

1 Instruction Set Architecture Specification, 2 Verification, and Validation using Algorithmic C 3 and ACL2

4 David S. Hardin
5 Collins Aerospace, USA
6 david.hardin@collins.com

7 — Abstract —

8 It is common practice during Instruction Set Architecture (ISA) development to create an ISA
9 simulator, usually in C/C++. We describe an experiment in which we implement such an ISA
10 simulator for a derivative of a popular ISA written in a subset of Algorithmic C, to allow for the
11 verification of binary programs targeting that ISA, as well as to aid in the validation of the ISA
12 model via simulated execution of test programs on the model. Algorithmic C defines C++ header
13 files that enable compilation to both hardware and software platforms, providing support for the
14 peculiar bit widths employed, for example, in floating-point hardware design. We utilize a toolchain,
15 due to Russinoff and O’Leary, that provides a translation from a restricted subset of Algorithmic C
16 to the Common Lisp subset supported by the ACL2 theorem prover. This toolchain, called RAC, is
17 documented in Russinoff’s recent book on floating-point hardware verification. We create an ISA
18 simulator in this C++ subset, use RAC to translate this simulator code to ACL2, produce small
19 binary programs for the ISA that we use to validate the simulator, and utilize the ACL2 Codewalker
20 decompilation-into-logic facility to prove those programs correct.

21 **2012 ACM Subject Classification** Hardware → Hardware validation

22 **Keywords and phrases** Formal verification, theorem proving, instruction set architecture, decompilation into logic

23 **Digital Object Identifier** 10.4230/LIPIcs...


24 **1** Introduction

25 in his *tour de force* book on floating-point hardware verification [4], David Russinoff details
26 a 20-year quest to provide mathematical proofs for complex arithmetic hardware utilizing an
27 automated theorem prover, namely ACL2. Russinoff begins by presenting a formalization of
28 modular arithmetic in standard mathematical notation, but backed by collections of ACL2
29 lemmas, called *books* in ACL2 parlance. The *RTL* books are named after the “Register
30 Transfer Logic” artifacts that hardware designers produce using Hardware Description
31 Languages (HDLs), and are designed for reasoning at the RTL level. After developing
32 progressively more complex arithmetic circuits, Russinoff concludes his text with the complete
33 verification of representative ARM floating-point RTL for addition, multiplication, fused
34 multiply/add, division and square root. All of the sources and tools described in Russinoff’s
35 book are available as part of the ACL2 distribution, so the curious reader can reproduce his
36 verification results.

37
38 The HDL that Russinoff employs in his book is a subset of Algorithmic C. Algorithmic C¹
39 entails a set of freely available C++ header files providing support for hardware development,
40 including the peculiar bit widths utilized in floating-point design, and enables compilation to
41 both hardware and software platforms. The Restricted Algorithmic C toolchain, or *RAC*,
42 evolved from a similar toolchain for SystemC called MASC [3]. RAC accepts Algorithmic C

¹ Available at <https://www.mentor.com/hls-lp/downloads/ac-datatypes>





■ **Figure 1** Restricted Algorithmic C (RAC) toolchain and Codewalker.

```

struct leg64St {
    ui10 pc;
    array<ui64, REG_SZ>regs;
    array<ui32, CMEM_SZ>cmem;
    array<ui64, DMEM_SZ>dmem;
    ui8 opcode;
    ui8 op1; [...] }

leg64St leg64step(leg64St s) {
    return do_Inst(nextInst(s)); }

leg64St leg64steps(leg64St s, uint n) {
    for (uint i=n; i>0; i--) {
        s = leg64step(s); }
    return s; }
  
```

■ **Figure 2** Fragment of the LEG64 ISA state and step functions in Restricted Algorithmic C.

43 source, then translates it to S-expression forms acceptable to ACL2. A simplified view of the
 44 RAC toolchain is shown in Fig. 1; for a complete depiction, see Fig. V.1 of [4].

