

*Dialectical refinement:
Rescuing programming from the logicians*

Richard Bornat
Professor of Computer Programming
School of Engineering and Information Sciences
Middlesex University

5th November 2008

There are two programming problems

There are two programming problems

- ▶ The novice programming problem

There are two programming problems

- ▶ The novice programming problem
- ▶ The expert programming problem

There are two programming problems

- ▶ The novice programming problem
- ▶ The expert programming problem

Both are interesting.

There are two programming problems

- ▶ The novice programming problem
- ▶ The expert programming problem

Both are interesting. This talk is about the second problem.

There are two programming problems

- ▶ The novice programming problem
- ▶ The expert programming problem

Both are interesting. This talk is about the second problem.

Most programming ‘languages’, *especially* including Java, address the expert programming problem.

Programming as a creative act

Programming as a creative act

I take 'programming' to be the whole business: conception, design, construction, verification, testing . . .

Programming as a creative act

I take 'programming' to be the whole business: conception, design, construction, verification, testing . . .

I'm very interested in 'correct' programs: ones with specification and verification.

Programming as a creative act

I take ‘programming’ to be the whole business: conception, design, construction, verification, testing . . .

I’m very interested in ‘correct’ programs: ones with specification and verification.

And I’m interested in *inventing* correct programs.

Programming as a creative act

I take ‘programming’ to be the whole business: conception, design, construction, verification, testing . . .

I’m very interested in ‘correct’ programs: ones with specification and verification.

And I’m interested in *inventing* correct programs.

(I claim that) invention always involves trial and error.

Programming as a creative act

I take ‘programming’ to be the whole business: conception, design, construction, verification, testing . . .

I’m very interested in ‘correct’ programs: ones with specification and verification.

And I’m interested in *inventing* correct programs.

(I claim that) invention always involves trial and error.

Can trial and error be ‘logical’? Or is it illogical guesswork?

Two ways to program

Two ways to program

I can invent a program, in a blinding flash of creativity

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it (*of course* I wouldn't test it ...)

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it (*of course* I wouldn't test it ...). Logicians call this *ad-hoc programming*.

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it (*of course* I wouldn't test it ...). Logicians call this *ad-hoc programming*.

Before the verification I have to think up a logical specification. I don't need the specification to start programming.

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it (*of course* I wouldn't test it ...). Logicians call this *ad-hoc programming*.

Before the verification I have to think up a logical specification. I don't need the specification to start programming. Or do I? ...

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it (*of course* I wouldn't test it ...). Logicians call this *ad-hoc programming*.

Before the verification I have to think up a logical specification. I don't need the specification to start programming. Or do I? ...

Logicians would like us to start with the specification

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it (*of course* I wouldn't test it ...). Logicians call this *ad-hoc programming*.

Before the verification I have to think up a logical specification. I don't need the specification to start programming. Or do I? ...

Logicians would like us to start with the specification; write a magical (unimplementable) program which satisfies the specification

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it (*of course* I wouldn't test it ...). Logicians call this *ad-hoc programming*.

Before the verification I have to think up a logical specification. I don't need the specification to start programming. Or do I? ...

Logicians would like us to start with the specification; write a magical (unimplementable) program which satisfies the specification; then *refine* the magical program step-by-step towards something that can be used

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it (*of course* I wouldn't test it ...). Logicians call this *ad-hoc programming*.

Before the verification I have to think up a logical specification. I don't need the specification to start programming. Or do I? ...

Logicians would like us to start with the specification; write a magical (unimplementable) program which satisfies the specification; then *refine* the magical program step-by-step towards something that can be used; using refinement steps which preserve the property that the program satisfies the specification.

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it (*of course* I wouldn't test it ...). Logicians call this *ad-hoc programming*.

Before the verification I have to think up a logical specification. I don't need the specification to start programming. Or do I? ...

Logicians would like us to start with the specification; write a magical (unimplementable) program which satisfies the specification; then *refine* the magical program step-by-step towards something that can be used; using refinement steps which preserve the property that the program satisfies the specification.

The logicians' way of programming produces a result that is *correct by construction*.

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it (*of course* I wouldn't test it ...). Logicians call this *ad-hoc programming*.

Before the verification I have to think up a logical specification. I don't need the specification to start programming. Or do I? ...

Logicians would like us to start with the specification; write a magical (unimplementable) program which satisfies the specification; then *refine* the magical program step-by-step towards something that can be used; using refinement steps which preserve the property that the program satisfies the specification.

The logicians' way of programming produces a result that is *correct by construction*.

Ad-hoc programs are rarely specified, even more rarely verified.

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it (*of course* I wouldn't test it ...). Logicians call this *ad-hoc programming*.

Before the verification I have to think up a logical specification. I don't need the specification to start programming. Or do I? ...

Logicians would like us to start with the specification; write a magical (unimplementable) program which satisfies the specification; then *refine* the magical program step-by-step towards something that can be used; using refinement steps which preserve the property that the program satisfies the specification.

The logicians' way of programming produces a result that is *correct by construction*.

Ad-hoc programs are rarely specified, even more rarely verified.

Can we defend program invention?

Two ways to program

I can invent a program, in a blinding flash of creativity; and then verify it (*of course* I wouldn't test it ...). Logicians call this *ad-hoc programming*.

Before the verification I have to think up a logical specification. I don't need the specification to start programming. Or do I? ...

Logicians would like us to start with the specification; write a magical (unimplementable) program which satisfies the specification; then *refine* the magical program step-by-step towards something that can be used; using refinement steps which preserve the property that the program satisfies the specification.

The logicians' way of programming produces a result that is *correct by construction*.

Ad-hoc programs are rarely specified, even more rarely verified.

Can we defend program invention? Can we give it a nicer name?

Dialectical programming

Dialectical programming

Lakatos, in “Proofs and Refutations”, shows the evolution over time of Euler’s conjecture about polyhedral solids

Vertices – *Edges* + *Faces* = 2 (e.g for a cube $V = 8, E = 12, F = 6$).

Dialectical programming

Lakatos, in “Proofs and Refutations”, shows the evolution over time of Euler’s conjecture about polyhedral solids

Vertices – *Edges* + *Faces* = 2 (e.g for a cube $V = 8, E = 12, F = 6$).

The *conjecture* leads to a *proof*

Dialectical programming

Lakatos, in “Proofs and Refutations”, shows the evolution over time of Euler’s conjecture about polyhedral solids

Vertices – *Edges* + *Faces* = 2 (e.g for a cube $V = 8, E = 12, F = 6$).

The *conjecture* leads to a *proof* which is challenged by a *counter-example* (e.g. a polyhedron with a tunnel, a polyhedron with a central cavity)

Dialectical programming

Lakatos, in “Proofs and Refutations”, shows the evolution over time of Euler’s conjecture about polyhedral solids

Vertices – *Edges* + *Faces* = 2 (e.g for a cube $V = 8, E = 12, F = 6$).

The *conjecture* leads to a *proof* which is challenged by a *counter-example* (e.g. a polyhedron with a tunnel, a polyhedron with a central cavity) which leads to a refinement of the conjecture, a new proof, a deeper understanding.

Dialectical programming

Lakatos, in “Proofs and Refutations”, shows the evolution over time of Euler’s conjecture about polyhedral solids

Vertices – *Edges* + *Faces* = 2 (e.g for a cube $V = 8, E = 12, F = 6$).

The *conjecture* leads to a *proof* which is challenged by a *counter-example* (e.g. a polyhedron with a tunnel, a polyhedron with a central cavity) which leads to a refinement of the conjecture, a new proof, a deeper understanding.

