

Translation Validation for a Verified OS Kernel

Thomas Sewell¹ Magnus Myreen² Gerwin Klein¹

¹UNSW & NICTA

²University of Cambridge

12 Feb 2013











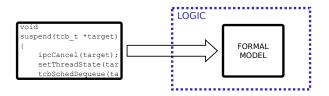




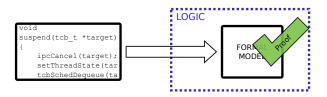




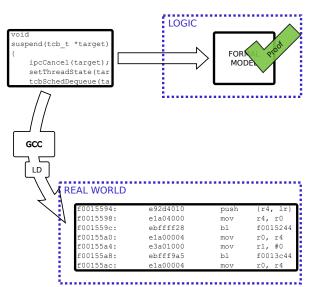




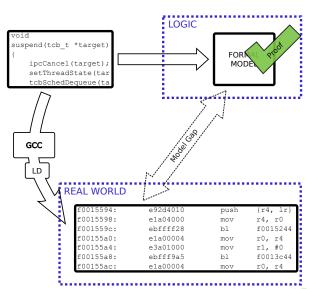






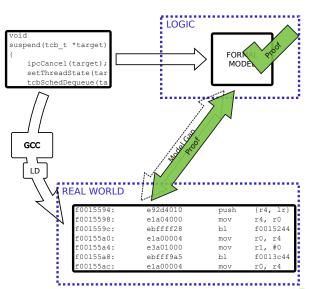




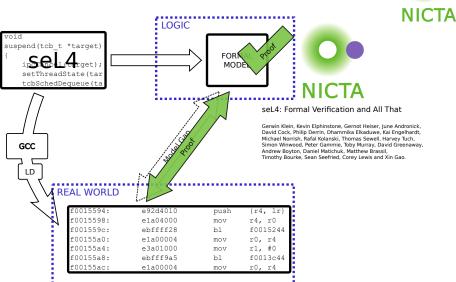




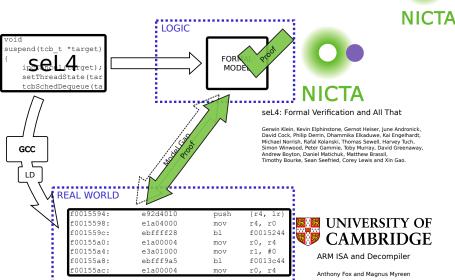












Overview



This talk in one bullet point:

 We can prove the binary refines the formal model, for seL4's verified components and gcc-4.5.1 -01.



This talk in one bullet point:

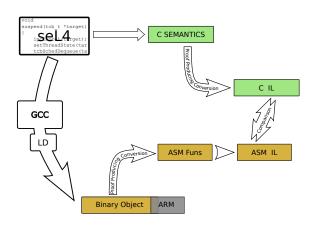
 We can prove the binary refines the formal model, for seL4's verified components and gcc-4.5.1 -01.

Talk contents:

- Structure of proof and approach.
 - Alternative approaches.
 - C Semantics & C Standard.
 - Challenges.
 - Restrictions & workarounds.

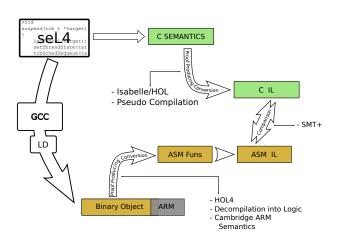
Approach





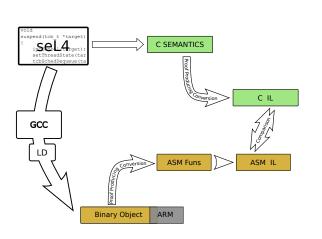
Approach

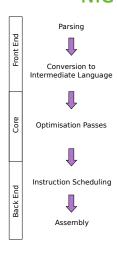




Approach







C Program Semantics



```
void
suspend(tcb_t *target)
{
    ipcCancel(target);
    setThreadState(tar
    tcbSchedDequeue(ta
C SEMANTICS
```

C Program Semantics



Maps syntax of C to a deeply embedded language in Isabelle/HOL with an operational semantics.

```
if (...) {...} \Rightarrow IF (...) THEN ... FI

f (1, 2); \Rightarrow CALL f_'proc (1, 2);;

x ++; \Rightarrow

('x :== 'x + 1);;

*p = *q; \Rightarrow
```

(h_val 'q 'mem) 'mem

mem :== h_upd 'p

C Program Semantics



Maps syntax of C to a deeply embedded language in Isabelle/HOL with an operational semantics.

