R4 have not been allocated (the second two bits: 00), and that the stack pointer is currently eight-byte aligned (the remaining three bits: 000). A transition between states represents the placement of a single argument. Since arguments of different types and sizes impose different demands on the machine's resources, we may find more than one transition leaving a particular state. In our example, q_8 has three transitions even though two of them

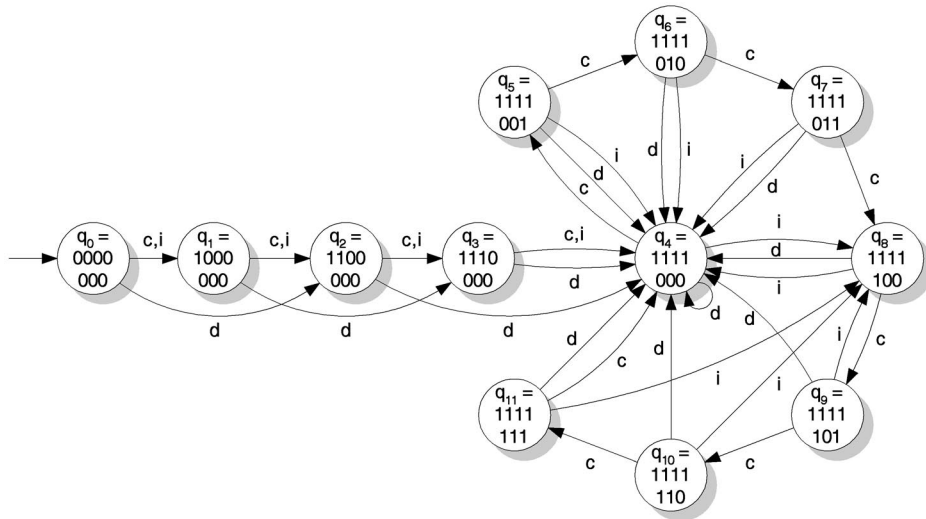


Fig. 2. P-FSA for transmission of parameters for a simple calling convention.

(`int` and `double`) have the same target state (q_4). This duplication is required since the output from mapping an `int` is different from the output from mapping a `double`.

In most language implementations, the runtime stack is modeled as an infinite or boundless resource. Modeling the allocation of an infinite resource using an FSA poses a problem, however. As mentioned previously, the state represents which resources have been allocated. For finite resources, this is easily accomplished by maintaining a bit vector. When a resource no longer may be used, the associated bit is set. For an infinite resource, this scheme cannot work if we hope to use an FSA since this would require a bit vector of infinite length. To simplify the problem, we impose a restriction on infinite resources: Their allocation must be contiguous.¹ Thus, for an infinite resource $I = \{i_1, i_2, \dots\}$, we can store the allocation state by maintaining an index p whose value corresponds to the index of the first available resource in I . Because the allocation of I must be contiguous, p partitions the resources since a resource i_j is unavailable if $j < p$ or available if $j \geq p$. For instance, if the stack is the infinite resource, p can be considered the stack pointer.

Nevertheless, we still have a problem. Although, for a particular machine, the value of p must be finite, the resulting FSA could have as many as 2^{32} stack allocation states for a 32-bit machine. However, we can significantly reduce this number by observing that the decision of where to place a parameter in memory is not based on p , but rather on alignment restrictions. For our example, we care only if the next available memory location is 1, 4, or 8-byte aligned. Consequently, we can capture the allocation state of the machine with three bits that distinguish the memory allocation states. We call these the *distinguishing* bits for infinite resource allocation.

1. P-FSAs could also handle a convention whose stack allocation was not continuous by treating the stack as a finite resource. Although this approach might require a very large bit-vector, any implementation of such a convention would have to use a bit-vector (or equivalent) to properly implement the convention. We know of no such compiler or convention.

Handling pass-by-value structures creates an analogous problem. Structures of different sizes allocate different amounts of space. Hence, each structure of a different size impacts the state of resource allocation differently. This implies that each P-FSA state requires an infinite number of exiting transitions; one for every different structure size. Fortunately, since only the “alignment state” of the stack pointer is of interest, we need only include transitions for structures that leave the P-FSA in a different state. So, for a convention that requires structures to be passed in 8-byte aligned memory locations, all structures of size n , where $n \bmod 8 = 1$, share the same transition out of a given state because they leave the alignment, p , in the same state. Therefore, the number of transitions leaving a state is limited by the alignment restrictions of the machine.

Placement functions are described in terms of finite resources, infinite resources, and selection criteria. A set of finite resources $R = \{r_1, r_2, \dots, r_n\}$ is used to represent machine registers, while an infinite resource $I = \{i_1, i_2, \dots\}$ ² is used to represent the stack. The *selection criteria* $C = \{c_1, c_2, \dots, c_m\}$ correspond to characteristics about arguments (such as their type and size) that the calling convention uses to select the appropriate location for a value. We encode the signature of a procedure with a tuple $w \in (C^*, C^*)$. The first element of the tuple contains zero or more³ return value criteria, while the second element contains zero or more parameter criteria. Each state q in the automaton is labeled according to the allocation state that it represents. The label includes a bit vector v of size n that encodes the allocation of each of the finite resources in R . Additionally, to express the state of allocation for the stack, we include d , the distinguishing bits that indicate the state of stack alignment. So, a state label is a string vd that indicates the resource allocation state. In our example convention, $n = 4$ and the length of the string d ($|d|$) is 3. So, each state is labeled by a string from the language $\{0, 1\}^4 \{0, 1\}^3$. The output of the machine is a string $s \in P$, where

2. This can easily be extended to model more than one infinite resource.
3. Supports languages that allow multiple return values.

TABLE 1
Definition of λ for Example P-FSA

λ	q_0	q_1	q_2	q_3	q_4	q_5	q_6	q_7	q_8	q_9	q_{10}	q_{11}
char	R1	R2	R3	R4	000	001	010	011	100	101	110	111
int	R1	R2	R3	R4	m_1^a	m_2^b	m_2	m_2	m_2	m_1	m_1	m_1
double	R1 R2	R2 R3	R3 R4	m_3^c	m_3	m_3	m_3	m_3	m_3	m_3	m_3	m_3

$$^a m_1 = 000\ 001\ 010\ 011$$

$$^b m_2 = 100\ 101\ 110\ 111$$

$$^c m_3 = 000\ 001\ 010\ 011\ 100\ 101\ 110\ 111$$

$$P = R \cup \{0, 1\}^{|d|}$$

which contains the placement information.