This is intentionally reminiscent of Hegel’s dialectic “Thesis plus antithesis yields synthesis”.

Dialectical programming

Lakatos, in “Proofs and Refutations”, shows the evolution over time of Euler’s conjecture about polyhedral solids

Vertices – *Edges* + *Faces* = 2 (e.g for a cube $V = 8, E = 12, F = 6$).

The *conjecture* leads to a *proof* which is challenged by a *counter-example* (e.g. a polyhedron with a tunnel, a polyhedron with a central cavity) which leads to a refinement of the conjecture, a new proof, a deeper understanding.

This is intentionally reminiscent of Hegel’s dialectic “Thesis plus antithesis yields synthesis”. (But we are simple people, we won’t bother with Hegel.)

Dialectical programming

Lakatos, in “Proofs and Refutations”, shows the evolution over time of Euler’s conjecture about polyhedral solids

$Vertices - Edges + Faces = 2$ (e.g. for a cube $V = 8, E = 12, F = 6$).

The *conjecture* leads to a *proof* which is challenged by a *counter-example* (e.g. a polyhedron with a tunnel, a polyhedron with a central cavity) which leads to a refinement of the conjecture, a new proof, a deeper understanding.

This is intentionally reminiscent of Hegel’s dialectic “Thesis plus antithesis yields synthesis”. (But we are simple people, we won’t bother with Hegel.)

Can the dialectic reach places that logical refinement cannot?

Dialectical programming

Lakatos, in “Proofs and Refutations”, shows the evolution over time of Euler’s conjecture about polyhedral solids

$Vertices - Edges + Faces = 2$ (e.g. for a cube $V = 8, E = 12, F = 6$).

The *conjecture* leads to a *proof* which is challenged by a *counter-example* (e.g. a polyhedron with a tunnel, a polyhedron with a central cavity) which leads to a refinement of the conjecture, a new proof, a deeper understanding.

This is intentionally reminiscent of Hegel’s dialectic “Thesis plus antithesis yields synthesis”. (But we are simple people, we won’t bother with Hegel.)

Can the dialectic reach places that logical refinement cannot? **I think so.**

The single-place buffer

The single-place buffer

In an aeroplane you have *sensors* – e.g. temperature, pressure, airflow – and you have *monitors* – e.g. black-box recorders, autopilots, display screens.

The single-place buffer

In an aeroplane you have *sensors* – e.g. temperature, pressure, airflow – and you have *monitors* – e.g. black-box recorders, autopilots, display screens.

We want a *simple* interface between a monitor and a sensor.

The single-place buffer

In an aeroplane you have *sensors* – e.g. temperature, pressure, airflow – and you have *monitors* – e.g. black-box recorders, autopilots, display screens.

We want a *simple* interface between a monitor and a sensor.

The simplest is a single-place buffer: the sensor writes to it, the monitor reads.

The single-place buffer

In an aeroplane you have *sensors* – e.g. temperature, pressure, airflow – and you have *monitors* – e.g. black-box recorders, autopilots, display screens.

We want a *simple* interface between a monitor and a sensor.

The simplest is a single-place buffer: the sensor writes to it, the monitor reads.

The sensor may write faster than the monitor reads (so there may be missed values).

The single-place buffer

In an aeroplane you have *sensors* – e.g. temperature, pressure, airflow – and you have *monitors* – e.g. black-box recorders, autopilots, display screens.

We want a *simple* interface between a monitor and a sensor.

The simplest is a single-place buffer: the sensor writes to it, the monitor reads.

The sensor may write faster than the monitor reads (so there may be missed values).

The monitor may read faster than the sensor writes (so there may be repeated values).

The single-place buffer

In an aeroplane you have *sensors* – e.g. temperature, pressure, airflow – and you have *monitors* – e.g. black-box recorders, autopilots, display screens.

We want a *simple* interface between a monitor and a sensor.

The simplest is a single-place buffer: the sensor writes to it, the monitor reads.

The sensor may write faster than the monitor reads (so there may be missed values).

The monitor may read faster than the sensor writes (so there may be repeated values).

But the monitor *must* always get complete values, not half-written values, not half of one value and half of the next.

How it can go wrong

How it can go wrong

The sensor is a clock, displaying minutes m and seconds s in two shared integer variables.

How it can go wrong

The sensor is a clock, displaying minutes m and seconds s in two shared integer variables.

Each second it does

```
if  $s = 59$  then  $s := 0; m := m + 1$  else  $s := s + 1$  fi
```

How it can go wrong

The sensor is a clock, displaying minutes m and seconds s in two shared integer variables.

Each second it does

```
if  $s = 59$  then  $s := 0; m := m + 1$  else  $s := s + 1$  fi
```

Between $s := 0$ and $m := m + 1$ the clock is slow (by 60 seconds). If the monitor reads at that instant it sees a half-written value.

How it can go wrong

The sensor is a clock, displaying minutes m and seconds s in two shared integer variables.

Each second it does

```
if  $s = 59$  then  $s := 0; m := m + 1$  else  $s := s + 1$  fi
```

Between $s := 0$ and $m := m + 1$ the clock is slow (by 60 seconds). If the monitor reads at that instant it sees a half-written value.

If the sensor does $m := m + 1; s := 0$ then the clock can be seen as fast (by 59 seconds).

Dijkstra's solution: atomic accesses

Dijkstra's solution: atomic accesses

local $b = \text{null}$ in

$$\left(\begin{array}{l} \dots \\ \text{write}(w) \hat{=} \langle b := w \rangle \end{array} \parallel \begin{array}{l} \dots \\ \text{read}() \hat{=} \text{local } y \text{ in} \\ \quad \langle y := b \rangle; \\ \quad \text{return } y \\ \text{ni} \end{array} \right)$$

Dijkstra's solution: atomic accesses

local $b = \text{null}$ in

$$\left(\begin{array}{l} \dots \\ \text{write}(w) \hat{=} \langle b := w \rangle \end{array} \parallel \begin{array}{l} \dots \\ \text{read}() \hat{=} \text{local } y \text{ in} \\ \quad \langle y := b \rangle; \\ \quad \text{return } y \\ \text{ni} \end{array} \right)$$

Atomicity means indivisibility. Interleaving will do.

Dijkstra's solution: atomic accesses

local $b = \text{null}$ in

$$\left(\begin{array}{l} \dots \\ \text{write}(w) \hat{=} \langle b := w \rangle \end{array} \parallel \begin{array}{l} \dots \\ \text{read}() \hat{=} \text{local } y \text{ in} \\ \quad \langle y := b \rangle; \\ \quad \text{return } y \\ \text{ni} \end{array} \right)$$

Atomicity means indivisibility. Interleaving will do.

Atomicity can be done with 'semaphores' as on the railways (block signalling).

Dijkstra's solution: atomic accesses

local $b = \text{null}$, $ws = \langle \rangle.\text{null}$, $rs = \langle \rangle$ in

$$\left(\begin{array}{l} \dots \\ \text{write}(w) \hat{=} \\ \quad \langle b := w; ws := ws.w \rangle \end{array} \parallel \begin{array}{l} \dots \\ \text{read}() \hat{=} \\ \quad \text{local } y \text{ in} \\ \quad \quad \langle y := b; rs := rs.y \rangle; \\ \quad \text{return } y \\ \text{ni} \end{array} \right)$$

Atomicity means indivisibility. Interleaving will do.

Atomicity can be done with 'semaphores' as on the railways (block signalling).

For specification purposes I add a couple of auxiliary variables ws and rs .

