



COMP4161: Advanced Topics in Software Verification

$P \parallel Q$

Gerwin Klein, June Andronick, Ramana Kumar, Miki Tanaka
S2/2017

data61.csiro.au



Content



- Intro & motivation, getting started [1]

- Foundations & Principles
 - Lambda Calculus, natural deduction [1,2]
 - Higher Order Logic [3^a]
 - Term rewriting [4]

- Proof & Specification Techniques
 - Inductively defined sets, rule induction [5]
 - Datatypes, recursion, induction [6, 7]
 - Hoare logic, proofs about programs, invariants [8^b, 9]
 - (mid-semester break)
 - C verification [10]
 - CakeML, Isar [11^c]
 - Concurrency [12]

^aa1 due; ^ba2 due; ^ca3 due

Program verification so far



If the following true?

$\{x = 0\}$

$y := x;$

$x := x + 1;$

$\{x = 1 \wedge y = 0\}$

YES!

Program verification with concurrency



Is it still true?

$$\begin{array}{l} \{x = 0\} \\ y := x; \\ x := x + 1; \\ \{x = 1 \wedge y = 0\} \end{array} \quad || \quad x := 4$$

NO!

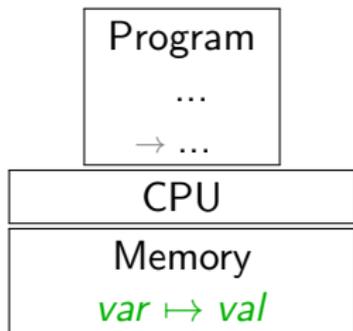
Program verification so far



So far we have assumed **sequential execution**

$\{x = 0\}$	$x \mapsto 0$	$y \mapsto -$
$y := x;$	$x \mapsto 0$	$y \mapsto 0$
$x := x + 1;$	$x \mapsto 1$	$y \mapsto 0$
$\{x = 1 \wedge y = 0\}$		

i.e. a single thread of execution accessing the memory state

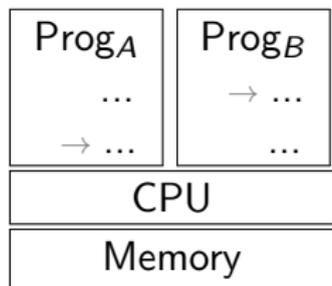


This is not always the case!

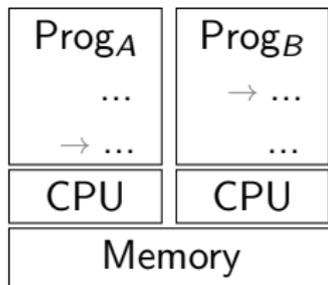
Types of concurrency



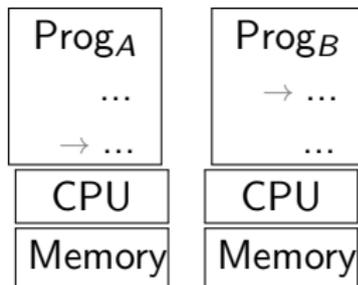
Multithreading



Multicore



Distributed



All need communication and synchronisation mechanisms

Shared memory
Interleaved execution

Shared memory
Parallel execution

Message passing

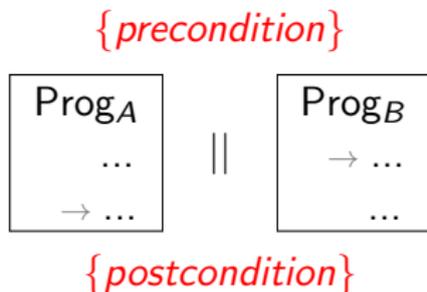
Here: we'll look at shared-memory concurrency

(and we'll ignore further complications such as caches, weak memory ...)