Since the P-FSA produces output on transitions, we have a Mealy machine [7]. We define a P-FSA, M , as a six-tuple⁴ $M = (Q, \Sigma, \Delta, \delta, \lambda, q_0)$, where:

- Q is the set of states with labels $\{0, 1\}^n \{0, 1\}^{|d|}$ representing the allocation state of machine resources.
- The input alphabet $\Sigma = C$ is the set of selection criteria.
- The output alphabet $\Delta = P$ is the set of memory location strings.
- The transition function $\delta : Q \times \Sigma \rightarrow Q$.
- The output function $\lambda : Q \times \Sigma \rightarrow \Delta^+$.
- q_0 is the state labeled by $0^n w$ where $|w| = |d|$, and w is the initial state of d .

We also define $\hat{\delta} : Q \times \Sigma^* \rightarrow Q$ and $\hat{\lambda} : Q \times \Sigma^* \rightarrow \Delta^*$, which are just string versions⁵ of δ and λ , respectively. So, for our example, we have

$$M = (Q, \{\text{char}, \text{int}, \text{double}\}, \{R1, R2, R3, R4\} \cup \{0, 1\}^3, \delta, \lambda, q_0),$$

where Q and δ are shown in Fig. 2 and λ is defined in Table 1. Note that we have modified the traditional definition of λ to allow multiple symbols to be output on a single transition. This reflects the fact that arguments can be located in more than one resource. For example, in state q_5 on an `int`, Table 1 indicates that M produces the string of four symbols 100 101 110 111 that designates four bytes that are four-byte aligned, but are not eight-byte aligned.

The signature:

```
int f(double, double, char, int);
```

will take the P-FSA in Fig. 2 along the path

$$q_0 \rightarrow q_2 \rightarrow q_4 \rightarrow q_5 \rightarrow q_4$$

producing the string (R1 R2) (R3 R4) (000) (100 101 110 111) along the way. The parentheses in the output string are required to determine where the placement of one argument ends and the next argument's placement begins. From

4. We use the notation of Hopcroft and Ullman for finite state automata and regular expressions [8]. We use letters early in the alphabet (a, b, c) to denote single symbols. Letters late in the alphabet (w, x, y, z) will denote strings of symbols.

5. Defined by Hopcroft and Ullman [8].

the string, we can derive the placement of `f`'s arguments. The first `double` is placed in registers R1 and R2, the second in registers R3 and R4, the `char` at the stack location with offset zero, and the `int` at the stack location with offset four.

3.2 Completeness and Consistency in P-FSAs

In our experience, we have encountered many recurring difficulties in the calling convention code generators of optimizing compilers for RISC machines.⁶ There are three sources for these problems: the convention specification, the convention implementation, and the implementation process. We address each of these in the following paragraphs.

Many problems arise from the method of convention specification. Often, no specification exists at all. Instead, the native compiler uses a convention that must be extracted by reverse engineering. In the cases where a specification exists, it typically takes the form of written prose or a few general rules, e.g., our example description in Fig. 1. Such methods of specification have obvious deficiencies. Furthermore, even if we have an accurate method for specifying a convention, it still may be possible to describe conventions that are internally inconsistent or incomplete. For example, the convention may require that more than one procedure argument be placed in a particular resource. Another possibility is that the specification may omit rules for a particular data type or combination of data types.

Those problems that do not arise from the specification result from incorrect implementation of the convention. Many of the same problems in the specification process also plague the implementation. Many conventions have numerous rules and exceptions that must be reflected in the implementation. Another difficulty is the implementation may require the use of the convention in several different locations. Maintaining a correspondence between the various implementations can itself be a great source of errors. Finally, this problem is exacerbated by the fact that the implementation frequently undergoes incremental development. Rather than taking on the chore of implementing the entire convention at once, a single aspect of the convention, such as providing support for a single data type, is tackled. After successfully implementing this subset, the next increment is undertaken. During this

6. Unlike many CISC machines, RISC machines typically increase the complexity of calling conventions by requiring that procedures pass parameters in registers whenever possible.

process, some aspect of the first stage may break due to the interactions between the two pieces.

The result of these observations is that there are several properties that we would like to ensure about a specification and implementation. The preceding discussion motivates the following categories of questions:

- **Completeness:**
 - Does the specified convention handle any number of arguments?
 - Does the convention handle any combination of argument types?
- **Consistency:**
 - Does the convention map more than one argument to a single machine resource?
 - Do the caller and callee's implementations agree on the convention?

Many questions like these can be answered using P-FSAs. The following sections show how we can prove certain properties about P-FSAs that ensure desirable responses to the preceding questions.

3.2.1 Completeness

The completeness properties address how well the convention covers the possible input cases. A convention must handle any procedure signature. If we could guarantee that the convention was complete, or covered the input set, then we could answer the completeness questions posed in the previous section. We can determine if a convention is complete by looking at the resulting P-FSA. For example, will the convention work for any combination of argument types? The answer lies in the P-FSA transitions. For the convention to be complete, each state $q \in Q$ must have $\delta(q, c)$ defined for all $c \in C$. Should there exist some state $q \in Q$ and criteria $c \in C$ such that $\delta(q, c)$ is undefined, then, having arrived in state Q on input w , the machine would fail to accept any input string whose prefix was wc . Thus, there would be some signature whose placement could not be determined by the P-FSA. Since all correct P-FSAs must accept all strings in C^* , we can easily detect any P-FSA that implements an incomplete convention by looking for states with missing transitions.

3.2.2 Consistency

The consistency properties address whether the convention is internally and externally consistent. A convention is internally consistent if there is no machine resource that can be assigned to more than one argument. A convention is externally consistent if the caller and callee agree on the locations of transmitted values. In our model, we *detect* internal inconsistency and *prevent* external inconsistency.

