Formal Verification of x86 Machine-Code Programs

Computer Architecture and Program Analysis



Shilpi Goel shigoel@cs.utexas.edu Department of Computer Science The University of Texas at Austin

Software and Reliability

Can we rely on our software systems?

Recent example of a serious bug:

CVE-2016-5195 or "Dirty COW"



- Privilege escalation vulnerability in Linux
- E.g.: allowed a user to write to files intended to be read only
- Copy-on-Write (COW) breakage of private read-only memory mappings
- Existed since around v2.6.22 (2007) and was fixed on Oct 18, 2016

Software Formal Verification: proving or disproving that the *implementation* of a program meets its *specification* using mathematical techniques

Software Formal Verification: proving or disproving that the *implementation* of a program meets its *specification* using mathematical techniques

Suppose you needed to count the number of 1s in the binary representation of a natural number (*population count*).

Specification:

```
popcountSpec(v): [v: natural number]
if v <= 0 then
   return 0
else
   lsb = v & 1
   v = v >> 1
   return (lsb + popcountSpec(v))
endif
```

Specification:

```
popcountSpec(v): [v: natural number]
if v <= 0 then
   return 0
else
   lsb = v & 1
   v = v >> 1
   return (lsb + popcountSpec(v))
endif
```

Specification:

```
popcountSpec(v): [v: natural number]
if v <= 0 then
    return 0
else
    lsb = v & 1
    v = v >> 1
    return (lsb + popcountSpec(v))
endif
```

Implementation:

```
int popcount_32 (unsigned int v) {
    v = v - ((v >> 1) & 0x55555555);
    v = (v & 0x33333333) + ((v >> 2) & 0x33333333);
    v = ((v + (v >> 4) & 0xF0F0F0F) * 0x1010101) >> 24;
    return(v);
}
```

Specification:

```
popcountSpec(v): [v: natural number]
if v <= 0 then
   return 0
else
   lsb = v & 1
   v = v >> 1
   return (lsb + popcountSpec(v))
endif
```

Implementation:

```
int popcount_32 (unsigned int v) {
    v = v - ((v >> 1) & 0x55555555);
    v = (v & 0x33333333) + ((v >> 2) & 0x33333333);
    v = ((v + (v >> 4) & 0xF0F0F0F) * 0x1010101) >> 24;
    return(v);
}
```

Do the specification and implementation behave the same way for all inputs?

Suppose you needed to check if a given natural number is a power of 2.

Specification:

```
isPowerOfTwoSpec(x): [x: natural number]
if x == 0 then
   return 0
else
   if x == 1 then
      return 1
   else
      if remainder(x,2) == 0 then
         return isPowerOfTwoSpec(x/2)
      else
         return 0
      endif
   endif
endif
```

Can you trust your specification?

Can you trust your specification?

Correctness of isPowerOfTwoSpec:

- 1. If is PowerOfTwoSpec(v) returns 1, then there exists a natural number n such that $v = 2^n$.
- 2. If v = 2ⁿ, where n is a natural number, then isPowerOfTwoSpec(v) returns 1.

Can you trust your specification?

Correctness of isPowerOfTwoSpec:

- 1. If is PowerOfTwoSpec(v) returns 1, then there exists a natural number n such that $v = 2^n$.
- 2. If v = 2ⁿ, where n is a natural number, then isPowerOfTwoSpec(v) returns 1.

Implementation:

```
bool powerOfTwo (long unsigned int v) {
    bool f;
    f = v && !(v & (v - 1));
    return f;
}
```

Can you trust your specification?

Correctness of isPowerOfTwoSpec:

- 1. If is PowerOfTwoSpec(v) returns 1, then there exists a natural number n such that $v = 2^{n}$.
- 2. If v = 2ⁿ, where n is a natural number, then isPowerOfTwoSpec(v) returns 1.

Implementation:

```
bool powerOfTwo (long unsigned int v) {
    bool f;
    f = v && !(v & (v - 1));
    return f;
}
```

Do the specification and implementation behave the same way for all inputs?