Partial semantics to explain undefined behaviour.

C Standard Semantics



Aside: Why not just trust the compiler?

The ptr_valid assertion used in Guard is subtle.

The **object rule** says that a pointers may come from arithmetic within an object, & and malloc.

C Standard Semantics



Aside: Why not just trust the compiler?

The ptr_valid assertion used in Guard is subtle.

The **object rule** says that a pointers may come from arithmetic within an object, & and malloc.

What about casts from numbers?

C Standard Semantics



Aside: Why not just trust the compiler?

The ptr_valid assertion used in Guard is subtle.

The **object rule** says that a pointers may come from arithmetic within an object, & and malloc.

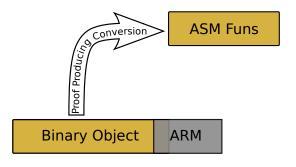
What about casts from numbers?

There are multiple interpretations of the C language.

- NICTA seL4: Liberal, portable assembler, soundy.
 - Strict aliasing rule but not object rule.
- CompCert: Conservative.

Decompilation





Example Decompilation



```
uint avg (uint i, uint j) {
  return (i + j) / 2;
}
<avg>:
avg+0
         e0810000
                       add r0, r1, r0
                                          // add r1 to r0
avg+4
         e1a000a0
                       lsr r0, r0, #1
                                          // shift r0 right
avg+8
         e12fff1e
                                          // return
                       bx lr
```

Example Decompilation



```
uint avg (uint i, uint j) {
  return (i + j) / 2;
}
<avg>:
avg+0
         e0810000
                      add r0, r1, r0
                                        // add r1 to r0
avg+4
         e1a000a0
                      lsr r0, r0, #1
                                        // shift r0 right
         e12fff1e
avg+8
                      bx lr
                                         // return
```

avg
$$(r_0, r_1)$$
 = let $r_0 = r_1 + r_0$ in
let $r_0 = r_0 >>> 1$ in
 r_0

Example Decompilation



avg
$$(r_0, r_1)$$
 = let $r_0 = r_1 + r_0$ in let $r_0 = r_0 >>> 1$ in r_0

```
  \{ \text{R0 } r_0 * \text{R1 } r_1 * \text{R14 } lr * \text{PC } p \} 
  p : \text{e0810000 e1a000a0 e12fff1e} 
  \{ \text{R0 } (\text{avg } (r_0, r_1)) * \text{R1} * \text{R14} * \text{PC } lr \} 
  \text{enfield OS Kernel}
```

Challenges for Decompilation



```
uint avg8 (uint i1, i2, i3, i4, i5, i6, i7, i8) {
  return (i1+i2+i3+i4+i5+i6+i7+i8) / 8;
}
```

Challenges for Decompilation



```
uint avg8 (uint i1, i2, i3, i4, i5, i6, i7, i8) {
  return (i1+i2+i3+i4+i5+i6+i7+i8) / 8;
<avg8>:
e0811000
             add r1, r1, r0
e0811002
             add r1, r1, r2
                                  // load
e59d2000
             ldr r2, [sp]
e0811003
             add r1, r1, r3
e0810002
             add r0, r1, r2
e99d000c
                                  // load
             ldmib sp, {r2, r3}
e0800002
             add r0, r0, r2
e0800003
             add r0, r0, r3
e59d300c
             ldr r3, [sp, #12]
                                 // load
e0800003
             add r0, r0, r3
e1a001a0
             lsr r0, r0, #3
e12fff1e
             bx lr
```

Stack and Heap



Aside: Hiding stack accesses mean they must not be aliased.

Our C semantics forbids pointers to the stack.