Dijkstra's solution: atomic accesses

local $b = \text{null}$, $ws = \langle \rangle.\text{null}$, $rs = \langle \rangle$ in

$$\left(\begin{array}{l} \dots \\ \text{write}(w) \hat{=} \\ \quad \langle b := w; ws := ws.w \rangle \end{array} \parallel \begin{array}{l} \dots \\ \text{read}() \hat{=} \\ \quad \text{local } y \text{ in} \\ \quad \quad \langle y := b; rs := rs.y \rangle; \\ \quad \text{return } y \\ \text{ni} \end{array} \right)$$

Atomicity means indivisibility. Interleaving will do.

Atomicity can be done with 'semaphores' as on the railways (block signalling).

For specification purposes I add a couple of auxiliary variables ws and rs .

The specification is that the read sequence destuttered is a subsequence of ws and the last element of ws is always in b .

$$[rs] \preceq ws \wedge ws_{\Omega} = b.$$

Dijkstra's solution: atomic accesses

local $b = \text{null}$, $ws = \langle \rangle.\text{null}$, $rs = \langle \rangle$ in

$$\left(\begin{array}{l} \dots \\ \text{write}(w) \hat{=} \\ \quad \langle b := w; ws := ws.w \rangle \end{array} \parallel \begin{array}{l} \dots \\ \text{read}() \hat{=} \\ \quad \text{local } y \text{ in} \\ \quad \quad \langle y := b; rs := rs.y \rangle; \\ \quad \text{return } y \\ \quad \text{ni} \end{array} \right)$$

Atomicity means indivisibility. Interleaving will do.

Atomicity can be done with 'semaphores' as on the railways (block signalling).

For specification purposes I add a couple of auxiliary variables ws and rs .

The specification is that the read sequence destuttered is a subsequence of ws and the last element of ws is always in b .

$$[rs] \preceq ws \wedge ws_{\Omega} = b.$$

But atomicity means waiting, and waiting isn't *simple* or even *certain*.

Cut down waiting with a two-slot buffer

Cut down waiting with a two-slot buffer

local $b = \text{null}$, $ws = \langle \rangle.\text{null}$, $rs = \langle \rangle$ in

(\dots	\dots)
	$\text{write}(w) \hat{=}$	$\text{read}() \hat{=}$	
	$\langle b := w; ws := ws.w \rangle$	$\text{local } y \text{ in}$	
		$\langle y := b; rs := rs.y \rangle;$	
		$\text{return } y$	
		ni	

local $c[2] = (\text{null}, \text{null})$, $l = 0$, $ws = \langle \rangle.\text{null}$, $rs = \langle \rangle$ in

(\dots	\dots)
	$\text{write}(w) \hat{=}$	$\text{read}() \hat{=}$	
	$\langle c[!l] := w; l := !l; \rangle$	$\text{local } y \text{ in}$	
	$\langle ws := ws.w \rangle$	$\langle y := c[l]; rs := rs.y \rangle;$	
		$\text{return } y$	
		ni	

Cut down waiting with a two-slot buffer

local $b = \text{null}$, $ws = \langle \rangle.\text{null}$, $rs = \langle \rangle$ in

(...	...)
	write(w) $\hat{=}$	read() $\hat{=}$	
	$\langle b := w; ws := ws.w \rangle$	local y in	
		$\langle y := b; rs := rs.y \rangle$;	
		return y	
		ni	

local $c[2] = (\text{null}, \text{null})$, $l = 0$, $ws = \langle \rangle.\text{null}$, $rs = \langle \rangle$ in

(...	...)
	write(w) $\hat{=}$	read() $\hat{=}$	
	$\langle c[!l] := w; l := !l; \rangle$	local y in	
	$ws := ws.w$	$\langle y := c[l]; rs := rs.y \rangle$;	
		return y	
		ni	

Data refinement!

Cut down waiting with a two-slot buffer

local $b = \text{null}$, $ws = \langle \rangle.\text{null}$, $rs = \langle \rangle$ in

$$\left(\begin{array}{l} \dots \\ \text{write}(w) \hat{=} \\ \quad \langle b := w; ws := ws.w \rangle \end{array} \parallel \begin{array}{l} \dots \\ \text{read}() \hat{=} \\ \quad \text{local } y \text{ in} \\ \quad \quad \langle y := b; rs := rs.y \rangle; \\ \quad \text{return } y \\ \text{ni} \end{array} \right)$$

local $c[2] = (\text{null}, \text{null})$, $l = 0$, $ws = \langle \rangle.\text{null}$, $rs = \langle \rangle$ in

$$\left(\begin{array}{l} \dots \\ \text{write}(w) \hat{=} \\ \quad \langle c[!l] := w; l := !l; \\ \quad \quad ws := ws.w \rangle \end{array} \parallel \begin{array}{l} \dots \\ \text{read}() \hat{=} \\ \quad \text{local } y \text{ in} \\ \quad \quad \langle y := c[l]; rs := rs.y \rangle; \\ \quad \text{return } y \\ \text{ni} \end{array} \right)$$

Data refinement! Now $[rs] \preceq ws \wedge ws_{\Omega} = c[l]$.

Simplify atomicity in the writer

local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

(...	...)
	write(w) $\hat{=}$ $\langle c[!l] := w; l := !l;$ $ws := ws.w \rangle$	read() $\hat{=}$ local y in $\langle y := c[l]; rs := rs.y \rangle;$ return y ni	

Simplify atomicity in the writer

local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

(...		...)
	write(w) $\hat{=}$		read() $\hat{=}$	
	$\langle c[!l] := w; l := !l; \rangle$		local y in	
	$\langle ws := ws.w \rangle$		$\langle y := c[l]; rs := rs.y \rangle;$	
			return y	
			ni	

local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

(...		...)
	write(w) $\hat{=}$		read() $\hat{=}$	
	local wt in		local y in	
	$\langle\langle wt := !l \rangle\rangle;$		$\langle y := c[l]; rs := rs.y \rangle;$	
	$\langle c[wt] := w \rangle;$		return y	
	$\langle\langle l := wt; ws := ws.w \rangle\rangle$		ni	
	ni			

Simplify atomicity in the writer

local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

(...	...)
	write(w) $\hat{=}$	read() $\hat{=}$	
	$\langle c[l] := w; l := !l; \rangle$	local y in	
	$\langle ws := ws.w \rangle$	$\langle y := c[l]; rs := rs.y \rangle;$	
		return y	
		ni	

local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

(...	...)
	write(w) $\hat{=}$	read() $\hat{=}$	
	local wt in	local y in	
	$\langle\langle wt := !l \rangle\rangle;$	$\langle y := c[l]; rs := rs.y \rangle;$	
	$\langle c[wt] := w \rangle;$	return y	
	$\langle\langle l := wt; ws := ws.w \rangle\rangle$	ni	
	ni		

Stores naturally serialise small reads and writes.