Inspection of a Program's Behavior

• Testing:

× Exhaustive analysis is infeasible

Formal Verification:

- Wide variety of techniques
 - Lightweight: e.g., checking if array indices are within bounds
 - Heavyweight: e.g., proving functional correctness

Example: Pop-Count Program

popcount 64:				
89 fa	mov	%edi,%edx		
89 d1	mov	%edx,%ecx		
d1 e9	shr	%ecx		
81 e1 55 55 55 55	and	\$0x55555555,%ecx		
29 ca	sub	%ecx,%edx		
89 d0	mov	%edx,%eax		Functional Correctness:
c1 ea 02	shr	\$0x2,%edx		PAV = papaquatSpaq(v)
25 33 33 33 33	and	\$0x33333333,%eax		RAX = popeounespec(v)
81 e2 33 33 33 33	and	\$0x33333333,%edx		
01 c2	add	%eax,%edx		T
89 d0	mov	%edx,%eax		specification function
c1 e8 04	shr	\$0x4,%eax		specification function
01 c2	add	%eax,%edx		
48 89 18	mo∨	%rd1,%rax		
48 CI E8 20	SNT	\$0x20,%Tax	popco	<pre>puntSpec(v):</pre>
81 e2 01 01 01 01 80 c1	and	\$0X1010101,%eax	Γ	incided int]
	niov	%eax,%ecx	ן בי. נ	Insigned inc]
	and			
29 c8	sub	%ecx %eax	if v	< - 0 then
89 c1	mov	%eax %ecx		
c1 e8 02	shr	\$0x2,%eax	l re	eturn Ø
81 e1 33 33 33 33	and	\$0x33333333.%ecx		
25 33 33 33 33	and	\$0x33333333,%eax	else	
01 c8	add	%ecx,%eax	l	sb = v & 1
89 c1	mov	%eax,%ecx		
c1 e9 04	shr	\$0x4,%ecx		= V $>>$ 1
01 c8	add	%ecx,%eax	l re	eturn (lsb + popcountSpec(v))
25 Of Of Of Of	and	\$0xf0f0f0f,%eax	ondif	
69 d2 01 01 01 01	imul	\$0x1010101,%edx,%edx	enari	-
69 c0 01 01 01 01	imul	\$0x1010101,%eax,%eax		
c1 ea 18	shr	\$0x18,%edx		
c1 e8 18	shr	\$0x18,%eax		
01 d0	add	%edx,%eax		
c3	retq			

Case Study: Pop-Count Program

```
(defthm x86-popcount-64-symbolic-simulation
  (implies
  (and (x86p x86)
        (equal (model-related-error x86) nil)
        (unsigned-byte-p 64 n)
        (equal n (read 'register *rdi* x86))
        (equal *popcount-64-program*
               (read 'memory
                     (address-range
                      (read 'pc x86)
                      (len *popcount-64-program*))
                     x86)))
   (equal (read 'register *rax* (x86-run *num-of-steps* x86))
          (popcountSpec n))))
```

- Build a mathematical or *formal* model of programs
- Prove *theorems* about this model in order to establish program properties

Build a mathematical or *formal* model of programs

ISA model

• Prove *theorems* about this model in order to establish program properties

- Build a mathematical or *formal* model of programs
- ISA model
- Prove *theorems* about this model in order to establish program properties

- Defines the machine language
- Specification of state (registers, memory), machine instructions, instruction encodings, etc.

- Build a mathematical or *formal* model of programs
- ISA model
- Prove *theorems* about this model in order to establish program properties

- Defines the machine language
- Specification of state (registers, memory), machine instructions, instruction encodings, etc.
- An ISA model specifies the behavior of each machine instruction in terms of *effects made to the processor state*.

- Build a mathematical or *formal* model of programs ISA model
- Prove *theorems* about this model in order to establish program properties

- Defines the machine language
- Specification of state (registers, memory), machine instructions, instruction encodings, etc.
- An ISA model specifies the behavior of each machine instruction in terms of *effects made to the processor state*.
- All high-level programs compile down to machine-code programs.
 - A program is just a sequence of machine instructions.