To detect internal inconsistencies, we again turn to the P-FSA. If the convention only used finite resources, detecting a cycle in the P-FSA would be sufficient to detect the error. However, when infinite resources are introduced, so are cycles. We cannot have an internal inconsistency for an infinite resource since p is defined to be monotonically increasing. We detect finite resource inconsistencies in the following manner: An inconsistency can occur when there

is a transition from some state q_j to q_k , where bit i in the finite bit vector is 1 in q_j , but 0 in q_k . At this point, M has lost the information that resource r_i was already allocated. We can detect this change by comparing all pairs of bit vectors v_1, v_2 such that v_1 labels q_j , v_2 labels q_k , and $\delta(q_j, c) = q_k$ for some $c \in C$. To do the comparison, we compute

$$v_3 = (v_1 \oplus v_2) \wedge v_1.$$

$v_1 \oplus v_2$ selects all bits that differ between v_1 and v_2 . We logically AND this with v_1 to determine if any set bits change value. Thus, if v_3 has any bit set, we have an inconsistency.

Our convention specification language prevents external inconsistencies in the calling convention. A convention specification only defines the argument transmission locations once. Although both the caller and the callee must make use of this information, the specification does not duplicate the information. Since we only have a single definition of argument locations, we only construct a single P-FSA to model the placement mapping. This single P-FSA is used in both the caller and callee. Thus, we prevent external inconsistencies by requiring the caller and callee use the same implementation for the placement mapping.

4 CONSTRUCTION OF DIAGNOSTIC PROGRAMS

Using P-FSAs as an implementation foundation for a compiler enables all of the static analyses described in the previous section. However, when a compiler does not use a P-FSA in its implementation, we can still leverage off the P-FSA formalism to increase the implementation's robustness through systematic testing.

4.1 Test Vector Selection

To test a compiler's implementation of a calling convention, we must select a set of programs to compile and run. To exercise the calling convention, each test program must contain a caller and a callee procedure. For the purpose of testing the proper transmission of program values between procedures, the signature of the callee uniquely identifies a test case. Thus, two different programs whose callees' signatures match perform the same test. Therefore, the problem of generating test cases reduces to the problem of selecting signatures to test.

Selecting which procedure signatures to test is a difficult problem. Because the set of signatures, $S = \{(C^*, C^*)\}$, is infinite, one cannot test all signatures. However, since we can model the function that computes the placement of arguments as an FSA, there must be a finite number of states in an implementation to be tested. This is the case for any implementation, including those that do not explicitly use FSAs to model the placement function.

The problem of confirming that an implementation properly places procedure arguments is equivalent to experimentally determining if the implementation behaves as described by the P-FSA state table. This problem is known as the *checking experiment problem* from finite-automata theory [9], [10]. There are numerous approaches to this problem, most of which are based on transition testing. Transition testing forces the implementation to

TABLE 2
P-FSA Profiles for Several Calling Conventions

Machine	Alloc. Vector Bits	Mem. Part. Bits	$ Q $	$ \delta $	$ \Sigma $	Longest Acyclic Path
DEC VAX	0	0	1	3	3	0
M68020 (Sun)	0	2	4	24	6	3
SPARC (Sun)	6	3	9	90	10	8
M88100 (Motorola)	8	3	72	720	10	15
MIPS R3000 (DEC)	6	3	70	772	25	11

undergo all the transitions that are specified in the specification FSA.

An obvious first approach to generating test vectors using the P-FSA specification is to generate all vectors whose paths through the FSA are acyclic and those whose path ends in a cycle.⁷ This solution insures that each state q is visited and each transition $\delta(q, a)$ is traversed. However, the number of such paths for an FSA is $O(|\Sigma|^{|Q|})$. Table 2, which contains profiles for five P-FSAs generated automatically from CCL specifications, demonstrates that the acyclic method is not feasible for complex conventions such as the MIPS and SPARC. A simpler approach is to guarantee that each transition is exercised at least once. This limits the number of test vectors to no more than $|Q| \cdot |\Sigma|$. However, this method results in poor coverage that does not inspire confidence in the test suite. For example, for the P-FSA in Fig. 2, the three signatures:

```
void f(double, double);
void f(int, int, int, int);
void f(int, double);
```

cover all int and double transitions leaving states q_{0-2} . This leaves the signature:

```
void f(double, int);
```

untested. Clearly, such a test should be included in the suite. To further illustrate the problem, consider the FSA specification shown in Fig. 3a. An erroneous implementation, shown in Fig. 3b, contains an extra state q'_1 that is reached on initial input "b." The two strings, "aaa" and "bbb" completely cover the specification FSA transitions. Unfortunately, these test vectors will not detect that the implementation has an additional (fault) state. Thus, it is not sufficient to include only test vectors that cover the transition set.

An alternative, which falls between the simple transition approach and the acyclic path approach, we call the *transition-pairing* approach. In transition pairing, we examine each state in the specification FSA. For each state, we include a test vector that covers each *pair* of entering and exiting transitions. This eliminates the faulty state detection problem illustrated in Fig. 3. To illustrate how, consider the test vectors this process generates: While examining state q_1 , transition-pairing will add the substrings "aa," "ab," "ba," and "bb" to the set of substrings used to generate test

7. We define a *path that ends in a cycle* to be a cyclic path w where the path w is acyclic.




Fig. 3. Example FSA where a fault will not be detected. (a) Specification FSA. (b) Implementation FSA.

vectors. Since the context in which these substrings are to be used is q_0 , they contribute prefixes to the test vector set. Upon exercising q_1 using the prefix "ba," the implementation FSA will generate incorrect output: 10 instead of 11. This difference can be identified and the faulty state detected.

In addition to such fault detection, transition-pairing provides tests that have a similar characteristic to the acyclic method: Transitions are tested in all the contexts that can be applied. Although there are many combinations that are not tested, they are similar to ones included in the set. For example, in the simple FSA pictured in Fig. 2, we could have a set of test vectors that includes the vector double double to exercise the state q_4 with the transition pair $((q_2, \text{double}), (q_4, \text{double}))$. Such a set would not need to include int int double double to cover the same transition pair.