Simplify atomicity in the writer

local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

$$\left(\begin{array}{l} \dots \\ \text{write}(w) \hat{=} \\ \quad \langle c[l] := w; l := !l; \\ \quad \quad ws := ws.w \rangle \\ \dots \end{array} \parallel \begin{array}{l} \dots \\ \text{read}() \hat{=} \\ \quad \text{local } y \text{ in} \\ \quad \quad \langle y := c[l]; rs := rs.y \rangle; \\ \quad \quad \text{return } y \\ \dots \\ \text{ni} \end{array} \right)$$

local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

$$\left(\begin{array}{l} \dots \\ \text{write}(w) \hat{=} \\ \quad \text{local } wt \text{ in} \\ \quad \quad \langle\langle wt := !l \rangle\rangle; \\ \quad \quad \langle c[wt] := w \rangle; \\ \quad \quad \langle\langle l := wt; ws := ws.w \rangle\rangle \\ \dots \\ \text{ni} \end{array} \parallel \begin{array}{l} \dots \\ \text{read}() \hat{=} \\ \quad \text{local } y \text{ in} \\ \quad \quad \langle y := c[l]; rs := rs.y \rangle; \\ \quad \quad \text{return } y \\ \dots \\ \text{ni} \end{array} \right)$$

Stores naturally serialise small reads and writes. Still

$[rs] \preceq ws \wedge ws_{\Omega} = c[l]$.

(Can't) simplify atomicity in the reader

(Can't) simplify atomicity in the reader

local $c[2] = (\text{null}, \text{null}), l = 0, ws = \langle \rangle.\text{null}, rs = \langle \rangle$ in

...	...
write(w) $\hat{=}$	read() $\hat{=}$
local wt in	local y in
$\langle\langle wt := !l \rangle\rangle;$	$\langle y := c[l]; rs := rs.y \rangle;$
$\langle c[wt] := w \rangle;$	return y
$\langle\langle l := wt; ws := ws.w \rangle\rangle$	ni
ni	

(Can't) simplify atomicity in the reader

local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

...	...
write(w) $\hat{=}$	read() $\hat{=}$
local wt in	local y in
$\langle\langle wt := !l \rangle\rangle$;	$\langle y := c[l]; rs := rs.y \rangle$;
$\langle c[wt] := w \rangle$;	return y
$\langle\langle l := wt; ws := ws.w \rangle\rangle$	ni
ni	

local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

...	...
write(w) $\hat{=}$	read() $\hat{=}$
local wt in	local y, rt in
$\langle\langle wt := !l \rangle\rangle$;	$\langle\langle rt := l \rangle\rangle$;
$\langle c[wt] := w \rangle$;	$\langle y := c[rt]; rs := rs.y \rangle$;
$\langle\langle l := wt; ws := ws.w \rangle\rangle$	return y
ni	ni

(Can't) simplify atomicity in the reader

local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

...	...
write(w) $\hat{=}$	read() $\hat{=}$
local wt in	local y in
$\langle\langle wt := !l \rangle\rangle$;	$\langle y := c[l]; rs := rs.y \rangle$;
$\langle c[wt] := w \rangle$;	return y
$\langle\langle l := wt; ws := ws.w \rangle\rangle$	ni
ni	

local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

...	...
write(w) $\hat{=}$	read() $\hat{=}$
local wt in	local y, rt in
$\langle\langle wt := !l \rangle\rangle$;	$\langle\langle rt := l \rangle\rangle$;
$\langle c[wt] := w \rangle$;	$\langle y := c[rt]; rs := rs.y \rangle$;
$\langle\langle l := wt; ws := ws.w \rangle\rangle$	return y
ni	ni

Looks plausible, but it's broken.

(Can't) simplify atomicity in the reader

local $c[2] = (\text{null}, \text{null}), l = 0, \text{ws} = \langle \rangle.\text{null}, \text{rs} = \langle \rangle$ in

<pre>... write(w) ≐ local wt in ⟨⟨wt := !l⟩⟩; ⟨c[wt] := w⟩; ⟨⟨l := wt; ws := ws.w⟩⟩ ni</pre>	<pre>... read() ≐ local y in ⟨y := c[l]; rs := rs.y⟩; return y ni</pre>
--	---

local $c[2] = (\text{null}, \text{null}), l = 0, \text{ws} = \langle \rangle.\text{null}, \text{rs} = \langle \rangle$ in

<pre>... write(w) ≐ local wt in ⟨⟨wt := !l⟩⟩; ⟨c[wt] := w⟩; ⟨⟨l := wt; ws := ws.w⟩⟩ ni</pre>	<pre>... read() ≐ local y, rt in ⟨⟨rt := l⟩⟩; ⟨y := c[rt]; rs := rs.y⟩; return y ni</pre>
--	---

Looks plausible, but it's broken. Still $\text{ws}_\Omega = c[l]$

(Can't) simplify atomicity in the reader

local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

(...)
	write(w) $\hat{=}$		read() $\hat{=}$
	local wt in		local y in
	$\langle\langle wt := !l \rangle\rangle$;		$\langle y := c[l]; rs := rs.y \rangle$;
	$\langle c[wt] := w \rangle$;		return y
	$\langle\langle l := wt; ws := ws.w \rangle\rangle$		ni
	ni		

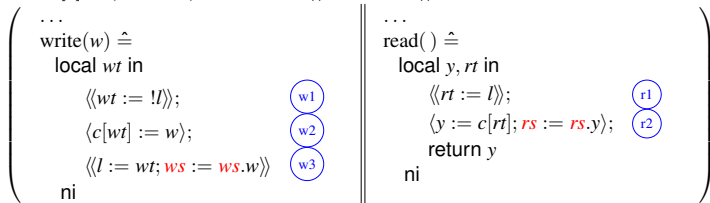
local $c[2] = (null, null), l = 0, ws = \langle \rangle.null, rs = \langle \rangle$ in

(...)
	write(w) $\hat{=}$		read() $\hat{=}$
	local wt in		local y, rt in
	$\langle\langle wt := !l \rangle\rangle$;		$\langle\langle rt := l \rangle\rangle$;
	$\langle c[wt] := w \rangle$;		$\langle y := c[rt]; rs := rs.y \rangle$;
	$\langle\langle l := wt; ws := ws.w \rangle\rangle$		return y
	ni		ni

Looks plausible, but it's broken. Still $ws_{\Omega} = c[l]$, but no longer $[rs] \preceq ws$.

What goes wrong?

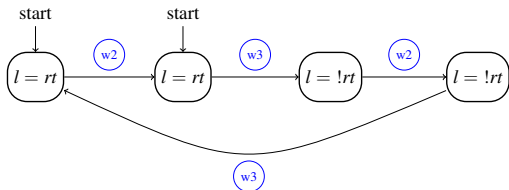
local $c[2] = (\text{null}, \text{null}), l = 0, ws = \langle \rangle.\text{null}, rs = \langle \rangle$ in



What goes wrong?

```
local c[2] = (null, null), l = 0, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
  local wt in
    ⟨⟨wt := !l⟩⟩;           (w1)
    ⟨c[wt] := w⟩;          (w2)
    ⟨⟨l := wt; ws := ws.w⟩⟩ (w3)
  ni
)
||
(
  ...
  read() ≐
  local y, rt in
    ⟨⟨rt := l⟩⟩;           (r1)
    ⟨y := c[rt]; rs := rs.y⟩; (r2)
  return y
  ni
)
```

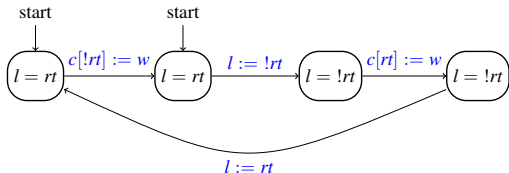
From the point of view of the reader, after $\langle\langle rt := l \rangle\rangle$, the writer behaves like this finite-state machine:



What goes wrong?

```
local c[2] = (null, null), l = 0, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
  local wt in
    ⟨⟨wt := !l⟩⟩;           (w1)
    ⟨c[wt] := w⟩;         (w2)
    ⟨⟨l := wt; ws := ws.w⟩⟩ (w3)
  ni
)
||
(
  ...
  read() ≐
  local y, rt in
    ⟨⟨rt := l⟩⟩;           (r1)
    ⟨y := c[rt]; rs := rs.y⟩; (r2)
  return y
  ni
)
```