Goal



We want to be able to reason about parallel composition of programs:



2 kinds of properties:

Safety:

“something bad does not happen”
(no bad state can be reached)
e.g. $\{x = 0\}$

Liveness:

“something good must happen”
(specific states must be reached)
e.g. the program terminates

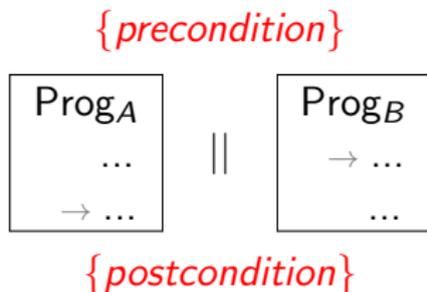
With concurrency: much harder!

With concurrency: new problems!

Goal



We want to be able to reason about parallel composition of programs:



Here:

- We focus on **safety** properties: postcondition holds **if reached**
- We will define parallel composition ($||$) as non-deterministic interleaving
- We go back to our minimal IMP language (forget about C and monads)

`datatype com = SKIP`

Program verification so far



If the following true?

$\{x = 0\}$

$y := x;$

$x := x + 1;$

$\{x = 1 \wedge y = 0\}$

YES!

Program verification with concurrency



Is it still true?

```
{x = 0}
y := x;           ||   x := 4
x := x + 1;
{x = 1 ∧ y = 0}
```

NO!

What is going wrong?

What do we need to change?

- to make sure we don't prove wrong statements!
- to allow us to prove true statements about concurrent programs

Program verification so far



How would we have proved this?

$$\{x = 0\} \implies \{x + 1 = 1 \wedge x = 0\}$$
$$y := x; \{x + 1 = 1 \wedge y = 0\}$$
$$x := x + 1;$$
$$\{x = 1 \wedge y = 0\}$$

Using Hoare logic rules!

$$\frac{\vdash \{P\} c_1 \{R\} \quad \vdash \{R\} c_2 \{Q\}}{\vdash \{P\} c_1; c_2 \{Q\}}$$
$$\frac{}{\vdash \{P[x \mapsto e]\} \quad x := e \quad \{P\}}$$

Why does this make it true? What does it mean that it's true?

It means:

If the program “ $y := x; x := x + 1$ ” is executed from a state satisfying $\{x = 0\}$ then, if it terminates, the resulting state satisfied $\{x = 1 \wedge y = 0\}$

Program verification so far



How would we have proved this?

$\{x = 0\} \implies \{x + 1 = 1 \wedge x = 0\}$
 $y := x; \{x + 1 = 1 \wedge y = 0\}$
 $x := x + 1;$
 $\{x = 1 \wedge y = 0\}$

Using Hoare logic rules!

$$\frac{\vdash \{P\} c_1 \{R\} \quad \vdash \{R\} c_2 \{Q\}}{\vdash \{P\} c_1; c_2 \{Q\}}$$
$$\frac{}{\vdash \{P[x \mapsto e]\} x := e \{P\}}$$

Why does this make it true? What does it mean that it's true?

It means:

$$\langle y := x; x := x + 1, \sigma \rangle \rightarrow \sigma' \wedge x \sigma = 0 \longrightarrow x \sigma' = 1 \wedge y \sigma' = 0$$

Where:

$$\frac{\langle c_1, \sigma \rangle \rightarrow \sigma' \quad \langle c_2, \sigma' \rangle \rightarrow \sigma''}{\langle c_1; c_2, \sigma \rangle \rightarrow \sigma''} \quad \frac{e \sigma = v}{\langle x := e, \sigma \rangle \rightarrow \sigma[x \mapsto v]}$$

Program verification with concurrency



$\{x = 0\}$
 ~~$y := x; \{x + 1 = 1 \wedge y = 0\}$~~ || $x := 4$
 $x := x + 1;$
 $\{x = 1 \wedge y = 0\}$

- Execution is interleaved
- Intermediate assertions can be interfered with
- Need a new reasoning framework!
- New syntax, new semantics,
new proof rules (proved sound w.r.t semantics), new VCG
- (1969: Hoare Logic (Tony Hoare))
- 1976: Owicki-Gries (Susan Owicki and David Gries)
- 1981: Rely-Guarantee (Cliff Jones)

Owicki-Gries framework



Intuition:

- Syntax: our IMP language + Parallel operator + Await operator
- Semantics:
 - ▶ $P \parallel Q$: pick one program and execute its current instruction
 - ▶ $AWAIT\ b\ DO\ c\ OD$: if guard is true execute c atomically
- Proof rules:
 - ▶ you prove *local correctness* (as before)
 - ▶ you prove *interference-freedom* (assertions not interfered with)

$\{is_even\ x\}$

$x := x + 1; \{is_even\ x + 1\} \parallel x := x + 2$

$x := x + 1;$

$\{is_even\ x\}$

→ Needs a fully annotated program!