45 RAC imposes several restrictions on the Algorithmic C developer. The most significant
 46 of these is that all RAC function arguments must be pass-by-value, and all functions must
 47 be side-effect-free. Please refer to Chapter 15 of [4] for details.

48 2 ISA Specification in Restricted Algorithmic C

49 Our present challenge is to determine whether the RAC toolchain can be used to develop
 50 an Instruction Set Architecture (ISA) simulator, of the sort commonly crafted (usually in
 51 C/C++) during ISA design. The simulator should support proof-based verification, as well
 52 as validation by simulation. For this experiment, we have created a simple 64-bit ISA, quite
 53 similar to a commercially popular architecture, called LEG64. Figure 2 presents a fragment
 54 of the LEG64 ISA state vector (left side), and instruction-stepping functions (right side).²

55 RAC translates the state struct, including its arrays, to ACL2 records, with `ag` get
 56 function, and `as` set function. The fixed-width data types are enforced in ACL2 using the
 57 RTL (`bits y i j`) function, which returns a bit-slice of `y` between bit indices `i` and `j`. Loop
 58 bodies are translated into ACL2 recursive functions, with generated ACL2 `:measure` forms
 59 to aid in termination analysis. The resulting ACL2 functions are readable, if not pretty.

60 3 Validation by Simulation

61 To validate our RAC specification for the ISA, we can use the C++ executable for the ISA
 62 simulator, a hardware simulator, or the RAC-translated ACL2 simulator, as depicted in

² Complete sources available at <https://github.com/david-s-hardin/algoc-isa>

```

ui64 fact(ui64 n, ui64, acc) {
    if (n == 0) {
        return acc;
    } else {
        return fact(n-1, acc * n);
    }
}

.L3:  cmp r0, #0           ; r0 == 0?
      beq .L2             ; if so, done
      mul r1, r1, r0      ; r1 <- r1*r0
      sub r0, r0, #1      ; r0 <- r0-1
      b .L3               ; goto top
.L2:  mov r0, r1          ; r0 <- r1
      ret                 ; return

```

■ **Figure 3** Example test program for the LEG64 ISA simulator.

63 Fig. 1. In all cases, we validate the ISA simulator by loading the simulated code memory
64 with binary programs compiled for our target ISA, load the data memory and registers
65 appropriately for a given test, then step the simulator for the required number of steps, then
66 inspecting the relevant elements of the resulting state vector. The results of running tests
67 on the C++ executable of the ISA simulator can be checked against the translated ACL2
68 model as a validation check for the translation; ACL2’s speed aids in this effort.

69 **4 Verification by ACL2 Proof using Codewalker**

70 Once the LEG64 simulator has been translated to ACL2, we can reason about the translated
71 functions, using the RTL books, as well as other ACL2 books. We can also reason about
72 LEG64 binary code, using J Moore’s decompilation-into-logic tool, Codewalker³. We have
73 previously used Codewalker to prove properties of LLVM code [1].

74 As an example, consider a simple tail-recursive unsigned 64-bit factorial function (left of
75 Figure 3), compiled to yield the LEG64 code on the right of Figure 3. The input parameter,
76 `n`, is passed in using register `r0`, and the tail-recursive accumulator, `acc`, is passed in `r1`. The
77 function actually computes `acc * n!`, which is equal to `n!` when `acc = 1` (modulo 64 bits).
78 This result is then copied into register `r0` for return to the caller.

79 To begin the verification, we describe the LEG64 machine, its registers, memories, and
80 ISA simulator functions, to Codewalker, and provide it the binary of the code of Figure 3.
81 Codewalker then “walks” the binary code, guided by the ACL2 ISA simulator, and decompiles
82 each basic block into a *semantic function* summarizing the effect of that block on the ISA
83 state `s`, in the style of Magnus Myreen’s Ph.D. work [2]. Codewalker also generates a
84 per-block clock function, giving the number of instruction steps needed for a given input
85 condition. We can then instruct Codewalker to “project” the final value for a given register
86 out of a loop, producing a function that abstracts away the LEG64 state details. That
87 function can then be reasoned about, and/or executed in ACL2. For our factorial example,
88 Codewalker produces `(fn1-loop pc n acc)` for the final accumulator value, a function of
89 the ISA simulator program counter `pc`, input value `n`, and initial accumulator value `acc`. After
90 creating a “wrapper” function, `fn1`, we prove the correctness theorems of Figure 4, where `!`
91 is a non-tail-recursive infinite-precision factorial function, and `expt` is exponentiation.