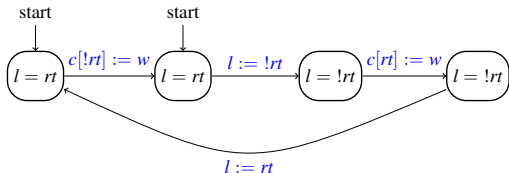
From the point of view of the reader, after $\langle\langle rt := l \rangle\rangle$, the writer behaves like this finite-state machine:



What goes wrong?

```
local c[2] = (null, null), l = 0, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
  local wt in
    ⟨⟨wt := !l⟩⟩;           (w1)
    ⟨c[wt] := w⟩;         (w2)
    ⟨⟨l := wt; ws := ws.w⟩⟩ (w3)
  ni
)
||
(
  ...
  read() ≐
  local y, rt in
    ⟨⟨rt := l⟩⟩;           (r1)
    ⟨y := c[rt]; rs := rs.y⟩; (r2)
  return y
  ni
)
```

From the point of view of the reader, after $\langle\langle rt := l \rangle\rangle$, the writer behaves like this finite-state machine:

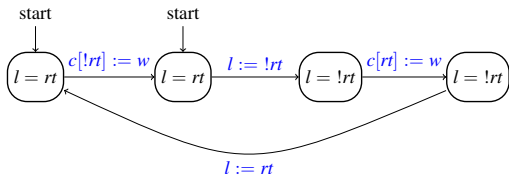


If the reader comes in at box 1 or 2 and reads at box 4, it will see the second value written

What goes wrong?

```
local c[2] = (null, null), l = 0, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
  local wt in
    ⟨⟨wt := !l⟩⟩;
    ⟨c[wt] := w⟩;
    ⟨⟨l := wt; ws := ws.w⟩⟩
  ni
)
||
(
  ...
  read() ≐
  local y, rt in
    ⟨⟨rt := l⟩⟩;
    ⟨y := c[rt]; rs := rs.y⟩;
    return y
  ni
)
```

From the point of view of the reader, after $\langle\langle rt := l \rangle\rangle$, the writer behaves like this finite-state machine:



If the reader comes in at box 1 or 2 and reads at box 4, it will see the second value written; if it then comes back quickly, it can see the first thing written!!

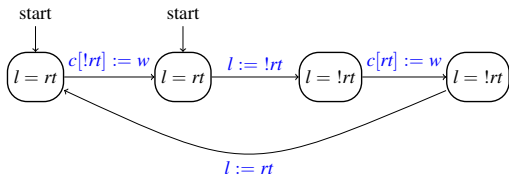
What goes wrong?

```

local c[2] = (null, null), l = 0, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
  local wt in
    ⟨⟨wt := !l⟩⟩;           (w1)
    ⟨c[wt] := w⟩;         (w2)
    ⟨⟨l := wt; ws := ws.w⟩⟩ (w3)
  ni
)
||
(
  ...
  read() ≐
  local y, rt in
    ⟨⟨rt := l⟩⟩;           (r1)
    ⟨y := c[rt]; rs := rs.y⟩; (r2)
  return y
  ni
)

```

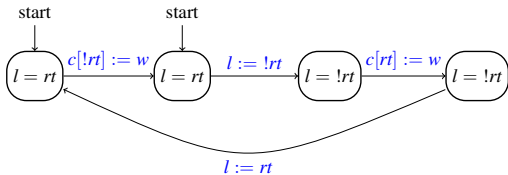
From the point of view of the reader, after $\langle\langle rt := l \rangle\rangle$, the writer behaves like this finite-state machine:



If the reader comes in at box 1 or 2 and reads at box 4, it will see the second value written; if it then comes back quickly, it can see the first thing written!! Also note $\langle c[rt] := w \rangle \parallel \langle y := c[rt] \rangle$.

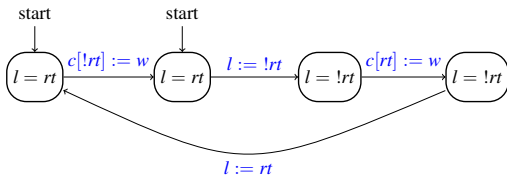
Can we repair it (1)?

All our problems (ordering, collisions) are caused by the third action



Can we repair it (1)?

All our problems (ordering, collisions) are caused by the third action

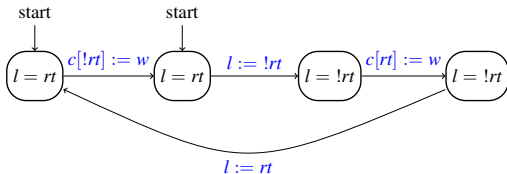


Can we detect when that action happens?

```
local c[2] = (null, null), l = 0, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
  local wt in
    ⟨wt := !l⟩;
    ⟨c[wt] := w⟩;
    ⟨l := wt; ws := ws.w⟩
  ni
  ||
  ...
  read() ≐
  local y, rt in
    ⟨rt := l⟩;
    ⟨y := c[rt]; rs := rs.y⟩;
  return y
  ni
)
```

Can we repair it (1)?

All our problems (ordering, collisions) are caused by the third action



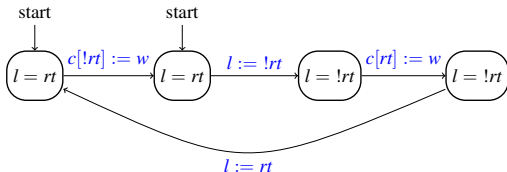
Can we detect when that action happens?

```
local c[2] = (null, null), l = 0, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
  local wt in
    ⟨wt := !l⟩;
    ⟨c[wt] := w⟩;
    ⟨l := wt; ws := ws.w⟩
  ni
  ||
  ...
  read() ≐
  local y, rt in
    ⟨rt := l⟩;
    ⟨y := c[rt]; rs := rs.y⟩;
    return y
  ni
)
```

No, because the writer can't tell the difference between the first and the third actions.

Can we repair it (2)?

All our problems (ordering, collisions) are caused by the third action

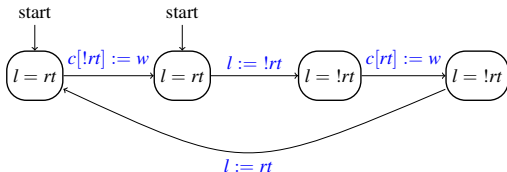


Can we detect when it *might* happen?

```
local c[2] = (null, null), l = 0, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
    local wt in
      ⟨wt := !l⟩;
      ⟨c[wt] := w⟩;
      ⟨l := wt; ws := ws.w⟩
  ni
  ||
  ...
  read() ≐
    local y, rt in
      ⟨rt := l⟩;
      ⟨y := c[rt]; rs := rs.y⟩;
      return y
  ni
)
```


Can we repair it (2)?

All our problems (ordering, collisions) are caused by the third action



Can we detect when it *might* happen?

```
local c[2] = (null, null), l = 0, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
  local wt in
    ⟨wt := !l⟩;
    ⟨c[wt] := w⟩;
    ⟨l := wt; ws := ws.w⟩
  ni
  ||
  ...
  read() ≐
  local y, rt in
    ⟨rt := l⟩;
    ⟨y := c[rt]; rs := rs.y⟩;
    return y
  ni
)
```

Yes: it becomes possible after the second action.