→ Needs a “small-step semantics” $\langle c, \sigma \rangle \rightarrow \langle c', \sigma' \rangle$

Owicki-Gries framework



Formally:

- Syntax: our IMP language + Parallel operator + Await operator
- Semantics:

$$\frac{\langle c_1, \sigma \rangle \rightarrow \langle c'_1, \sigma' \rangle}{\langle c_1 || c_2, \sigma \rangle \rightarrow \langle c'_1 || c_2, \sigma' \rangle} \quad \frac{\langle c_2, \sigma \rangle \rightarrow \langle c'_2, \sigma' \rangle}{\langle c_1 || c_2, \sigma \rangle \rightarrow \langle c_1 || c'_2, \sigma' \rangle}$$

- Hoare rules:

$$\frac{\{P_1\} c_1 \{Q_1\} \quad \{P_2\} c_2 \{Q_2\} \quad \textit{interfree} c_1 c_2 \quad \textit{interfree} c_2 c_1}{\{P_1 \wedge P_2\} c_1 || c_2 \{Q_1 \wedge Q_2\}}$$

Where

$\textit{interfree} c_1 c_2 \equiv$

$\forall p \in (\textit{assertions } c_1) \quad \forall (a \ c) \in (\textit{atomics } c_2) \quad \{p \wedge a\}c\{p\}$

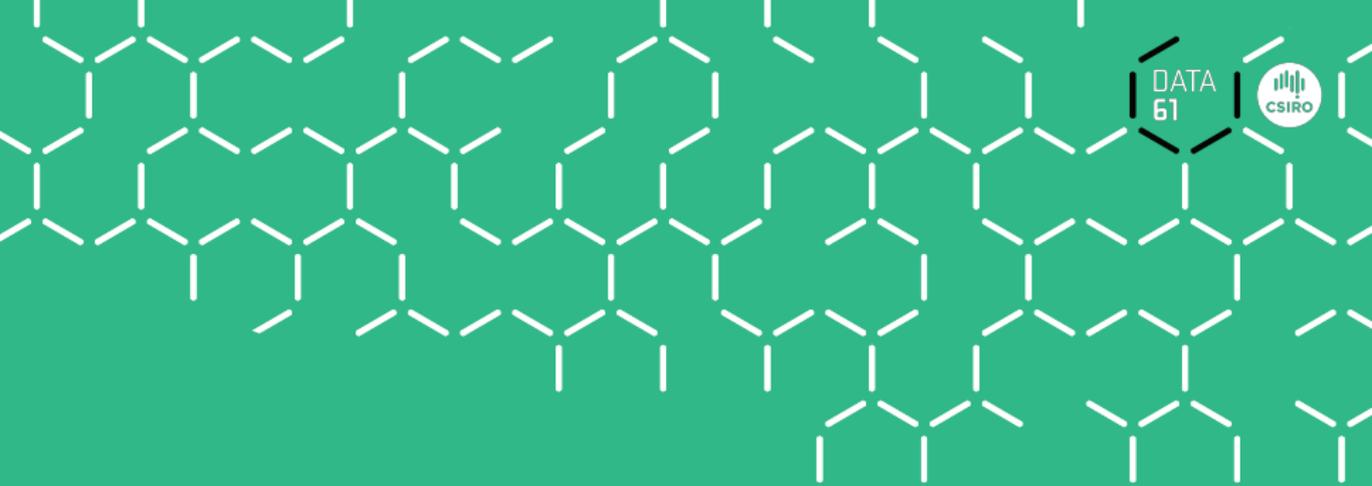
Owicki-Gries framework



- Quadratic explosion of proof obligations! (verification conditions)
- Not compositional
- Not complete: sometimes need auxiliary/ghost variables

