

Verification of an ML compiler

Lecture I:

An introduction to compiler verification

Marktoberdorf Summer School MOD 2017

Magnus O. Myreen, Chalmers University of Technology

Introduction

Your program crashes.

Where do you look for the fault?

- Do you look at *your source code*?
- Do *look at the code for the compiler* you used?

users want to rely on compilers

Verified compilers



What?

A verified compiler is a compiler that comes with a machine-checker proof. The proof states that the compiler *preserves the behaviour of source programs*.

Traditional compiler development relies on testing.
Compiler verification is considered too costly.



cost reduction?

All (unverified) compilers have bugs

“ Every compiler we tested was found to crash and also to silently generate wrong code when presented with valid input. ”

PLDI'11

Finding and Understanding Bugs in C Compilers

Xuejun Yang Yang Chen Eric Eide John Regehr

“ [The verified part of] CompCert is the only compiler we have tested for which Csmith cannot find wrong-code errors. This is not for lack of trying: we have devoted about six CPU-years to the task. ”

In this paper, we report the results of our bug-hunting study. Our first contribution is to advance the state of the art in compiler testing. Unlike previous tools, Csmith generates programs that cover a large subset of C while avoiding the unspecified behaviors that would destroy its ability to

was heavily patched; the base version of GCC and

We created Csmith, a randomized test-case generator that supports C99 and C11. Csmith generates test cases that

Motivations

Bugs in compilers are not tolerated by users

Bugs can be hard to find by testing

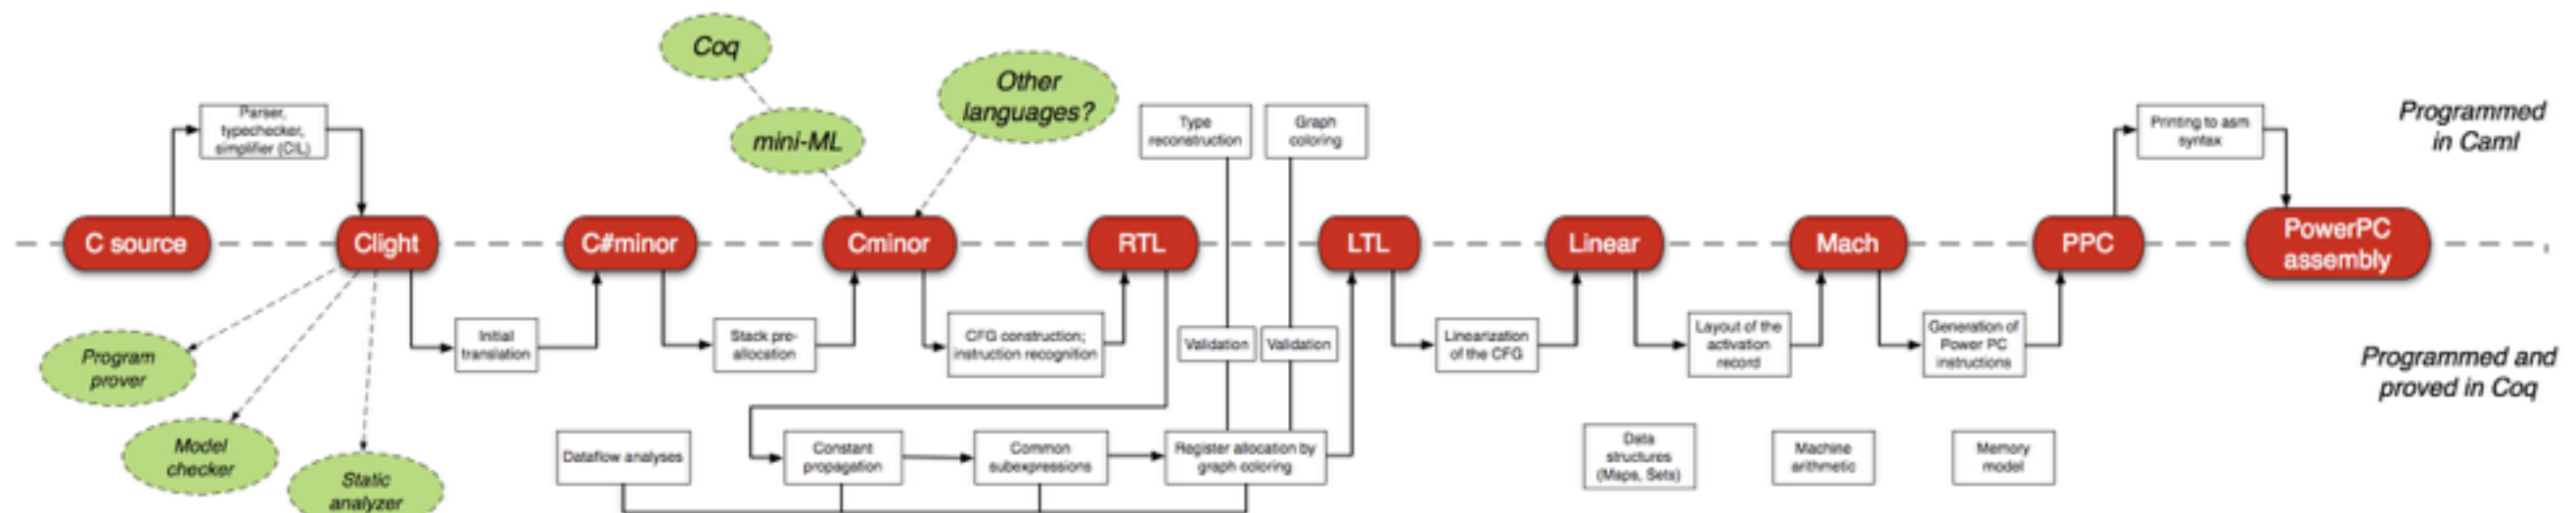
Verified compilers must be used for verification of source-level programs to imply guarantees at the level of verified machine code

Research question: how easy (cheap) can we make compiler verification?