We also eliminate padding, clearly separating:

- the heap, under user control.
- the stack, under compiler control.

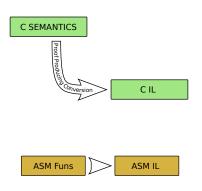
Enables a simple notion of correct compilation:

$$\forall (\mathit{in}, \mathit{in_heap}) \in \mathsf{domain}(\mathfrak{C}). \ \mathfrak{C}(\mathit{in}, \mathit{in_heap}) = \mathfrak{B}(\mathit{in}, \mathit{in_heap})$$

This would be difficult with higher level optimisations.

Conversion to Graph





Not going to discuss this in detail.

Graph Refinement

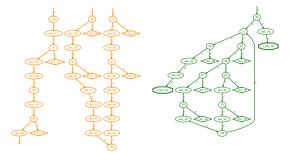


The proof of refinement between graphs involves two processes:

- A search process, which heuristically discovers a proof object.
- A check process, which checks the proof is sound.

This follows Pnueli's translation validation design.

Both processes use SMT solving extensively.



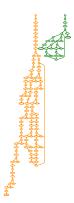




Proof objects contain:

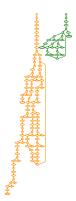
 An inlining of all needed function bodies into one space.





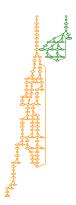
- An inlining of all needed function bodies into one space.
- Restrict rules, which observe that a given point in a loop may be reached only n times.





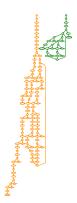
- An inlining of all needed function bodies into one space.
- Restrict rules, which observe that a given point in a loop may be reached only n times.
- Split rules, which observe that a C loop point is reached as often as a loop point in the binary.





- An inlining of all needed function bodies into one space.
- Restrict rules, which observe that a given point in a loop may be reached only n times.
- Split rules, which observe that a C loop point is reached as often as a loop point in the binary.
 - Checked by k-induction.



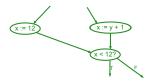


- An inlining of all needed function bodies into one space.
- Restrict rules, which observe that a given point in a loop may be reached only n times.
- Split rules, which observe that a C loop point is reached as often as a loop point in the binary.
 - Checked by k-induction.
 - Parameter eqs must relate enough of binary state to C state to relate events after the loop.



SMT problems generated contain:

- Fixed-length values and arithmetic: word32, +, -, <= etc.
- Arrays to model the heap: heap :: word30 => word32.
- If-then-else operators to handle multiple paths.

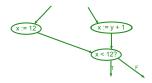


Validity assertions and needed inequalities:
 ptr1_valid & ptr2_valid ⇒ ptr1 > ptr2 + 7 ∨ ptr2 > ptr1 + 15.



SMT problems generated contain:

- Fixed-length values and arithmetic: word32, +, -, <= etc.
- Arrays to model the heap: heap :: word30 => word32.
- If-then-else operators to handle multiple paths.



Validity assertions and needed inequalities:
 ptr1_valid & ptr2_valid ⇒ ptr1 > ptr2 + 7 ∨ ptr2 > ptr1 + 15.

Strong compatibility with **SMTLIB2 QF_ABV**.



Strong similarity to **QF_ABV** category of the SMT competition.

We ran this experiment with Z3 (version 4.0) and SONOLAR (version 2012-06-14).

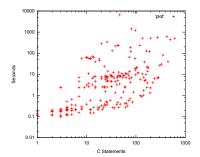
The solvers are efficient at producing both sat and unsat results, which is important in discovering and checking a proof.

Results



The proof rules and inlining heuristic mentioned are sufficient for seL4's verified code with gcc-4.5.1 -01.

Nested loops and some -02 loop optimisations are not yet handled.



Conclusions



Translation validation can scale up to substantial problem size, using naive approaches, for a carefully managed problem.

Supporting factors:

- Simple looping structure.
- C Semantics already at the level of bits and bytes.
- Clear separation of compiler and user control.
- Strong compatibility with SMT QF_ABV.

Software is available at http://www.ssrg.nicta.com.au/software/TS/graph-refine