$$\begin{array}{c} \{x = 0\} \\ x := x + 1; \parallel x := x + 1 \\ \{x = 2\} \end{array}$$

$$\begin{array}{c} \{x = 0 \wedge a_1 = 0 \wedge a_2 = 0\} \\ \{a_2 = 0 \wedge x = 0 \vee a_2 = 1 \wedge x = 1\} \parallel \{a_1 = 0 \wedge x = 0 \vee a_1 = 1 \wedge x = 1\} \\ \langle x := x + 1; a_1 := 1 \rangle \parallel \langle x := x + 1; a_2 := 1 \rangle \\ \{a_2 = 0 \wedge x = 1 \vee a_2 = 1 \wedge x = 2\} \parallel \{a_1 = 1 \wedge x = 1 \vee a_1 = 1 \wedge x = 2\} \\ \{x = 2\} \end{array}$$

A background pattern of white hexagons on a teal background, arranged in a staggered grid.

DATA
61



Demo

Rely-Guarantee?



Intuition:

- Syntax, semantics: as before (but no need for assertions)
- Proof rules:
 - ▶ each program is specified in isolation, **assuming a behavior of the “environment”** (other programs in parallel)
 - ▶ each program has: precondition, postcondition, **rely and guarantee**
 - ▶ rely and guarantee are **relations between 2 states**
 - ▶ **rely** expresses the maximum behavior of the environment (the interference that the program can tolerate)
 - ▶ **guarantee** expresses a maximum behavior promised to the environment

$$\begin{array}{c} c \\ \{P, R, G, Q\} \end{array} \parallel \begin{array}{c} c' \\ \{P', R', G', Q'\} \end{array}$$

$$\sigma_0 \xrightarrow{c} \sigma_1 \xrightarrow{c} \sigma_2 \xrightarrow{c'} \sigma_3 \xrightarrow{c} \sigma_4 \xrightarrow{c'} \sigma_5 \xrightarrow{c'} \sigma_6 \xrightarrow{c} \sigma_7$$

Rely-Guarantee?



Formally:

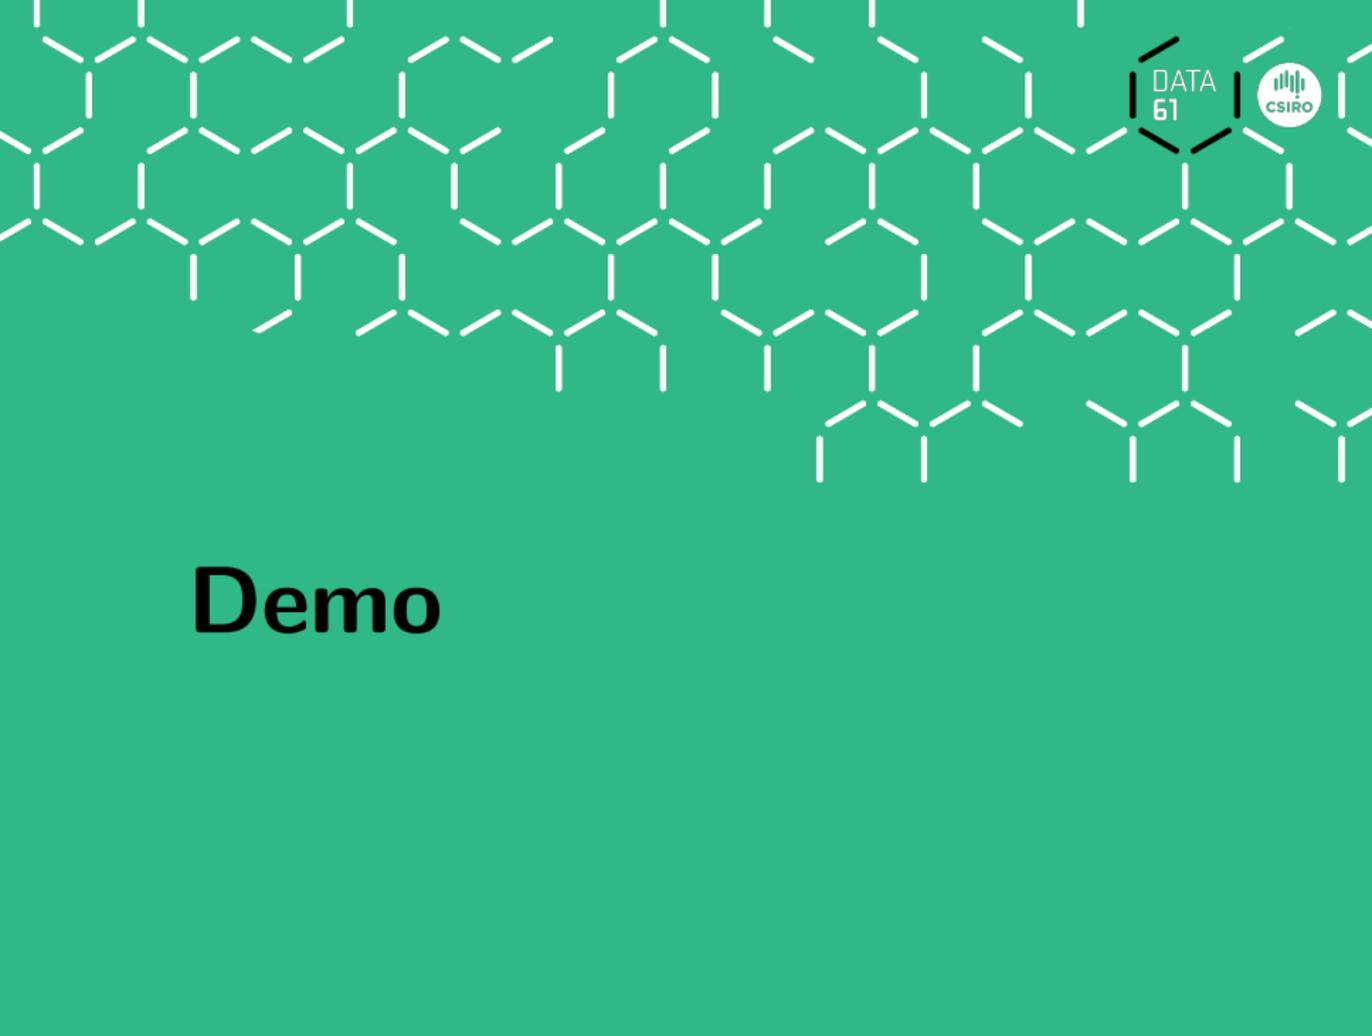
- Syntax, semantics: as before (but no need for assertions)
- Proof rules (examples):

$$\frac{P \subseteq \{s. f \ s \in Q\} \quad \{(s, t). P \ s \wedge (t = f \ s \vee t = s)\} \subseteq G \quad \text{stable } P \ R \quad \text{stable } Q}{\text{Basic } f\{P, R, G, Q\}}$$

$$\frac{c_1\{P_1, R_1, G_1, Q_1\} \quad c_2\{P_2, R_2, G_2, Q_2\} \quad G_1 \subseteq R_2 \quad G_2 \subseteq R_1}{c_1 || c_2\{P_1 \cap P_2, R_1 \cap R_2, G_1 \cup G_2, Q_1 \cap Q_2\}}$$

Where $\text{stable } P \ R = \forall \sigma \ \sigma'. (P\sigma \wedge R(\sigma, \sigma')) \rightarrow P\sigma'$
(doing an environment step before or after P should not make P invalid)

Intuition: the guarantee of one program is the rely of the other program

A background pattern of white hexagons on a dark teal background, arranged in a staggered grid.

DATA
61



Demo

We have seen today ...



- Need for new reasoning framework for parallel/concurrent programs
- Owicki-Gries
- Rely-Guarantee