A repair (back to correctness)

The writer signals when disaster becomes possible; the reader incorporates the signal in its answer.

local $c[2] = (\text{null}, \text{null}), l = 0, \text{ok}, \text{ws} = \langle \rangle.\text{null}, \text{rs} = \langle \rangle$ in

(...			...)
	write(w) $\hat{=}$			read() $\hat{=}$	
	local wt in			local y, rt in	
	$\langle\langle wt := !l \rangle\rangle$;			$\langle\langle ok := \text{true} \rangle\rangle$;	
	$\langle c[wt] := w \rangle$;			$\langle\langle rt := l \rangle\rangle$;	
	$\langle\langle l := wt; \text{ws} := \text{ws}.w \rangle\rangle$;			$\langle y := c[rt] \rangle$;	
	$\langle\langle ok := \text{false} \rangle\rangle$;			$\langle\langle rt = ok; \text{rs} := \text{rs}.(rt, y) \rangle\rangle$;	
	ni			return (rt, y)	
				ni	

A repair (back to correctness)

The writer signals when disaster becomes possible; the reader incorporates the signal in its answer.

```
local c[2] = (null, null), l = 0, ok, ws = ⟨⟩.null, rs = ⟨⟩ in
```

(...			...)
	write(<i>w</i>) $\hat{=}$			read() $\hat{=}$	
	local <i>wt</i> in			local <i>y</i> , <i>rt</i> in	
	⟨⟨ <i>wt</i> := ! <i>l</i> ⟩⟩;			⟨⟨ <i>ok</i> := true⟩⟩;	
	⟨ <i>c</i> [<i>wt</i>] := <i>w</i> ⟩;			⟨⟨ <i>rt</i> := <i>l</i> ⟩⟩;	
	⟨⟨ <i>l</i> := <i>wt</i> ; <i>ws</i> := <i>ws.w</i> ⟩⟩;			⟨ <i>y</i> := <i>c</i> [<i>rt</i>]⟩;	
	⟨⟨ <i>ok</i> := false⟩⟩;			⟨⟨ <i>rt</i> = <i>ok</i> ; <i>rs</i> := <i>rs</i> .(<i>rt</i> , <i>y</i>)⟩⟩;	
	ni			return (<i>rt</i> , <i>y</i>)	
				ni	

– and then we notice that we don't need atomic buffer accesses any more.

A repair (back to correctness)

The writer signals when disaster becomes possible; the reader incorporates the signal in its answer.

local $c[2] = (\text{null}, \text{null}), l = 0, ok, ws = \langle \rangle, null, rs = \langle \rangle$ in

(...			...)
	write(w) $\hat{=}$			read() $\hat{=}$	
	local wt in			local y, rt in	
	$\langle\langle wt := !l \rangle\rangle$;			$\langle\langle ok := \text{true} \rangle\rangle$;	
	$c[wt] := w$;			$\langle\langle rt := l \rangle\rangle$;	
	$\langle\langle l := wt; ws := ws.w \rangle\rangle$;			$y := c[rt]$;	
	$\langle\langle ok := \text{false} \rangle\rangle$;			$\langle\langle rt = ok; rs := rs.(rt, y) \rangle\rangle$;	
	ni			return (rt, y)	
				ni	

– and then we notice that we don't need atomic buffer accesses any more.

A repair (back to correctness)

The writer signals when disaster becomes possible; the reader incorporates the signal in its answer.

local $c[2] = (\text{null}, \text{null}), l = 0, \text{ok}, \text{ws} = \langle \rangle.\text{null}, \text{rs} = \langle \rangle$ in

(...	...)
	write(w) $\hat{=}$	read() $\hat{=}$	
	local wt in	local y, rt in	
	$\langle\langle wt := !l \rangle\rangle;$	$\langle\langle ok := \text{true} \rangle\rangle;$	
	$c[wt] := w;$	$\langle\langle rt := l \rangle\rangle;$	
	$\langle\langle l := wt; \text{ws} := \text{ws}.w \rangle\rangle;$	$y := c[rt];$	
	$\langle\langle ok := \text{false} \rangle\rangle;$	$\langle\langle rt = ok; \text{rs} := \text{rs}.(rt, y) \rangle\rangle;$	
	ni	return (rt, y)	
		ni	

– and then we notice that we don't need atomic buffer accesses any more.

If $\tilde{\text{rs}}$ is rs with the (false, -) results taken out and the true labels discarded, then we have $\lfloor \tilde{\text{rs}} \rfloor \preceq \text{ws} \wedge \text{ws}_\Omega = c[l]$.

Summary so far, and a criticism

Summary so far, and a criticism

We have gone from an atomic single-slot buffer (not wait-free)

Summary so far, and a criticism

We have gone from an atomic single-slot buffer (not wait-free) to an atomic double-slot buffer (ditto)

Summary so far, and a criticism

We have gone from an atomic single-slot buffer (not wait-free) to an atomic double-slot buffer (ditto) to a faulty not so completely atomic double-slot buffer (still not wait-free)

Summary so far, and a criticism

We have gone from an atomic single-slot buffer (not wait-free) to an atomic double-slot buffer (ditto) to a faulty not so completely atomic double-slot buffer (still not wait-free) to a working wait-free non-atomic double-slot buffer that tells us when it's succeeded.

Summary so far, and a criticism

We have gone from an atomic single-slot buffer (not wait-free) to an atomic double-slot buffer (ditto) to a faulty not so completely atomic double-slot buffer (still not wait-free) to a working wait-free non-atomic double-slot buffer that tells us when it's succeeded.

We certainly haven't proceeded by “a sequence of true understatements” (Lakatos);

Summary so far, and a criticism

We have gone from an atomic single-slot buffer (not wait-free) to an atomic double-slot buffer (ditto) to a faulty not so completely atomic double-slot buffer (still not wait-free) to a working wait-free non-atomic double-slot buffer that tells us when it's succeeded.

We certainly haven't proceeded by “a sequence of true understatements” (Lakatos); we have made at least one “false overstatement”;

Summary so far, and a criticism

We have gone from an atomic single-slot buffer (not wait-free) to an atomic double-slot buffer (ditto) to a faulty not so completely atomic double-slot buffer (still not wait-free) to a working wait-free non-atomic double-slot buffer that tells us when it's succeeded.

We certainly haven't proceeded by “a sequence of true understatements” (Lakatos); we have made at least one “false overstatement”; perhaps we have made a step of “exception barring”.

Summary so far, and a criticism

We have gone from an atomic single-slot buffer (not wait-free) to an atomic double-slot buffer (ditto) to a faulty not so completely atomic double-slot buffer (still not wait-free) to a working wait-free non-atomic double-slot buffer that tells us when it's succeeded.

We certainly haven't proceeded by "a sequence of true understatements" (Lakatos); we have made at least one "false overstatement"; perhaps we have made a step of "exception barring".

I see the step that includes the 'ok' variable as an example of thesis (program) plus antithesis (counter-example) yielding synthesis (repaired program).

Summary so far, and a criticism

We have gone from an atomic single-slot buffer (not wait-free) to an atomic double-slot buffer (ditto) to a faulty not so completely atomic double-slot buffer (still not wait-free) to a working wait-free non-atomic double-slot buffer that tells us when it's succeeded.

We certainly haven't proceeded by “a sequence of true understatements” (Lakatos); we have made at least one “false overstatement”; perhaps we have made a step of “exception barring”.

I see the step that includes the ‘*ok*’ variable as an example of thesis (program) plus antithesis (counter-example) yielding synthesis (repaired program).

