Sequoll: a framework for model checking binaries



Bernard Blackham, Gernot Heiser

School of Computer Science & Engineering, UNSW Software Systems Research Group, NICTA Sydney, Australia



Australian Government

Department of Broadband, Communications and the Digital Economy

Australian Research Council





SYDNEY





1) Griffith







Wednesday, 10 April 13

Motivation



 The desire for a *trustworthy* kernel to build reliable mixed-criticality real-time systems



Motivation



 The desire for a *trustworthy* kernel to build reliable multi-criticality real-time systems

- Using **seL4** to guarantee:
 - functional correctness through formal proof

(Klein et al., SOSP 2009)

 – timing constraints through sound WCET analysis (Blackham et al., RTSS 2011)

Motivation

- **NICTA**
- Current analysis uses annotations to specify:
 - loop counts
 - infeasible paths
- We want to reduce scope for errors in WCET analysis.

Results





With infeasible path information

© NICTA 2013

NIC⁻

seL4 is large



Large by WCET standards

C source

Binary (ARM)

~8,700 lines 316 functions 76 loops ~10,000 instructions

228 functions

56 loops

2,384 basic blocks

~400,000 basic blocks when inlined

























Pruning infeasible paths



Express constraints as one of:

a conflicts with b when called under f

or

a is consistent with b when called under f



NICTA

Verifying annotations

- How can we verify these manual annotations?
- For two basic blocks **a**, **b**:
 - show all paths execute ${\bf a}$ and ${\bf b}$ the same number of times
 - show all paths execute at most one of ${\boldsymbol{a}}$ or ${\boldsymbol{b}}$
- For some loop ℓ in the binary:
 - How many times can l iterate?

Why model checking?

- Many infeasible paths and even some loop bounds cannot be determined without program invariants
- Invariants are known to the formal proof
- Formal proof statements and invariants are a natural fit for model checkers

NICTA

Why model checking?

```
int count_bits(uint32_t x) {
    int c = 0;
    while (x != 0) {
        x &= x - 1; /* clear lowest bit */
        c++;
     }
    return c;
}
```

What loop bound?

© NICTA 2012

Wednesday, 10 April 13

From binary → model checker

Extracting instruction semantics

 Reused existing formalization of ARM ISA (Fox & Myreen, ITP 2010)

0

return

Computing a program slice

Find all nodes which can impact upon execution of that block through:

- data dependencies
- control dependencies

Reducing the CFG further

init(n) := 0
next(n) := case
n=0: 1
n=1 && cond_1: 1
n=1 && !cond_1: 2
n=2: 2

VAR

init(c) := 0
next(c) := case
n=0: 0
n=1: c + 1
else: c

n : -1..2; memRead : unsigned word[8]; cond 3 : boolean; r3 2 : unsigned word [32]; c : unsigned word[32]; DEFINE psrZ 3 := (((1) >= 32) ? 0ud32 cond 1 := !(psrZ 3); cond 2 := psrZ 1; r3 1 := r0 0; r3 3 := ((1) >= 32) ? Oud32 0 : psrZ 1 := (r0 0) = (0ud32 0);FROZENVAR r0 0 : unsigned word [32]; r14 0 : unsigned word [32]; ASSIGN next(cond 3) := case n=1: cond 5;n=0: cond 2;TRUE: cond 3; esac; next(r3_2) := case n=1: r3_3; n=0: r3 1;

TRUE: r3 2;

esac;

For loop counts:

Ask a model checker: is c < k? (and binary search for k)

For conflict constraints:

Add "visited" flag for nodes a and b

Ask a model checker to ensure that *a* and *b* are never both true

In CTL: AG ! (visitedA & visitedB)

For consistency constraints:

Add "count" variables for nodes a and b

Ask a model checker to ensure counts for *a* and *b* are equal (eventually)

In CTL: AF (countA = countB)

Loop bound verification

Can compute 18/32 loop bounds in seL4:

- 1 loop depends upon invariants in the proof
- 1 loop cannot be bounded due to complex exit conditions
 - model checker attempts to find the *smallest* loop bound
 - complex state space must be explored to deduce bound
- 12 loops, all identical in structure, cannot be bounded due to poor memory alias analysis

NIC

Loop bound verification

```
void f(uint32 t n)
  ł
    uint32 t i = 1 << n;
    if (i > 256)
      i = 256;
    while (i > 0) {
      i -= 4;
    }
  }
If n \ge 32, i is undefined
If n \leq 1, loop is infinite
```


Of 35 infeasible path constraints

-4 validated

- -1 shown untrue (oops!)
- -11 cannot be validated without better alias analysis
- -19 depend on kernel invariants

Results

Estimated worst-case execution time of seL4

With verified infeasible path information With all infeasible path information

© NICTA 2013

NIC

Research directions

- Integrate with proof invariants
- Automate infeasible path detection (WIP)
- Use a faster ISA formalization

Summary

Sequoll is able to:

- apply model checkers to reason about compiled ARM binaries
- validate manual infeasible path annotations
- compute "interesting" loop bounds
- (eventually) integrate formal proof with infeasible path information

reduce scope for errors in WCET analysis!

Download it!

http://www.ssrg.nicta.com.au/software/TS/wcet-tools

Bernard.Blackham@nicta.com.au