This method of test vector generation provides a complete coverage of transitions in the specification FSA. Further, the tests reflect the context sensitivity that transitions have. This allows for some erroneous state and transition detection, while significantly reducing the number of test vectors. The test vector sizes are significantly smaller than the acyclic method, while still providing a significant degree of confidence (Table 3).

An algorithm for generating transition-pair paths is shown in Fig. 4. The algorithm performs a depth-first search of the FSA state graph. Each time a transition (q, a) is encountered, it is marked. This mark indicates that all paths that go beyond (q, a) have been visited. When the algorithm

TABLE 3
Sizes of Test Suites for Various Selection Methods

Machine	Transition Paths	Transition-Pair Paths	Acyclic Paths
DEC VAX	3	12	3
M68020 (Sun)	24	324	96
SPARC (Sun)	224	7,434	$> 10^8$
M88100 (Motorola)	720	22,412	$> 10^8$
MIPS R3000 (DEC)	772	5,655	8×10^8

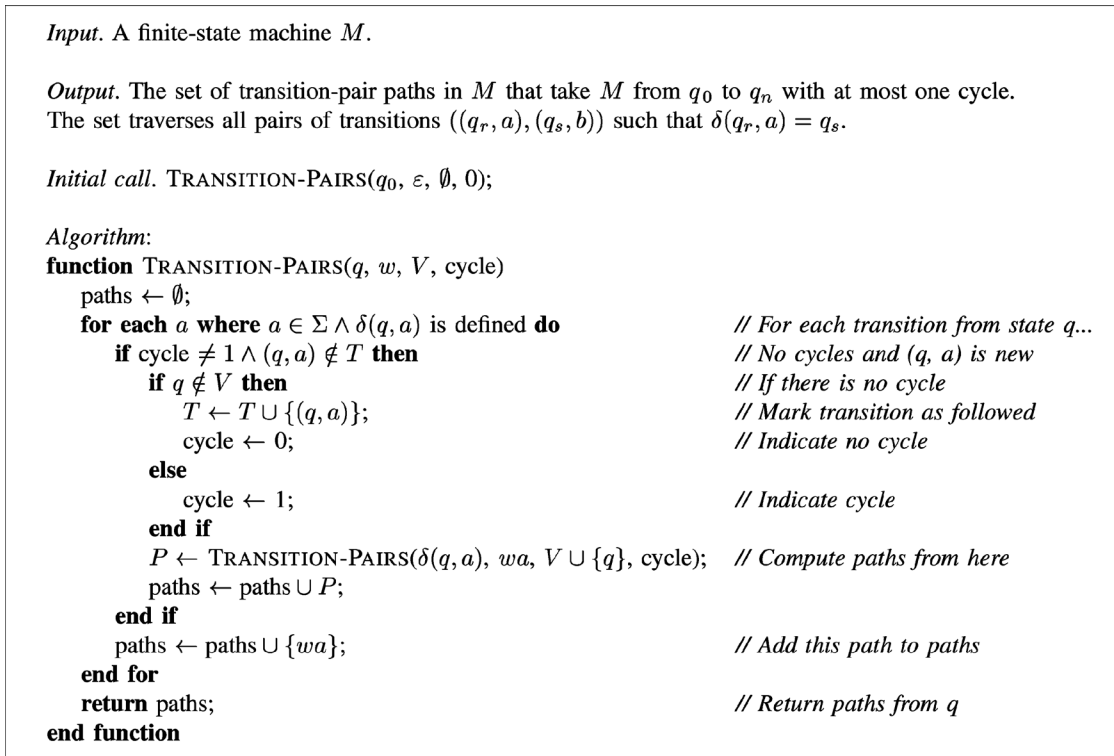


Fig. 4. Test vector generation algorithm.

reaches a state q_n on transition (q_m, a) , each transition (q_n, b) where $b \in \Sigma$ is visited whether or not it is marked. This causes all pairs of transitions $((q_m, a), (q_n, b))$ to be included. These pairs represent all combinations of one entering transition with all exiting transitions. Because the algorithm is depth-first, each entering transition is guaranteed to be visited. Thus, all combinations of entering and exiting transitions are included.

Work related to the automatic generation of test suites has received much attention recently in the area of conformance testing of network protocols [11]. The purpose of these suites is to determine if the implementation of a communication protocol adheres to the protocol's specification. Often, the protocol specification is provided as a finite-state machine. This has resulted in many methods of test selection, including the Transition tour, Partial W-method [12], Distinguishing Sequence Method [10], and Unique-Input-Output method [13]. These methods are derivatives of the checking experiment problem where an implementation is checked against a specification FSM [14]. Such techniques have also been used in the automatic verification of digital circuits [9], [15].

What distinguishes these methods from ours are the underlying assumptions concerning the characteristics of the implementation FSAs. Unlike theirs, our FSAs can have a large number of states and transitions. This significantly changes the nature of the solution to the problem. Furthermore, much of the problem that network conformance researchers are faced with is identifying which state the implementation FSA is in. A significant portion of their work focuses on generating test vectors that discover the state of the machine. Fortunately, we can always put our implementation FSA in the start state. Also, in their work, a

bound on the number of states in the implementation FSAs is assumed. Because we have no practical bound on the number of states in the implementation, their work is not applicable.

4.2 Test Case Generation

After selecting the appropriate test vectors or procedure signatures, the corresponding test cases must be realized. In our approach, we generate a separate test program for each test vector so that we can easily match any reported errors to the specific test vector.

A procedure call is composed of two pieces: the procedure call within the caller (the call site) and the body of the callee. Because they are implemented differently, these two pieces of code are typically generated in separate locations in a compiler. This natural separation is reflected in the way that we construct our test cases. Each test case is comprised of two files, one contains the caller, the other contains the callee. The two files are compiled and linked together. The programs are self-checking so that, if a procedure call fails, this event is reported by the test itself.