State of the art

CompCert

CompCert C compiler



Leroy et al. Source: <http://compcert.inria.fr/>

Compiles C source code to assembly.

Has **good performance numbers**

Proved correct in Coq.

<http://compcert.inria.fr/>

CakeML compiler

Compiles CakeML concrete syntax to machine code.

Proved correct in HOL4.

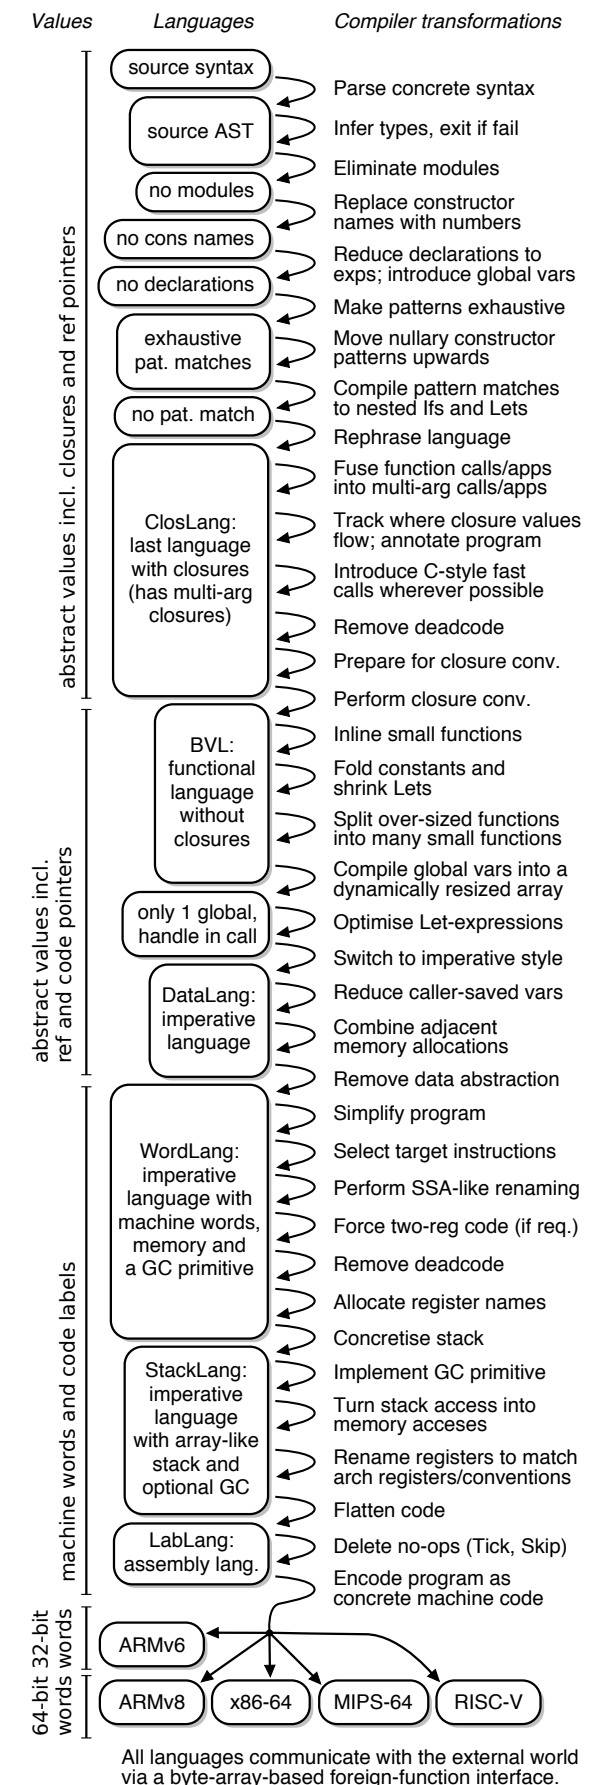
Has **mostly good performance numbers** (later lecture)

Known as the first verified compiler to be bootstrapped.

I'm one of the six developers behind version 2 (diagram to the right).

later lecture zooms in

<https://cakeml.org/>



robust, inflexible

proved to always
work correctly

Verified compilers

A spectrum

more flexible,
but can be fragile

produces a proof for each run

Proof-producing compilers

...

Pilsner

Fiat

CompCert C compiler

Cogent

CakeML compiler

*Translation validation for
a verified OS kernel*

CompCertTSO

These 4 lectures on

Verification of an ML compiler

will focus on key concepts rather than implementation details

The lectures:

Lecture 1: introduction to compiler verification

Lecture 2: data representation and garbage collection

Lecture 3: closures and call optimisations

Lecture 4: compiler bootstrapping



but feel free to ask

The lectures will not cover:

Verification of parsing and ML-style type inference

Optimisations (except call optimisations in lecture 3)

Modelling and verification of I/O

How to manage a large code base of proofs

The lectures:

Lecture 1: introduction to compiler verification

Lecture 2: data representation and garbage collection

Lecture 3: closures and call optimisations


Lecture 4: compiler bootstrapping

Let's get started!