But what use is the pair (false, *something*)? What can a user do but ignore *something* and try to read again?

A more honest repair

local $c[2] = (null, null), l = 0, ok, ws = \langle \rangle.null, rs = \langle \rangle$ in

(... write(w) $\hat{=}$ local wt in $\langle\langle wt := !l \rangle\rangle$; $c[wt] := w$; $\langle\langle l := wt; ws := ws.w \rangle\rangle$; $\langle\langle ok := false \rangle\rangle$; ni		... read() $\hat{=}$ local y, rt in do $\langle\langle ok := true \rangle\rangle$; $\langle\langle rt := l \rangle\rangle$; $y := c[rt]$; $\langle\langle rt = ok \rangle\rangle$ until rt ; $\langle\langle rs := rs.y \rangle\rangle$; return y ni)
---	--	--	---	---

A more honest repair

local $c[2] = (\text{null}, \text{null}), l = 0, \text{ok}, \text{ws} = \langle \rangle.\text{null}, \text{rs} = \langle \rangle$ in

(...		...)
	write(w) $\hat{=}$		read() $\hat{=}$	
	local wt in		local y, rt in	
	$\langle\langle wt := !l \rangle\rangle;$		do	
	$c[wt] := w;$		$\langle\langle ok := \text{true} \rangle\rangle;$	
	$\langle\langle l := wt; \text{ws} := \text{ws}.w \rangle\rangle;$		$\langle\langle rt := l \rangle\rangle;$	
	$\langle\langle ok := \text{false} \rangle\rangle;$		$y := c[rt];$	
	ni		$\langle\langle rt = ok \rangle\rangle$	
			until $rt;$	
			$\langle\langle \text{rs} := \text{rs}.y \rangle\rangle;$	
			return y	
			ni	

Obviously not wait-free.

A more honest repair

local $c[2] = (\text{null}, \text{null}), l = 0, ok, ws = \langle \rangle.\text{null}, rs = \langle \rangle$ in

(... write(w) $\hat{=}$ local wt in $\langle\langle wt := !l \rangle\rangle$; $c[wt] := w$; $\langle\langle l := wt; ws := ws.w \rangle\rangle$; $\langle\langle ok := false \rangle\rangle$; ni		... read() $\hat{=}$ local y, rt in do $\langle\langle ok := true \rangle\rangle$; $\langle\langle rt := l \rangle\rangle$; $y := c[rt]$; $\langle\langle rt = ok \rangle\rangle$ until rt ; $\langle\langle rs := rs.y \rangle\rangle$; return y ni)
---	--	--	---	---

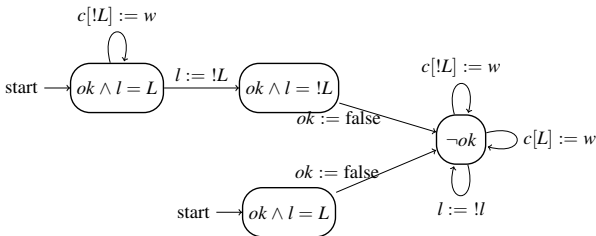
Obviously not wait-free. But otherwise repaired.

A more honest repair

local $c[2] = (\text{null}, \text{null}), l = 0, \text{ok}, \text{ws} = \langle \rangle, \text{null}, \text{rs} = \langle \rangle$ in

<pre> ... write(w) $\hat{=}$ local wt in $\langle\langle wt := !l \rangle\rangle$; $c[wt] := w$; $\langle\langle l := wt; \text{ws} := \text{ws}.w \rangle\rangle$; $\langle\langle \text{ok} := \text{false} \rangle\rangle$; ni </pre>	<pre> ... read() $\hat{=}$ local y, rt in do $\langle\langle \text{ok} := \text{true} \rangle\rangle$; $\langle\langle \text{rt} := l \rangle\rangle$; $y := c[rt]$; $\langle\langle \text{rt} = \text{ok} \rangle\rangle$ until rt; $\langle\langle \text{rs} := \text{rs}.y \rangle\rangle$; return y ni </pre>
---	---

Obviously not wait-free. But otherwise repaired. Finite-state machine now



Try three slots

```
local c[2] = (null, null), d = null, l = 0, ok, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
    local wt in
      ⟨⟨wt := !l⟩⟩;
      c[wt] := w;
      ⟨⟨l := wt; ws := ws.w⟩⟩;
      ⟨d := w⟩;
      ⟨⟨ok := false⟩⟩;
  ni
  ||
  ...
  read() ≐
    local y, rt in
      ⟨⟨ok := true⟩⟩;
      ⟨⟨rt := l⟩⟩;
      y := c[rt];
      ⟨⟨rt = ok⟩⟩;
      if ¬rt then ⟨y := d⟩ else skip fi;
      ⟨⟨rs := rs.y⟩⟩;
      return y
  ni
)
```

The writer can first write in a side-channel, then signal that mayhem approaches. If it gets the signal, the reader uses the side-channel.

Try three slots

```
local c[2] = (null, null), d = null, l = 0, ok, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
    local wt in
      ⟨⟨wt := !l⟩⟩;
      c[wt] := w;
      ⟨⟨l := wt; ws := ws.w⟩⟩;
      ⟨d := w⟩;
      ⟨⟨ok := false⟩⟩;
  ni
  ||
  ...
  read() ≐
    local y, rt in
      ⟨⟨ok := true⟩⟩;
      ⟨⟨rt := l⟩⟩;
      y := c[rt];
      ⟨⟨rt = ok⟩⟩;
      if ¬rt then ⟨y := d⟩ else skip fi;
      ⟨⟨rs := rs.y⟩⟩;
      return y
  ni
)
```

The writer can first write in a side-channel, then signal that mayhem approaches. If it gets the signal, the reader uses the side-channel.

Perhaps this will work ...

Try three slots

```
local c[2] = (null, null), d = null, l = 0, ok, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
    local wt in
      ⟨⟨wt := !l⟩⟩;
      c[wt] := w;
      ⟨⟨l := wt; ws := ws.w⟩⟩;
      ⟨⟨wt := ok⟩⟩;
      if wt then ⟨d := w⟩; ⟨ok := false⟩
        else skip fi
    ni
  ...
  read() ≐
    local y, rt in
      ⟨⟨ok := true⟩⟩;
      ⟨⟨rt := l⟩⟩;
      y := c[rt];
      ⟨⟨rt = ok⟩⟩;
      if ¬rt then ⟨y := d⟩ else skip fi;
      ⟨⟨rs := rs.y⟩⟩;
      return y
    ni
)
```

The writer can first write in a side-channel, then signal that mayhem approaches. If it gets the signal, the reader uses the side-channel.

Perhaps this will work ... but it's more likely to work if the writer only writes when the reader is asking for it.

Try three slots

```
local c[2] = (null, null), d = null, l = 0, ok, ws = ⟨⟩.null, rs = ⟨⟩ in
```

(... write(w) $\hat{=}$ local wt in ⟨⟨wt := !l⟩⟩; c[wt] := w; ⟨⟨l := wt; ws := ws.w⟩⟩; ⟨⟨wt := ok⟩⟩; if wt then d := w; ⟨⟨ok := false⟩⟩ else skip fi ni)	... read() $\hat{=}$ local y, rt in ⟨⟨ok := true⟩⟩; ⟨⟨rt := l⟩⟩; y := c[rt]; ⟨⟨rt = ok⟩⟩; if ¬rt then y := d else skip fi; ⟨⟨rs := rs.y⟩⟩; return y ni
---	---	--	---	--

The writer can first write in a side-channel, then signal that mayhem approaches. If it gets the signal, the reader uses the side-channel.