Fig. 5 shows the compiler conformance test process. One file is compiled by the compiler-under-test (CUT), while the other is compiled by the reference compiler. The reference compiler operationally defines the procedure-calling convention (its implementation is defined to be correct). The resulting object files are linked together and run. Results of the test are checked by the conformance verifier and given to the test conductor. The test conductor tallies the results of all tests for a test suite and generates a conformance report. Although this process uses two compilers, the same process may still be used if a reference compiler is not available. However, this will weaken the conformance verifier's

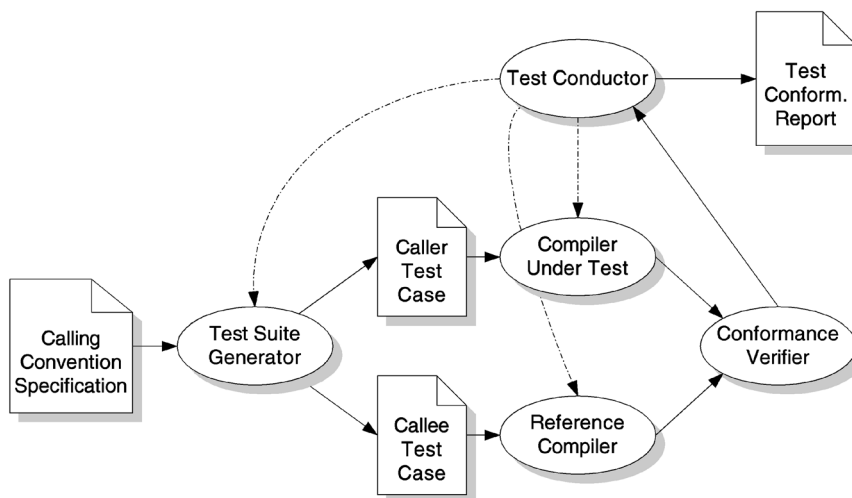


Fig. 5. The compiler conformance test process.

ability to automatically diagnose errors as discussed in the next section.

Figs. 6 and 7 show an example test case for the C signature `void (int, double, struct(2))8`. The caller loads each argument with randomly selected bytes. However, the values of these bytes have an important property: Each contiguous set of two bytes is unique. Thus, for a string B of m bytes, for all indexes $0 < i \leq m$, there exists no index $0 < j \leq m$ and $j \neq i$ such that $B[j+k] = B[i+k]$ for all $0 \leq k < 2$. We can easily guarantee this property for all strings B whose length is no greater than $65,536$ (2^{16}) bytes. Since the likelihood of using an argument list of size greater than 64 Kbytes is small, this is sufficient to guarantee that any two bytes passed between procedures are unique. This makes it easier to identify if an argument has been shifted or misplaced. The callee receives the values and checks them against the expected values. If the values do not match, an error condition is signaled.

As one might expect, the generation of good test cases from selected signatures is language dependent. One convention used in the C programming language is *varargs*. *varargs* is a standard for writing procedures that accept variable length argument lists. The proper implementation of *varargs* in a C compiler is difficult. For each test case that we generate, we also generate a *varargs* version to verify that this standard convention is implemented correctly.

4.3 Automatic Diagnosis of Errors

Generation of good tests is only a part of the testing process. If a test fails, the problem must be diagnosed and a solution developed. In this section, we discuss how the second step, diagnosis, can be partially automated.

As discussed earlier, the conformance verifier links a caller and callee together and runs the resulting program. When both a reference compiler and CUT are used, this results in four distinct caller-callee pairs whose results we call an *outcome*. We can glean more information than a single test can supply by considering the composite result that an outcome provides. When a test fails, the results of two other tests can help isolate the fault. For example, in the

outcome shown in Fig. 8, the CUT/reference test has failed. Since the CUT/CUT test failed, but the reference/CUT test passed, this indicates the fault is in the CUT caller.

This method of isolating errors by swapping different components makes it possible to automatically diagnose common errors. Since there are only 16 outcome configurations, each outcome can be hand-analyzed once and the results tabulated, as shown in Table 4. Several diagnoses deserve mention. First, although the reference compiler is considered the authority, there are six cases where the reference can be determined to be faulty. Second, the four outcome configurations where only a single test fails are not possible. This cannot occur with a single test failure since we assume each component uses a single convention.⁹ Finally, for two of the cases, we not only can isolate the location of the fault, but we can identify the nature of the error. This occurs in outcomes where two conflicting conventions have been discovered.

The combination of test vector selection and automatic diagnosis proves to be a powerful debugging tool. As tests are generated, run, and analyzed, patterns of errors tend to emerge. We have found that the patterns themselves suggest the nature of the problem. For example, finding that an error occurred for every signature that included a struct of size greater than seven bytes might suggest an alignment problem. More sophisticated patterns can exist and with knowledge of the calling convention can significantly help the developer correct faults.

4.4 Test Results

We used our technique for selecting test vectors to test several compilers on several target machines. Several errors

9. Appel observes that such outcomes actually are possible [16]. In his counterexample, the CUT caller implements a different convention than the reference compiler, but the CUT callee implements *both* conventions. In this scenario, the fault is detected in the CUT/reference test, but not in either the CUT/CUT or the reference/reference tests. Although such a case is possible, the probability of a callee implementing two different conventions that do not conflict, i.e., using the same register for two different purposes, is small. The benefits, in terms of diagnostic ability, of considering such a case as invalid far outweigh any accuracy gained by labeling it a valid outcome. Finally, if such a case were to occur, it would still be detected; it just could not be automatically diagnosed.

8. We denote a structure whose size is n bytes as `struct(n)`.


```

typedef union {
    unsigned char bytes[8];
    double dbl;
} dblcvt;

typedef struct struct_2 {
    unsigned char field[2];
} struct_2;

void test_function_callee();

void test_function_caller()
{
    dblcvt cvt;
    static struct_2 struct_2_arg_3 = {
        { 0x34,0x8f,}
    };
    double double_arg_2;

    cvt.bytes[0] = 0xe2;
    cvt.bytes[1] = 0xed;
    cvt.bytes[2] = 0xab;
    cvt.bytes[3] = 0xad;
    cvt.bytes[4] = 0x67;
    cvt.bytes[5] = 0x31;
    cvt.bytes[6] = 0xee;
    cvt.bytes[7] = 0x7;
    double_arg_2 = cvt.dbl;

    test_function_callee(0x440260971,
                          double_arg_2,
                          struct_2_arg_3);
}

```

Fig. 6. Code generated for caller.

were found in C compilers on the MIPS. In this section, we present these results.