Lecture 1: introduction to compiler verification

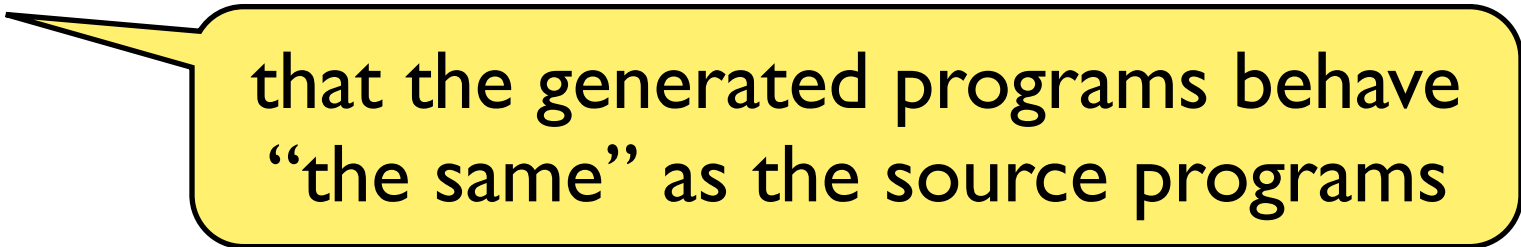
Questions

Should we write a new compiler for the verification?



yes, because then it has
nice invariants

What should we prove about it?



that the generated programs behave
“the same” as the source programs

Do we need to use mechanised proof?



yes

Proving a compiler correct

Ingredients:

- a **formal logic** for the proofs
- **accurate models** of
 - the **source** language
 - the **target** language
 - the **compiler** algorithm

does this correspond
with reality?

requires *testing*

Tools:

- a **proof assistant** (software)

Method:

- prove simulation theorem using **induction**

Ingredient 1:

Formal logic

Formal logic

These lectures will use classical higher-order logic (HOL)

HOL = lambda calculus with simple ML-like types

Allows for quantification over functions and relations.

Example 1: Skolem

$$\vdash (\forall x. \exists y. P\ x\ y) \iff \exists f. \forall x. P\ x\ (f\ x)$$

Example 2: complete induction over nat:

$$\vdash (\forall n. (\forall m. m < n \Rightarrow P\ m) \Rightarrow P\ n) \Rightarrow \forall n. P\ n$$

Example 3: definition of list permutation:

$$\vdash \text{PERM } L_1\ L_2 \iff \forall x. \text{FILTER } ((=)\ x)\ L_1 = \text{FILTER } ((=)\ x)\ L_2$$

User-definitions in HOL

Datatypes:

```
datatype e = Var string
          | Num int
          | Add e e
          | Assign string e
```

```
datatype r = Rval int
          | Rbreak
          | Rfail
```

Recursive functions (that **terminate** or are **tail-recursive**):

```
FILTER P [] = []
```

```
FILTER P (h::t) = if P h then h::FILTER P t else FILTER P t
```

Inductive relations (also co-inductive relations):

$$(S1) \frac{\begin{array}{c} (t_1, s) \Downarrow_t (Rval\ n_1, s_1) \\ (t_2, s_1) \Downarrow_t r \end{array}}{(Seq\ t_1\ t_2, s) \Downarrow_t r}$$

$$(S2) \frac{\begin{array}{c} (t_1, s) \Downarrow_t (r, s_1) \\ \neg is_Rval\ r \end{array}}{(Seq\ t_1\ t_2, s) \Downarrow_t (r, s_1)}$$

Ingredient 2: Language models

Syntax

The abstract syntax is defined as a datatype.

A toy **source** language:

```
t = Dec string t
    | Exp e
    | Break
    | Seq t t
    | If e t t
    | For e e t
```

evaluates an expression

Break aborts a loop

For-loop

```
e = Var string
    | Num int
    | Add e e
    | Assign string e
```

expressions can have
side-effects (c.f. ML)

Syntax

The abstract syntax is defined as a datatype.

A toy *source* language:

```
t = Dec string t
    | Exp e
    | Break
    | Seq t t
    | If e t t
    | For e e t

e = Var string
    | Num int
    | Add e e
    | Assign string e
```