92 Finally, we prove that the execution of the loop of the factorial subroutine, loaded into
93 code memory of the simulated LEG64, using the stepping function `leg64steps`, written in
94 Algorithmic C and then translated by RAC to ACL2, for the number of steps given by the

³ Codewalker is part of the ACL2 standard distribution, in `books/projects/codewalker`

```

(defthm fn1-loop-is-acc-*-n!
  (implies
    (and (natp n)(natp acc)
         (< n (expt 2 64))
         (< acc (expt 2 64)))
    (equal (fn1-loop 0 n acc)
           (bits (* acc (! n)) 63 0))))

(defun fn1 (n) (fn1-loop 0 n 1))

(defthm fn1-is-n!
  (implies
    (and (natp n) (< n (expt 2 64)))
    (equal (fn1 n)
           (bits (! n) 63 0))))

```

■ **Figure 4** Correctness theorems for the projected loop result function from Codewalker.

```

(defthm reg-1-of-fact-loop-is-acc-*-n!
  (implies
    (and (fact-routine-loadedp s) (integerp (ag 0 (ag 'regs s)))
         (<= 1 (ag 0 (ag 'regs s))) (< (ag 0 (ag 'regs s)) (expt 2 64))
         (integerp (ag 1 (ag 'regs s))) (< (ag 1 (ag 'regs s)) (expt 2 64))
         (= (ag 'pc s) 0))
    (equal (ag 1 (ag 'regs (leg64steps s (acl2::clk-0 s))))
           (bits (* (ag 1 (ag 'regs s)) (! (ag 0 (ag 'regs s)))) 63 0))))

```

■ **Figure 5** Correctness of the factorial loop when executed on the LEG64 ISA simulator.

95 Codewalker-generated clock function `clk-0`, yields the expected result, as shown in Figure 5.

96 5 Conclusion and Future Work

97 We have used Restricted Algorithmic C to create an ISA-level simulator for a representative
 98 CPU, and successfully employed the RAC toolchain, initially developed to verify floating
 99 point hardware, to translate this simulator to ACL2. We produced small binary programs
 100 for the ISA that we used to validate the simulator, then utilized the ACL2 Codewalker
 101 decompilation-into-logic facility to prove those programs correct.

102 Future work includes refining the RAC toolchain to allow tail recursion, support for ACL2
 103 typed records and/or stobj's, as well as investigating the verification of hardware/software
 104 codesigns. Codewalker is an experimental tool, and while its ability to comprehend new
 105 instruction sets and automatically create semantic functions, clock functions, etc., is impres-
 106 sive, it is difficult to configure. Additionally Codewalker is only useful for single subroutines;
 107 future work is needed to improve usability and scalability.

108 — References —

- 109 1 David S. Hardin. Reasoning about LLVM code using Codewalker. In *Proceedings of the 13th*
 110 *International Workshop on the ACL2 Theorem Prover and its Applications*, volume 192, pages
 111 79–92. EPTCS, 2015. doi:10.4204/EPTCS.192.7.
- 112 2 Magnus O. Myreen. *Formal verification of machine-code programs*. PhD thesis, University of
 113 Cambridge, 2009.
- 114 3 John W. O’Leary and David M. Russinoff. Modeling algorithms in SystemC and ACL2.
 115 In *Proceedings of the 12th International Workshop on the ACL2 Theorem Prover and its*
 116 *Applications*, volume 152, pages 145–162. EPTCS, 2014. doi:10.4204/EPTCS.152.12.
- 117 4 David M. Russinoff. *Formal Verification of Floating-Point Hardware Design: A Mathematical*
 118 *Approach*. Springer, 2018.