We selected several C compilers that generate code for the MIPS architecture (a DECStation Model 5000/125). These included the native compiler supplied by DEC, two versions of Fraser and Hanson's *lcc* compiler [17], [18], several versions of GNU's *gcc* [19], and a previous version of our own C compiler, *vpcc/vpo*, that used a hand-coded calling sequence generator [20]. Although we feel that this technique is extremely valuable throughout the compiler development cycle, we believe that it would be fairest to evaluate its effectiveness in finding errors in young implementations of compilers. Where possible, we have used early versions of these compilers. These versions, called *legacy* compilers, represent younger implementations that more accurately exhibit bugs found in initial releases of compilers. However, each of these compilers is a production-quality compiler that has been widely used for years. Finding any bugs in their implementations is still a significant challenge.

In testing the compilers, we checked for two types of conformance: internal and external. Compiler *A* internally conforms if code that it generates for a caller can properly call code for a callee that it generated. We denote this using

```

typedef union {
    unsigned char bytes[8];
    double dbl;
} dblcvt;

typedef struct struct_2 {
    unsigned char field[2];
} struct_2;

void test_function_callee(long_arg_1,
                          double_arg_2, struct_2_arg_3)
{
    long long_arg_1;
    double double_arg_2;
    struct_2 struct_2_arg_3;

    dblcvt cvt;

    if(long_arg_1 != 0x440260971) {
        fprintf(stderr, "Bad long_arg_1\n");
        exit(1);
    }
    cvt.bytes[0] = 0xe2;
    cvt.bytes[1] = 0xed;
    cvt.bytes[2] = 0xab;
    cvt.bytes[3] = 0xad;
    cvt.bytes[4] = 0x67;
    cvt.bytes[5] = 0x31;
    cvt.bytes[6] = 0xee;
    cvt.bytes[7] = 0x7;
    if(double_arg_2 != cvt.dbl) {
        fprintf(stderr, "Bad double_arg_2\n");
        exit(1);
    }
    if(struct_2_arg_3.field[0] != 0x34) {
        fprintf(stderr, "Element 0 is bad in "
                "struct_2_arg_3.field\n");
        exit(1);
    }
    if(struct_2_arg_3.field[1] != 0x8f) {
        fprintf(stderr, "Element 1 is bad in "
                "struct_2_arg_3.field\n");
        exit(1);
    }
}

```

Fig. 7. Code generated for callee.

$A \mapsto A$. Compiler *A* externally conforms if its caller code can call another compiler *B*'s callee code and *vice versa* ($A \mapsto B$ and $B \mapsto A$). Thus, the callees and callers are compiled using each of the compilers under test. This results in n object versions for n compilers. Each caller version is then linked with the callee that was generated by the same compiler. This results in the n tests necessary to verify internal conformance for this test case. To establish external conformance, we could naively link each caller to each callee, which would yield $2n^2$ tests. However, we can do better. Recognizing that procedure call (\mapsto) is symmetric, we can easily reduce this to n^2 (since if $A \mapsto B$, then $B \mapsto A$). Furthermore, procedure call is also transitive, so if $A \mapsto B$ and $B \mapsto C$, then $A \mapsto C$. This reduces the number to $2n - n$ as

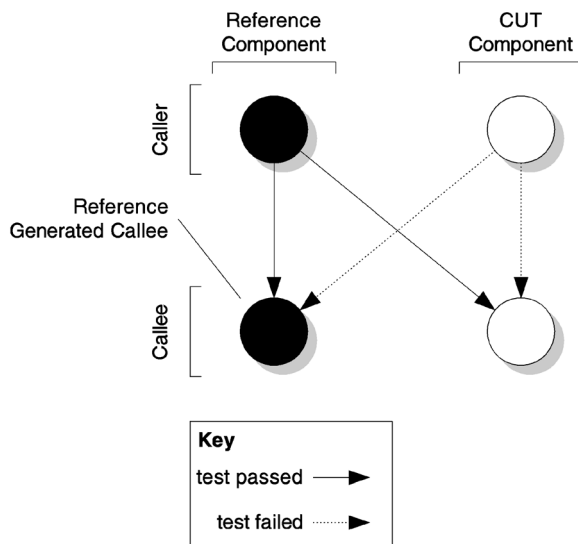


Fig. 8. An example outcome.

pictured in Fig. 9. Each compiler's caller is linked to the reference compiler's callee. This facilitates the isolation of which compiler does not conform when an error is detected.

The results of running both internal and external tests on the compiler set for the MIPS are shown in Table 5. We found both internal and external conformance errors in all of the tested compilers. Table 5 reports internal and external errors separately. Within each class, the number of actual tests that failed and the number of faults that caused failure are indicated.¹⁰ The numbers reported in the fault columns indicate the approximate number of actual coding errors resulting in test failures. These numbers are only approximate. We tried, as best we could, to glean this information from the results of tests. More accurate numbers can only be obtained by examining the compiler's source.

4.4.1 Standard Procedure Calls

Internal conformance errors were found in two versions of *gcc*. *gcc* 1.38 failed 24 tests that focus on passing structures in registers. Structures between nine and 12 bytes in size (three words) are not properly passed starting in the second argument register. *gcc* 2.4.5 fails a single test. The fault occurs with procedures with the signature:

```
void (struct(1), struct(1), struct(1));
```

gcc 2.4.5 fails to even compile a procedure with this signature.¹¹ The fact that *gcc* 2.1 does not have this error indicates that the error was *introduced* after version 2.1. This supports our conjecture that such a method of automatic testing is extremely useful throughout the development and maintenance life-cycle of a compiler.

External conformance errors were more prevalent. *gcc* 1.38 does not properly pass 1-byte structures in registers. *gcc* 1.38 and 2.4.5 cannot pass a structure in the third argument register when that structure is followed by another structure. *vpcc/vpo* has two faults: 1) Structures are not passed properly in registers and 2) 1 to 4-byte structures

are not passed in memory correctly if they are immediately followed by another structure.