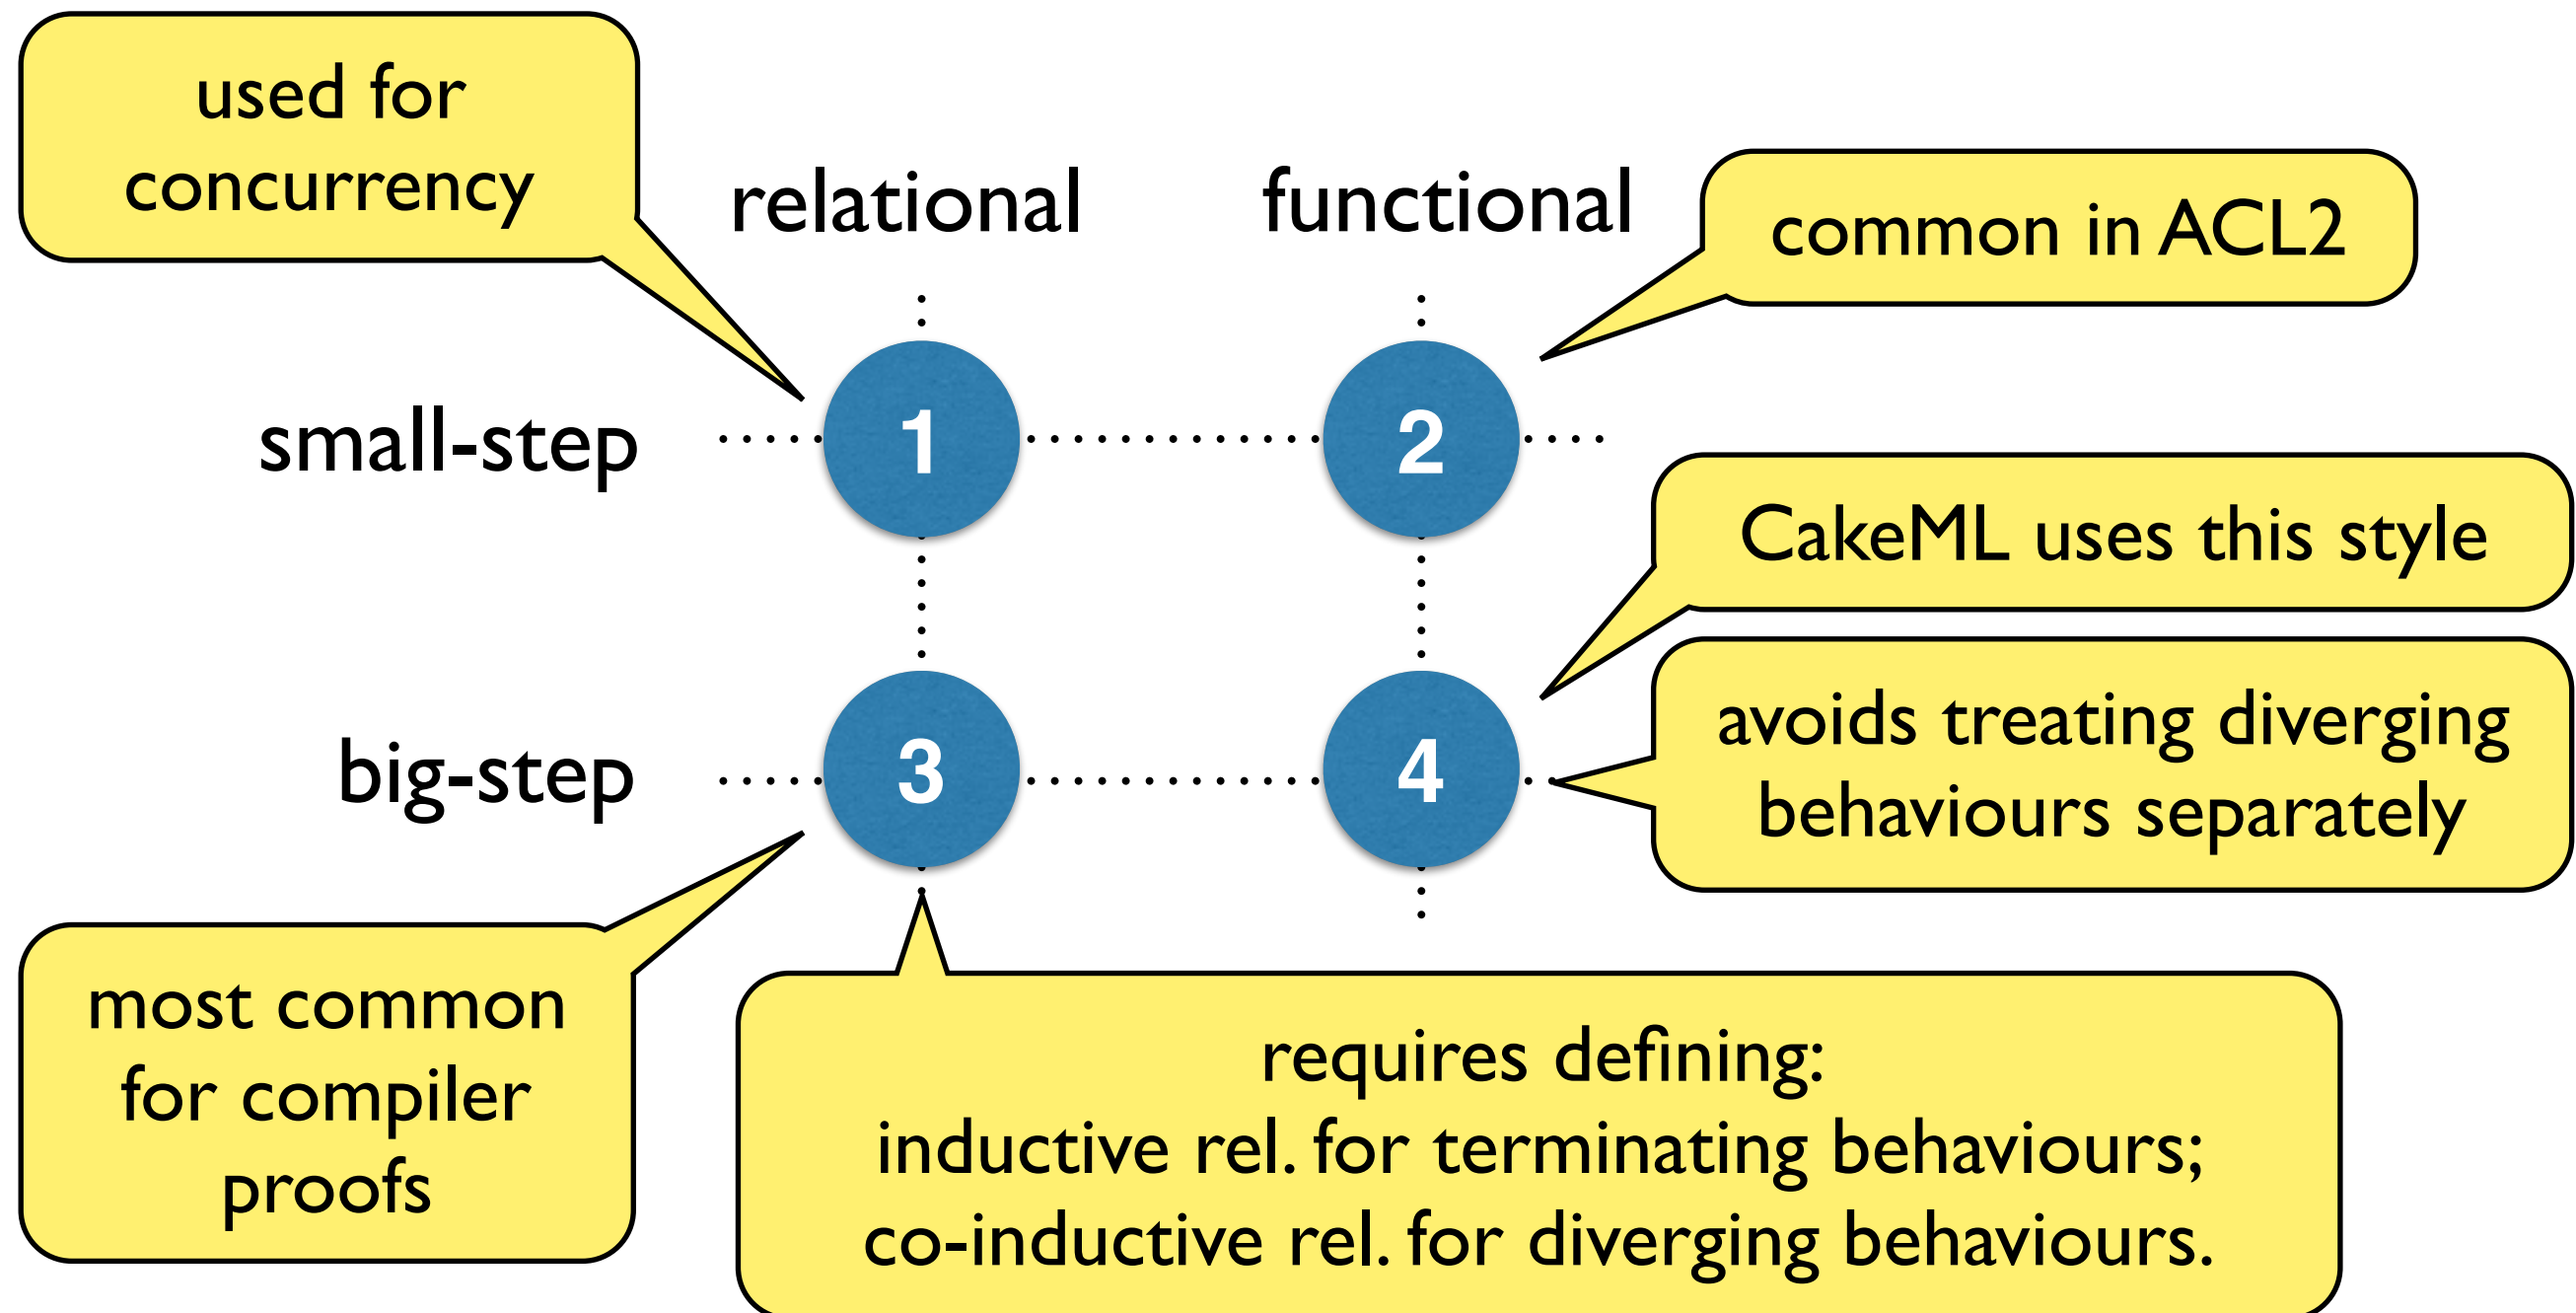
A toy *target* language:

```
inst
    = Add reg reg reg
    | Int reg int
    | Jmp offset
    | JmpIf reg offset
```