- Build a mathematical or *formal* model of programs ISA model
 - Prove *theorems* about this model in order to establish program properties

- Defines the machine language
- Specification of state (registers, memory), machine instructions, instruction encodings, etc.
- An ISA model specifies the behavior of each machine instruction in terms of *effects made to the processor state*.
- All high-level programs compile down to machine-code programs.
 A program is just a sequence of machine instructions.
- We can reason about a program by *inspecting the cumulative effects of its constituent instructions on the machine state*.

Why Not Use Abstract Machine Models?

Abstract vs. Concrete Machine Models

Machine Models



Assembly



Data

1) char 2) int, float 3) double 4) struct, array 5) pointer

1) byte

Control

- 1) loops
- 2) conditionals
- 3) switch
- 4) proc. call
- 5) proc. return

- 1) branch/jump
- 2) call
- 3) ret
- 3) 4-byte long word4) contiguous byte allocation

2) 2-byte word

5) address of initial byte

Why Not Use Abstract Machine Models?



Why x86 Machine-Code Verification?

• Why not high-level code verification?

- × Sometimes, high-level code is unavailable (e.g., malware)
- × High-level verification frameworks do not address compiler bugs
 - Verified/verifying compilers can help
 - × But these compilers typically generate inefficient code
- × Need to build verification frameworks for many high-level languages
- Why x86?
 - ✓ x86 is in widespread use

Goal: *Specify* and *verify* properties of x86 programs

Goal: *Specify* and *verify* properties of x86 programs



- **Program property:** statement about a program's behavior
 - One state, set of states, relationship between a set of final & initial states

Goal: *Specify* and *verify* properties of x86 programs





- **Program property:** statement about a program's behavior
 - One state, set of states, relationship between a set of final & initial states
- **Program's computation:** how the execution of each instruction transforms one state to another

Goal: *Specify* and *verify* properties of x86 programs



- **Program property:** statement about a program's behavior
 - One state, set of states, relationship between a set of final & initial states
- **Program's computation:** how the execution of each instruction transforms one state to another
- **Symbolic Executions:** a final (or next) x86 state is described in terms of *symbolic updates* made to the initial x86 state
 - Allows consideration of many, if not all, possible executions at once

Formal Tool Used: ACL2 Theorem-Proving System

- ACL²: A Computational Logic for Applicative Common Lisp
 - Programming Language
 - Mathematical Logic
 - Mechanical Theorem Prover



- See <u>ACL2 Home Page</u> for more details.
 - Extensive documentation!
 - ACL2 Research Group located at GDC 7S

x86 ISA Model

x86 ISA Model



A Run of the x86 Interpreter that executes k instructions

Interpreter-Style Operational Semantics: x86 ISA model is a machinecode interpreter written in ACL2's formal logic

- **x86 State:** specifies the components of the ISA
- Run Function: takes n steps or terminates early if an error occurs
- Step Function: fetches, decodes, and executes one instruction
- Instruction Semantic Functions: specifies instructions' behavior

Run Function

Recursively defined interpreter that specifies the x86 model

```
run(n, x86):
if n == 0 then
  return x86
else
  if model-related error encountered then
    return x86
  else
    run(n - 1, step(x86))
  end if
end if
```

Step Function

State-transition function that corresponds to the execution of a single x86 instruction

```
step(x86):
pc = rip(x86)
[prefixes, opcode, ..., imm] = Fetch-and-Decode(pc, x86)
case opcode:
    #x00 -> add-semantic-fn(prefixes, ..., imm, x86)
    ....
    #xFF -> inc-semantic-fn(prefixes, ..., imm, x86)
```

Instruction Semantic Functions

- A semantic function describes the effects of executing an instruction.
 - Input: x86 state and decoded parts of the instruction
 - Output: next x86 state
- Every instruction has its own semantic function.

add-semantic-fn(prefixes, ..., imm, x86): operand1 = getOperand1(prefixes, ..., imm, x86) operand2 = getOperand2(prefixes, ..., imm, x86) resultSum = fix(operand1 + operand2, ...) resultFlags = computeFlags(operand1, operand2, result, x86) x86 = updateState(resultSum, dst, resultFlags) return x86