4.4.2 Variadic Procedure Calls

Procedures that take variable-length argument lists (variadic functions) are written using one of the two standard header files: `stdarg.h` (for traditional C) and `stdarg.h` (for ANSI C). The following paragraphs detail the results of calling callees that are implemented using `stdarg/stdarg`.

When running test cases that contained variadic functions whose first argument was a `double`, we found that none of the compilers, including the reference compiler, properly implemented the calling convention. Version 2 releases of *gcc* managed to avoid this problem at the expense of interoperability; their generated callees do not conform to the established calling convention.

We also tested several compilers targeted to the SPARC architecture. On the SPARC, the test suite generator produces 12,034 tests. Using our automated testing infrastructure, we tested three mature compilers and one research compiler. The mature compilers were *cc*—the native C compiler supplied by Sun Microsystems (Sun WorkShop Compiler C SPARC Version 5.000), *gcc* (version 2.95.2), and *lcc/vpo*—a compiler built using the Zephyr compiler infrastructure (*lcc* version 4.1 and *vpo* version 2.0). The research compiler was built using the Scale compiler infrastructure [21]. We tested a compiler included with version 1.7, which is the third release of the compiler. We shall refer to this compiler as *scale*. The results of running both internal and external tests on the compiler set for the SPARC are given in Table 6.

The *scale* compiler failed a large number of tests. Inspection of the test report showed that the internal failures were because *scale* did not handle variadic functions. There appeared to be two different faults: 1) *scale* threw an exception when trying to produce code for the callee and 2) the compiler was able to generate both a caller and a callee, but there appeared to be a mismatch between the convention implemented by the caller and the callee. For the additional 442 failed external compliance tests, *scale* generated caller code that was incompatible with *cc* callee code. In this case, *scale* generated incorrect caller code for functions with signatures of the following form:

```
void (int, int, int, int, int, int,
      struct(1, 2, 3, 5, 6, 7));
```

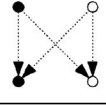
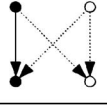
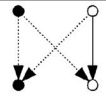
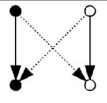
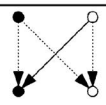
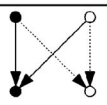
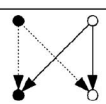
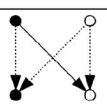
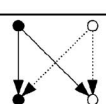
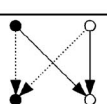
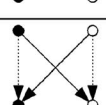
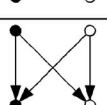
The test cases represented by these signatures test a boundary condition—the first six arguments are passed in registers (`%o0-%o5`) and the structure is placed on the stack. Notice that both 4-byte and 8-byte structures are successfully passed. This suggests the type of fault that has occurred: Fetching the last bytes of structures that are not multiples of four is implemented incorrectly. The test failures also demonstrate a type of fault that is not likely to be covered by a hand-generated test suite.

From these results, obviously the state-of-the-art in compiler testing is inadequate. All of the compilers we tested had undergone rigorous testing. However, hand development of test suites is an arduous and, itself, error-prone task. Furthermore, because these tests are target specific, they must be revisited with each retargeting of the compiler. In contrast, by using automatic test generators that are target sensitive, compilers can quickly be validated

10. These numbers include tests of both standard procedure calls and variadic procedure calls.

11. The error returned by *gcc* 2.4.5 was: `gcc: Internal compiler error: program ccl got fatal signal 4.`

TABLE 4
All Outcome Configurations

Outcome	Diagnosis	Outcome	Diagnosis
	Faults in at least two components.		Faults in both components of the CUT.
	Faults in both components of reference compiler.		CUT implements wrong convention (does not externally conform with the reference).
	Fault in the reference compiler's caller. Fault in the CUT's callee.		Fault in the CUT's callee.
	Fault in the reference compiler's caller.		Fault in the reference compiler's callee. Fault in CUT's caller.
	Fault in the CUT's caller.		Fault in reference compiler's callee.
	Two conventions. One shared between the reference compiler's callee and CUT's caller, and vice versa.		No faults detected.

Key

test passed →

test failed →

Reference

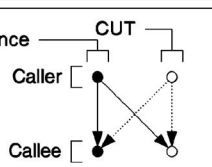
Caller [●]

Callee [●]

CUT

Caller [○]

Callee [○]






Fig. 9. Determining conformance of n compilers.

TABLE 5
Results of Running the MIPS Test Suite on Several Compilers

Compiler	Internal		External	
	Failed Tests	Faults	Failed Tests	Faults
cc (native)	2,346	1	2,346	1
gcc (1.38)	2,370	2	2,567	3
gcc (2.1)	0	0	2,346	1
gcc (2.4.5)	1	1	2,374	3
lcc (1.9)	0	0	0	0
lcc (3.3)	2,407	2	2,407	2
vpcc/vpo	2,346	1	486	3
Total	9,470	7	12,526	13

TABLE 6
Results of Running the SPARC Test Suite on Several Compilers

Compiler	Internal		External	
	Failed Tests	Faults	Failed Tests	Faults
cc (native)	0	0	0	0
gcc (2.95)	0	0	0	0
lcc/vpo	0	0	0	0
scale	5,925	2	6,367	3
Total	5,925	2	6,367	3

before each release. Although the ratio of failed tests to faults (1,000:1) may seem high, these tests are generated, run, and analyzed automatically. It is not necessary to examine each of the failed tests. Instead, a single failed test can be examined and corrected. The suite can then be run again to determine if faults still exist (or if the fix has introduced new faults).

5 CONCLUSIONS

Building compilers that generate correct code continues to be a difficult problem. Current implementations of calling sequence generators often contain errors. This comes from the lack of a formal model and implementation mechanism that can guarantee completeness and consistency properties. We have presented such a formal model, called P-FSAs, for procedure-calling conventions that can ensure these properties. A P-FSA that models a convention can be automatically constructed from the convention's specification. During construction, the convention can be analyzed to determine if it is complete and consistent. The resulting P-FSA can then be directly used as an implementation of the convention in an application.