an assembly program is a
list of instructions

Operational Semantics — the options

There are 4 main variants one can choose from.



Towards a functional big-step semantics (I)

We define an interpreter for the source *in Standard ML*:

```
fun lookup y [] = NONE
  | lookup y ((x,v)::xs) = if y = x then SOME v else lookup y xs
```

```
fun run_e s (Var x) =
  (case lookup x s of
    NONE => (Rfail,s)
  | SOME v => (Rval v,s))
| run_e s (Num i) = (Rval i,s)
| run_e s (Assign (x, e)) =
  (case run_e s e of
    (Rval n1, s1) => (Rval n1, (x,n1)::s1)
  | r => r)
| run_e s (Add (e1, e2)) = ...
```

state is a simple assoc-list

assignments update the state

expression evaluation returns Rval or Rfail and new state

Towards a functional big-step semantics (2)

... and the evaluation of statements (type t):

```
fun run_t s (Exp e) = run_e s e
  | run_t s (Dec (x, t)) = run_t ((x,0)::s) t
  | run_t s Break = (Rbreak, s)
  | run_t s (Seq (t1, t2)) =
    (case run_t s t1 of
      (Rval _, s1) => run_t s1 t2
    | r => r)
  | run_t s (If (e, t1, t2)) =
    (case run_e s e of
      (Rval n1, s1) => run_t s1 (if n1 = 0 then t2 else t1)
    | r => r)
```

Break causes an Rbreak

Rbreak skips Seq-code

Towards a functional big-step semantics (3)

... so far everything we wrote in ML could be defined in HOL.

```
| run_t s (For (e1, e2, t)) =  
  (case run_e s e1 of  
    (Rval n1, s1) =>  
      if n1 = 0 then (Rval 0, s1) else  
        (case run_t s1 t of  
          (Rval _, s2) =>  
            (case run_e s2 e2 of  
              (Rval _, s3) => run_t s3 (For (e1, e2, t))  
              | r => r)  
            | (Rbreak, s2) => (Rval 0, s2)  
            | r => r)  
          | r => r)
```

... but now this recursive call
introduces potential
non-termination

Why is non-termination an issue?

Suppose HOL allowed definitions of non-terminating functions, e.g.

$$f\ n = f\ n + 1$$

then $f\ n - f\ n = f\ n + 1 - f\ n$

and $0 = 1$

Towards a functional big-step semantics (4)

We can force our functions to terminate by inserting a clock.