Obtaining the x86 ISA Specification



Running tests on x86 machines



x86 State

Focus: Intel's 64-bit mode

Sixteen 64-bit Registers				
	se Registers	i	2^64 -1	
Six 16-bit Registers	Segment Regis	ters		
64-bits 64-bits	RFLAGS Regist	er 1 Pointer Regis ¹	ter)	
PU Registers				
Eight 80-bi Registers	Floating-Poir Data Registe	0		
<u> </u>	16 bits 16 bits 16 bits bits bits	Control Regi Status Regis Tag Register Opcode Regi FPU Instruct FPU Data (O	ster ter ster (11-bits) ion Pointer Reg perand) Pointe	gister Register
MMX Registers Eight 64-bit Registers MM		X Registers		
KMM Registers				
Sixteen Regi	n 128-bit isters		XMM Regist	ers
	3	2-bits	MXCSR Regis	ter



x86 State

Focus: Intel's 64-bit mode



Figure 3-2. 64-Bit Mode Execution Environment

Source: Intel Manuals Figure 2-2. System-Level Registers and Data Structures in IA-32e Mode

Model Validation

How can we know that our model faithfully represents the x86 ISA? Validate the model to increase trust in the applicability of formal analysis



Symbolic Execution

Supporting Symbolic Execution



Rules (theorems) describing interactions between these reads and writes to the x86 state enable *symbolic execution* of programs.

non-interference



non-interference



non-interference































Symbolic Execution

These read-over-write and write-over-write lemmas operate on **symbolic expressions** that describe the program's behavior.

```
(implies
 (preconditions loc val x86)
 (let ((old-rbx (read 'register *rbx* x86))
        (old-pc (read 'pc x86)))
      (equal
        (x86-run (clk) x86)
        (write 'register *rax* old-rbx
        (write 'pc (+ 18 old-pc)
              (write 'memory loc val x86))))))
```

Also, we can project out *relevant* parts of the resulting state.

Conclusions

What I Haven't Talked About Today...

1. How to prove theorems using a mechanical theorem prover

- Useful to reason about both hardware and software
- Fall'16 Grad-level Course: *Programming Languages*
- Spring'17 Grad-level Course: Recursion and Induction
- 2. Supervisor-mode features of the x86 ISA
 - Useful for developing and analyzing kernel programs
 - An advanced architecture class
 - An OS class

Opportunities for Future Research

Operating System Verification	<i>User-friendly Program Analysis</i>
detect reliance on non-portable or	automate the discovery of
undefined behaviors	preconditions
<i>Multi-process/threaded Program</i>	Reasoning about the Memory
<i>Verification</i>	System
reason about concurrency-related	determine if caches are (mostly)
issues	transparent, as intended
<i>Firmware Verification</i>	<i>Micro-architecture Verification</i>
formally specify software/hardware	x86 ISA model serves as a build-to
interfaces	specification

Resources

We have exciting research and engineering projects in this area! Please feel free to email if you want to know more.

- See <u>ACL2 Home Page</u>
- Talk to people on GDC 7S
- See some publications

Publications

Shilpi Goel, Warren A. Hunt, Jr., and Matt Kaufmann. *Abstract Stobjs and Their Application to ISA Modeling*. In Proceedings of the ACL2 Workshop 2013, EPTCS 114, pp. 54-69, 2013

Shilpi Goel and Warren A. Hunt, Jr. *Automated Code Proofs on a Formal Model of the x86*. In Verified Software: Theories, Tools, Experiments (VSTTE'13), volume 8164 of Lecture Notes in Computer Science, pages 222–241. Springer Berlin Heidelberg, 2014

Shilpi Goel, Warren A. Hunt, Jr., Matt Kaufmann, and Soumava Ghosh. *Simulation and Formal Verification of x86 Machine-Code Programs That Make System Calls*. In Proceedings of the 14th Conference on Formal Methods in Computer-Aided Design (FMCAD'14), pages 18:91–98, 2014

Shilpi Goel, Warren A. Hunt, Jr., and Matt Kaufmann. *Engineering a Formal, Executable x86 ISA Simulator for Software Verification*. In Provably Correct Systems (ProCoS), 2015

Shilpi Goel. *Formal Verification of Application and System Programs Based on a Validated x86 ISA Model*. Ph.D. Dissertation, The University of Texas at Austin, 2016 [Source Code]

[Documentation]

x86isa in the ACL2+Community Books Manual