Although it is possible to automatically generate the calling sequence generator using P-FSAs, some work is required to retrofit an existing compilation system to use them. Fortunately, it is possible to reap the benefits of P-FSA without any modification of the compiler. Using automated compiler tools and testing, one can significantly increase the robustness of *any* compiler. We have combined these two techniques in a new way that further closes the gap between actual compiler implementations and the ever-sought-after correct compiler. By using a formal model of procedure-calling conventions, we have designed and implemented a technique that automatically identifies boundary test cases for calling sequence generators and diagnoses the nature of the fault. We then applied this technique to measure the conformance of a number of production-quality compilers for the MIPS and SPARC. This system identified a total of at least 23 faults in the tested compilers. These errors were significant enough to cause over 6,000 different test cases to fail. Clearly, this technique is effective at exposing and isolating faults in calling sequence generators of mature compilers. Undoubtedly, it would be even more effective during the initial development of a compilation system.

REFERENCES

- [1] M.W. Bailey and J.W. Davidson, "A Formal Model and Specification Language for Procedure Calling Conventions," *Proc. ACM SIGPLAN-SIGACT Symp. Principles of Programming Languages*, pp. 298-310, Jan. 1995.
- [2] S.C. Johnson and D.M. Ritchie, "The C Language Calling Sequence," Bell Labs, Year?
- [3] B.W. Kernighan and D.M. Ritchie, *The C Programming Language*, second ed. Prentice Hall, 1988.
- [4] W.L. Johnson, J.H. Porter, S.I. Ackley, and D.T. Ross, "Automatic Generation of Efficient Lexical Processors Using Finite State Techniques," *Comm. ACM*, vol. 11, no. 12, pp. 805-813, 1968.
- [5] T.A. Proebsting and C.W. Fraser, "Detecting Pipeline Structural Hazards Quickly," *Proc. ACM SIGPLAN-SIGACT Symp. Principles of Programming Languages*, pp. 280-286, 1994.
- [6] T. Müller, "Employing Finite Automata for Resource Scheduling," *Proc. 26th Ann. Int'l Symp. Microarchitecture*, pp. 12-20, 1993.
- [7] G.H. Mealy, "A Method for Synthesizing Sequential Circuits," *Bell System Technical J.*, vol. 35, no. 5, pp. 1045-1079, 1955.
- [8] J.E. Hopcroft and J.D. Ullman, *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 1979.
- [9] F.C. Hennie, "Fault Detecting Experiments for Sequential Circuits," *Proc. Fifth Ann. Symp. Switching Theory and Logical Design*, pp. 95-110, Nov. 1964.
- [10] Z. Kohavi, *Switching and Finite Automata Theory*, second ed. McGraw-Hill, 1978.
- [11] D.P. Sidhu and T.-K. Leung, "Formal Methods for Protocol Testing: A Detailed Study," *IEEE Trans. Software Eng.*, vol. 15, no. 4, pp. 413-426, Apr. 1989.
- [12] S. Fujiwara and G.v. Bochmann, F. Khendek, M. Amalou, and A. Ghedamsi, "Test Selection Based on Finite State Models," *IEEE Trans. Software Eng.*, vol. 17, no. 6, pp. 591-603, June 1991.
- [13] A.V. Aho, A.T. Dahbura, D. Lee, and M.U. Uyar, "An Optimization Technique for Protocol Conformance Test Generation Based on UIO Sequences and Rural Chinese Postman Tours," *IEEE Trans. Comm.*, vol. 39, pp. 1604-1615, Nov. 1991.
- [14] M. Yannakakis and D. Lee, "Testing Finite State Machines: Fault Detection," *J. Computer and System Sciences*, vol. 50, pp. 209-227, 1995.
- [15] R.C. Ho, C.H. Yang, M.A. Horowitz, and D.L. Dill, "Architecture Validation for Processors," *Proc. Int'l Symp. Computer Architecture*, pp. 404-413, 1995.
- [16] A.W. Appel, personal communication, May 1996.
- [17] C.W. Fraser and D.R. Hanson, "A Code Generation Interface for ANSI C," *Software—Practice and Experience*, vol. 21, no. 9, pp. 963-988, 1991.
- [18] C. Fraser and D. Hanson, *A Retargetable C Compiler: Design and Implementation*. Benjamin Cummings, 1995.
- [19] R.M. Stallman, *Using and Porting GNU CC (Version 2.0)*. Free Software Foundation, Inc., Feb. 1992.
- [20] M.E. Benitez and J.W. Davidson, "A Portable Global Optimizer and Linker," *Proc. SIGPLAN Conf. Programming Language Design and Implementation*, pp. 329-338, July 1988.
- [21] G.E. Weaver, B.D. Cahoon, J.E.B. Moss, K.S. McKinley, E.J. Wright, and J.H. Burrill, "The Common Language Encoding Form (CLEF) Design Document," Technical Report 97-58, Dept. of Computer Science, Univ. of Massachusetts, Amherst, Aug. 1997.



Mark W. Bailey received the BS degree in computer and information science from the University of Massachusetts in 1988. He received the MCS and PhD degrees in computer science from the University of Virginia in 1990 and 2000, respectively. He has been a member of the faculty of Hamilton College since 1997, where he is an assistant professor of computer science. His research interests include compilers, optimization, embedded systems, computer architecture, and computer security. He is author of numerous workshop, conference, and journal articles. He has been a member of organizing committees of international conferences in his field. He is a member of the IEEE, IEEE Computer Society, ACM SIGPLAN, and SIGCSE.



Jack W. Davidson received the BAS and MS degrees in computer science from Southern Methodist University in 1975 and 1977, respectively. He received the PhD degree in computer science from the University of Arizona in 1981. He has been a member of the faculty of the University of Virginia since 1982, where he is a professor of computer science. His main research interests include compilers, code generation, optimization, embedded systems, and computer architecture. He is author of more than 100 research articles in refereed conferences and journals articles as well as coauthor of two widely used introduction to programming textbooks. He has been a program/general chair or member of steering/program/organizing committee of many international conferences in his field and he was an associate editor of the *ACM Transactions on Programming Language and Systems* from 1994 to 2000. He is a member of the IEEE Computer Society, ACM SIGPLAN, SIGARCH, and SIGCSE.

► For more information on this or any computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.