ML: `| run_t s (For (e1, e2, t)) =`
 `...`
 `(Rval _, s3) => run_t s3 (For (e1, e2, t))`

HOL: `eval_t s (For e1 e2 t) =`
 `...`
 `(Rval _, s3) =>`
 `if s3.clock ≠ 0 then`
 `eval_t (dec_clock s3) (For e1 e2 t)`
 `else (Rtimeout, s3)`

in HOL, the state is a record containing a clock

which gets decremented

if it hits zero, then we abort with an *uncatchable* exception

Functional big-step semantics

Other parts do not require clock clutter:

HOL:

```
eval_t s (Exp e) = eval_e s e
eval_t s (Dec x t) = eval_t (st
eval_t s Break = (Rbreak,s)
eval_t s (Seq t1 t2) =
  case eval_t s t1 of
    (Rval _,s1) ⇒ eval_t s1 t2
  | r ⇒ r
eval_t s (If e t1 t2) =
  case eval_e s e of
    (Rval n1,s1) ⇒ eval_t s1 (if n1 = 0 then t2 else t1)
  | r ⇒ r
```

... since the size of
the input shrinks

Observational semantics

Compiler proof is to relate different *observational semantics*.

Without I/O, programs can only Terminate, Diverge or Crash.

terminates if there is a clock that is sufficient

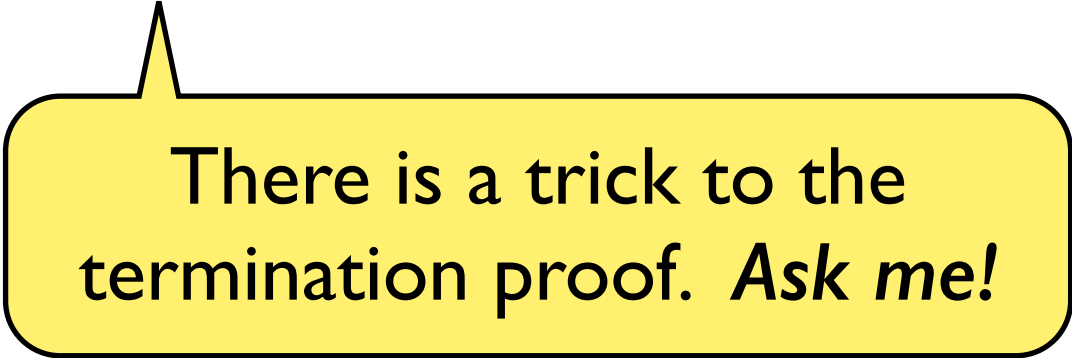
```
semantics  $t$  =  
if  $\exists c\ v\ s. \text{sem\_t } (s\_with\_clock\ c)\ t = (Rval\ v, s)$  then Terminate  
else if  $\forall c. \exists s. \text{sem\_t } (s\_with\_clock\ c)\ t = (Rtimeout, s)$  then Diverge  
else Crash
```

diverges if all clocks cause timeouts

crashes otherwise, e.g. when Rbreak propagates to the top

Exercise

Define a functional big-step semantics in your favourite prover (e.g. HOL4, Isabelle/HOL, Coq, ACL2)



There is a trick to the termination proof. *Ask me!*

Ingredient 3: Compiler function

Ingredient 3: Compiler function

Three phases:

phase 1: rewrite For to something simpler

phase 2: split expressions into instructions

phase 3: flatten to list of assembly instructions

Ingredient 3: Compiler function

Three phases:

phase I: rewrite For to something simpler

```
phase1 (Exp e) = Exp e
phase1 (Dec x t) = Seq (Exp (Assign x (Num 0))) (phase1 t)
phase1 Break = Break
phase1 (Seq t1 t2) = Seq (phase1 t1) (phase1 t2)
phase1 (If e t1 t2) = If e (phase1 t1) (phase1 t2)
phase1 (For e1 e2 t) = Loop (If e1 (Seq (phase1 t) (Exp e2)) Break)
```

where $\text{Loop } t = \text{For } (\text{Num } 1) (\text{Num } 1) t$

Makes all loops simple while-true loops.

Proving a compiler correct

Ingredients:

- a **formal logic** for the proofs
- **accurate models** of
 - the **source** language
 - the **target** language
 - the **compiler** algorithm

Tools:

- a **proof assistant** (software)

Method:

- prove simulation theorem using **induction**

Induction theorem

$\vdash \forall P.$

$(\forall s \ e. \ P \ s \ (\text{Exp } e)) \wedge$

$(\forall s \ x \ t. \ P \ (\text{store_var } x \ 0 \ s) \ t \Rightarrow P \ s \ (\text{Dec } x \ t)) \wedge$

$(\forall s. \ P \ s \ \text{Break}) \wedge$

$(\forall s \ t_1 \ t_2.$

$(\forall v_2 \ s_1 \ v_5.$

$(\text{eval_t } s \ t_1 = (v_2, s_1)) \wedge (v_2 = \text{Rval } v_5) \wedge$

$P \ s \ t_1 \Rightarrow$

$P \ s \ (\text{Seq } t_1 \ t_2)) \wedge$

$(\forall s \ e \ t_1 \ t_2.$

$(\forall v_2 \ s_1 \ n_1.$

$(\text{eval_e } s \ e = (v_2, s_1)) \wedge (v_2 = \text{Rval } n_1) \wedge$

$P \ s_1 \ (\text{if } n_1 = 0 \ \text{then } t_2 \ \text{else } t_1)) \Rightarrow$

$P \ s \ (\text{If } e \ t_1 \ t_2)) \wedge$

$(\forall s \ e_1 \ e_2 \ t.$

$(\forall v_4 \ s_1 \ n_1 \ v_3 \ s_2 \ v_7 \ v_2 \ s_3 \ v_5.$

$(\text{eval_e } s \ e_1 = (v_4, s_1)) \wedge (v_4 = \text{Rval } n_1) \wedge n_1 \neq 0 \wedge$

$(\text{eval_t } s_1 \ t = (v_3, s_2)) \wedge (v_3 = \text{Rval } v_7) \wedge$

$(\text{eval_e } s_2 \ e_2 = (v_2, s_3)) \wedge (v_2 = \text{Rval } v_5) \wedge s_3.\text{clock} \neq 0 \Rightarrow$

$P \ (\text{dec_clock } s_3) \ (\text{For } e_1 \ e_2 \ t)) \wedge$

$(\forall v_4 \ s_1 \ n_1.$

$(\text{eval_e } s \ e_1 = (v_4, s_1)) \wedge (v_4 = \text{Rval } n_1) \wedge n_1 \neq 0 \Rightarrow$

$P \ s_1 \ t) \Rightarrow$

$P \ s \ (\text{For } e_1 \ e_2 \ t)) \Rightarrow$

$\forall v \ v_1. \ P \ v \ v_1$

Has structure of functional
big-step interpreter.

One case for each program
construct.

Simulation proof

Particularly simple for phase I:

$$\vdash \forall s\ t. \text{eval_t } s\ (\text{phase1 } t) = \text{eval_t } s\ t$$

Requires only 8 lines of HOL4 proof script (mostly rewriting).

We can lift this to the observational semantics
with a one-line proof by rewriting:

$$\vdash \forall t. \text{semantics } (\text{phase1 } t) = \text{semantics } t$$

Simulation theorems *in general*

Usually, each compile phase requires a theorem of the form:

non-failing evaluation of source intermediate language (IL)

$\text{evaluate code } s1 = (res, s2) \wedge res \neq \mathbf{Rfail} \wedge$
 $\text{state_rel } s1 \ t1 \Rightarrow$

implies

$\exists t2. \text{evaluate (compile code) } t1 = (res, t2) \wedge$
 $\text{state_rel } s2 \ t2$

similar evaluation in target IL

... w.r.t. some relation between states (**state_rel**)

Simulation theorems *in general*

Usually, each compile phase requires a theorem of the form:

$\text{evaluate } code \ s1 = (res, s2) \wedge res \neq \mathbf{Rfail} \wedge$
 $\text{state_rel } s1 \ t1 \Rightarrow$

$\exists t2. \text{ evaluate } (\text{compile } code) \ t1 = (res, t2) \wedge$
 $\text{state_rel } s2 \ t2$

usually: state_rel keeps clocks in sync

variant: target can consume more clock ticks

Simulation theorems *in general*

Usually, each compile phase requires a theorem of the form:

$$\begin{aligned} &\text{evaluate } code \ s1 = (res, s2) \wedge res \neq \mathbf{Rfail} \wedge \\ &\text{state_rel } s1 \ t1 \Rightarrow \\ &\exists t2. \text{ evaluate } (\text{compile } code) \ t1 = (res, t2) \wedge \\ &\quad \text{state_rel } s2 \ t2 \end{aligned}$$

Sufficient to prove observational equivalence for both terminating and diverging runs.

Phase 2 simulation theorem

source IL

not failing

restrict language syntax
following phase 1

$\vdash \forall s \ e \ n \ res \ s_1.$

$(\text{eval_t } s \ e = (res, s_1)) \wedge res \neq \text{Rfail} \wedge \text{phase2_subset } e \wedge$
 $\text{possible_var_name } n \ s_1.\text{store} \wedge s_1.\text{clock} = 0 \wedge s_1.\text{store} \sqsubseteq t.\text{store} \wedge$
 $(t.\text{clock} = 0 \wedge \text{strlen } n \Rightarrow$

target IL

compiler

$\exists t_1.$

state relation

$(\text{eval_t } t \ (\text{flatten_t } e \ n) = (res, t_1)) \wedge$
 $s_1.\text{store} \sqsubseteq t_1.\text{store} \wedge (t_1.\text{clock} = s_1.\text{clock}) \wedge$
 $\text{possible_var_name } n \ s_1.\text{store} \wedge$

$\forall k \ v.$

$\text{possible_var_name } k \ s.\text{store} \wedge t_max \ e < \text{strlen } k \wedge$
 $\text{strlen } k < \text{strlen } n \wedge (\text{lookup } t.\text{store } k = \text{SOME } v) \Rightarrow$
 $(\text{lookup } t_1.\text{store } k = \text{SOME } v)$

extra property

Composing top-level theorems

Each phase maintains observational equivalence:

$$\text{semantics } (\text{phase1 } t) = \text{semantics } t$$

$$\text{semantics } t \neq \text{Crash} \wedge \text{phase2_subset } t \Rightarrow \\ (\text{semantics } (\text{phase2 } t) = \text{semantics } t)$$

$$\text{semantics } t \neq \text{Crash} \wedge \text{phase3_subset } t \Rightarrow \\ (\text{asm_semantics } (\text{phase3 } 0 \ 0 \ t) = \text{semantics } t)$$

observational semantics of target assembly

Here: $\text{compile } t = \text{phase3 } 0 \ 0 \ (\text{phase2 } (\text{phase1 } t))$

Composing top-level theorems

Result:

lemma: correct syntax implies no Crashes

for ML: type-correct program implies no Crash

$\vdash \forall t. \text{syntax_ok } t \Rightarrow (\text{asm_semantics } (\text{compile } t) = \text{semantics } t)$

where $\text{compile } t = \text{phase3 } 0 \ 0 \ (\text{phase2 } (\text{phase1 } t))$

What we learnt

Ingredients: formal logic, compiler, language semantics

Tools: proof assistant

Method: *using functional big-step semantics* it suffices to prove theorems of the form:

$$\text{evaluate } \text{code } s1 = (res, s2) \wedge res \neq \text{Rfail} \wedge \\ \text{state_rel } s1 \ t1 \Rightarrow$$
$$\exists t2. \text{evaluate } (\text{compile } \text{code}) \ t1 = (res, t2) \wedge \\ \text{state_rel } s2 \ t2$$

in order to prove observational equivalence, i.e.

$$\vdash \forall t. \text{syntax_ok } t \Rightarrow (\text{asm_semantics } (\text{compile } t) = \text{semantics } t)$$

Extra slides

Comparing functional with relation big-step

In the functional version, Seq was specified by:

```
eval_t s (Seq t1 t2) =  
  case eval_t s t1 of  
    (Rval _, s1) => eval_t s1 t2  
  | r => r
```

In the relational version, Seq is specified using *four rules*:

$$(S1) \frac{(t_1, s) \Downarrow_t (\text{Rval } n_1, s_1) \quad (t_2, s_1) \Downarrow_t r}{(\text{Seq } t_1 \ t_2, s) \Downarrow_t r}$$

$$(S2) \frac{(t_1, s) \Downarrow_t (r, s_1) \quad \neg \text{is_Rval } r}{(\text{Seq } t_1 \ t_2, s) \Downarrow_t (r, s_1)}$$

$$(S1') \frac{(t_1, s) \Uparrow_t}{(\text{Seq } t_1 \ t_2, s) \Uparrow_t}$$

$$(S2') \frac{(t_1, s) \Downarrow_t (\text{Rval } n_1, s_1) \quad (t_2, s_1) \Uparrow_t}{(\text{Seq } t_1 \ t_2, s) \Uparrow_t}$$

Induction from relational big-step

$$\begin{aligned}
 &\vdash \dots \wedge \dots \wedge \dots \wedge \dots \wedge \dots \wedge \dots \wedge \dots \wedge \dots \wedge \dots \wedge \\
 &(\forall s \ s_1 \ e_1 \ e_2 \ t. \\
 &\quad (e_1, s) \Downarrow_e (\text{Rval } 0, s_1) \Rightarrow P (\text{For } e_1 \ e_2 \ t, s) (\text{Rval } 0, s_1)) \wedge \\
 &(\forall s \ s_1 \ e_1 \ e_2 \ t \ r. \\
 &\quad (e_1, s) \Downarrow_e (r, s_1) \wedge \neg \text{is_Rval } r \Rightarrow P (\text{For } e_1 \ e_2 \ t, s) (r, s_1)) \wedge \\
 &(\forall s \ s_1 \ s_2 \ s_3 \ e_1 \ e_2 \ t \ n_1 \ n_2 \ n_3 \ r. \\
 &\quad (e_1, s) \Downarrow_e (\text{Rval } n_1, s_1) \wedge n_1 \neq 0 \wedge P (t, s_1) (\text{Rval } n_2, s_2) \wedge \\
 &\quad (e_2, s_2) \Downarrow_e (\text{Rval } n_3, s_3) \wedge P (\text{For } e_1 \ e_2 \ t, s_3) r \Rightarrow \\
 &\quad P (\text{For } e_1 \ e_2 \ t, s) r) \wedge \\
 &(\forall s \ s_1 \ s_2 \ e_1 \ e_2 \ t \ n_1. \\
 &\quad (e_1, s) \Downarrow_e (\text{Rval } n_1, s_1) \wedge n_1 \neq 0 \wedge P (t, s_1) (\text{Rbreak}, s_2) \Rightarrow \\
 &\quad P (\text{For } e_1 \ e_2 \ t, s) (\text{Rval } 0, s_2)) \wedge \\
 &(\forall s \ s_1 \ s_2 \ s_3 \ e_1 \ e_2 \ t \ n_1 \ n_2 \ r. \\
 &\quad (e_1, s) \Downarrow_e (\text{Rval } n_1, s_1) \wedge n_1 \neq 0 \wedge P (t, s_1) (\text{Rval } n_2, s_2) \wedge \\
 &\quad (e_2, s_2) \Downarrow_e (r, s_3) \wedge \neg \text{is_Rval } r \Rightarrow \\
 &\quad P (\text{For } e_1 \ e_2 \ t, s) (r, s_3)) \wedge \\
 &(\forall s \ s_1 \ s_2 \ e_1 \ e_2 \ t \ n_1 \ r. \\
 &\quad (e_1, s) \Downarrow_e (\text{Rval } n_1, s_1) \wedge n_1 \neq 0 \wedge P (t, s_1) (r, s_2) \wedge \\
 &\quad r \neq \text{Rbreak} \Rightarrow \\
 &\quad P (\text{For } e_1 \ e_2 \ t, s) (r, s_2)) \Rightarrow \\
 &\forall ts \ rs. \ ts \Downarrow_t rs \Rightarrow P \ ts \ rs
 \end{aligned}$$

It has one rule for each case in the relation

Six cases for For!

Observational semantics with I/O

Defining the observational semantics when there is I/O.

$$\begin{aligned} \text{semantics } t \text{ input } (\text{Terminate } io_trace) &\iff \\ \exists c \ nd \ i \ s. & \\ \quad (\text{sem_t } (\text{init_st } c \ nd \ \text{input}) \ t = (\text{Rval } i, s)) \ \wedge & \\ \quad (\text{FILTER ISL } s.io_trace = io_trace) & \\ \text{semantics } t \text{ input } \text{Crash} &\iff \\ \exists c \ nd \ r \ s. & \\ \quad (\text{sem_t } (\text{init_st } c \ nd \ \text{input}) \ t = (r, s)) \ \wedge & \\ \quad ((r = \text{Rbreak}) \vee (r = \text{Rfail})) & \\ \text{semantics } t \text{ input } (\text{Diverge } io_trace) &\iff \\ \exists nd. & \\ \quad (\forall c. \ \exists s. \ \text{sem_t } (\text{init_st } c \ nd \ \text{input}) \ t = (\text{Rtimeout}, s)) \ \wedge & \\ \quad (io_trace = & \\ \quad \vee c. & \\ \quad \quad \text{fromList} & \\ \quad \quad (\text{FILTER ISL } (\text{SND } (\text{sem_t } (\text{init_st } c \ nd \ \text{input}) \ t)).io_trace)) & \end{aligned}$$