Perhaps this will work ... but it's more likely to work if the writer only writes when the reader is asking for it.

And then I notice that they alternate, and I can use *non-atomic* read and write.

Try three slots

```
local c[2] = (null, null), d = null, l = 0, ok, ws = ⟨⟩.null, rs = ⟨⟩ in
```

(...	...)
	write(w) $\hat{=}$	read() $\hat{=}$	
	local wt in	local y, rt in	
	⟨⟨wt := !l⟩⟩;	⟨⟨ok := true⟩⟩;	
	c[wt] := w;	⟨⟨rt := l⟩⟩;	
	⟨⟨l := wt; ws := ws.w⟩⟩;	y := c[rt];	
	⟨⟨wt := ok⟩⟩;	⟨⟨rt = ok⟩⟩;	
	if wt then d := w; ⟨⟨ok := false⟩⟩	if \neg rt then y := d else skip fi;	
	else skip fi	⟨⟨rs := rs.y⟩⟩;	
	ni	return y	
		ni	

The writer can first write in a side-channel, then signal that mayhem approaches. If it gets the signal, the reader uses the side-channel.

Perhaps this will work ... but it's more likely to work if the writer only writes when the reader is asking for it.

And then I notice that they alternate, and I can use *non-atomic* read and write. This is Harris's algorithm, rationally developed.

Recapitulation

```
local b = null in
  (
    ...
    write(w) ≐
      ⟨b := w⟩
  ||
    ...
    read() ≐
      local y in
        ⟨y := b⟩;
        return y
      ni
  )
```

– Dijkstra's single-place buffer.

Recapitulation

```
local  $b = \text{null}$ ,  $ws = \langle \rangle.\text{null}$ ,  $rs = \langle \rangle$  in  
  (  $\dots$  |  $\dots$   
    write( $w$ )  $\hat{=}$  | read()  $\hat{=}$   
       $\langle b := w; ws := ws.w \rangle$  | local  $y$  in  
                                |    $\langle y := b; rs := rs.y \rangle$ ;  
                                |   return  $y$   
                                |   ni  
                                | )
```

– with auxiliary ws and rs to show $\llbracket rs \rrbracket \preceq ws \wedge ws_{\Omega} = b$.

Recapitulation

```
local c[2] = (null, null), l = 0, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
  ⟨ c[l] := w; l := !l;
    ws := ws.w ⟩
  ||
  ...
  read() ≐
  local y in
  ⟨ y := c[l]; rs := rs.y ⟩;
  return y
  ni
)
```

– data refinement to two slots and $[rs] \preccurlyeq ws \wedge ws_{\Omega} = c[l]$.

Recapitulation

```
local c[2] = (null, null), l = 0, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
    local wt in
      ⟨wt := !l⟩;
      ⟨c[wt] := w⟩;
      ⟨l := wt; ws := ws.w⟩
  ni
  ||
  ...
  read() ≐
    local y in
      ⟨y := c[l]; rs := rs.y⟩;
      return y
  ni
)
```

– atomicity refinement in the writer; still $[rs] \preceq ws \wedge ws_{\Omega} = c[l]$.

Recapitulation

```
local c[2] = (null, null), l = 0, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
    local wt in
      ⟨wt := !l⟩;
      ⟨c[wt] := w⟩;
      ⟨l := wt; ws := ws.w⟩
  ni
  ||
  ...
  read() ≐
    local y, rt in
      ⟨rt := l⟩;
      ⟨y := c[rt]; rs := rs.y⟩;
      return y
  ni
)
```

– atomicity refinement in the reader; now $[rs] \not\approx ws \wedge ws_{\Omega} = c[l]$.

Recapitulation

```
local c[2] = (null, null), l = 0, ok, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
    local wt in
      ⟨⟨wt := !l⟩⟩;
      c[wt] := w;
      ⟨⟨l := wt; ws := ws.w⟩⟩;
      ⟨⟨ok := false⟩⟩;
  ni
  ||
  ...
  read() ≐
    local y, rt in
      ⟨⟨ok := true⟩⟩;
      ⟨⟨rt := l⟩⟩;
      y := c[rt];
      ⟨⟨rt = ok; rs := rs.(rt, y)⟩⟩;
  return (rt, y)
  ni
)
```

– exception barring; now $\llbracket \tilde{rs} \rrbracket \rightsquigarrow ws \wedge ws_{\Omega} = c[l]$.

Recapitulation

```
local c[2] = (null, null), l = 0, ok, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
  local wt in
    ⟨⟨wt := !l⟩⟩;
    c[wt] := w;
    ⟨⟨l := wt; ws := ws.w⟩⟩;
    ⟨⟨ok := false⟩⟩;
  ni
  ||
  ...
  read() ≐
  local y, rt in
    do
      ⟨⟨ok := true⟩⟩;
      ⟨⟨rt := l⟩⟩;
      y := c[rt];
      ⟨⟨rt = ok⟩⟩
    until rt;
    ⟨⟨rs := rs.y⟩⟩;
    return y
  ni
)
```

– more honest exception barring; once again $[rs] \preceq ws \wedge ws_{\Omega} = c[l]$, but no longer wait-free.

Recapitulation

local $c[2] = (\text{null}, \text{null}), d = \text{null}, l = 0, ok, ws = \langle \rangle, null, rs = \langle \rangle$ in

<pre>... write(w) ≐ local wt in ⟨wt := !l⟩; c[wt] := w; ⟨l := wt; ws := ws.w⟩; ⟨wt := ok⟩; if wt then d := w; ⟨ok := false⟩ else skip fi ni</pre>	<pre>... read() ≐ local y, rt in ⟨ok := true⟩; ⟨rt := l⟩; y := c[rt]; ⟨rt = ok⟩; if ¬rt then y := d else skip fi; ⟨rs := rs.y⟩; return y ni</pre>
---	---

– three slots; $[rs] \preccurlyeq ws \wedge ws_{\Omega} = c[l]$; wait-free.

Recapitulation

```
local c[2] = (null, null), d = null, l = 0, ok, ws = ⟨⟩.null, rs = ⟨⟩ in
(
  ...
  write(w) ≐
  local wt in
    ⟨⟨wt := !l⟩⟩;
    c[wt] := w;
    ⟨⟨l := wt; ws := ws.w⟩⟩;
    ⟨⟨wt := ok⟩⟩;
    if wt then d := w; ⟨⟨ok := false⟩⟩
    else skip fi
  ni
  ||
  ...
  read() ≐
  local y, rt in
    ⟨⟨ok := true⟩⟩;
    ⟨⟨rt := l⟩⟩;
    y := c[rt];
    ⟨⟨rt = ok⟩⟩;
    if ¬rt then y := d else skip fi;
    ⟨⟨rs := rs.y⟩⟩;
    return y
  ni
)
```

– three slots; $[rs] \preceq ws \wedge ws_{\Omega} = c[l]$; wait-free.

Proof available on application.

A programmer's instinct demands ...

```
local c[2] = (null, null), d = null, l = 0, ok, ws = ⟨⟩, null, rs = ⟨⟩ in
```

(... write(<i>w</i>) $\hat{=}$ local <i>wt</i> in ⟨⟨ <i>wt</i> := ! <i>l</i> ⟩⟩; <i>c</i> [<i>wt</i>] := <i>w</i> ; ⟨⟨ <i>l</i> := <i>wt</i> ; <i>ws</i> := <i>ws.w</i> ⟩⟩; ⟨⟨ <i>wt</i> := <i>ok</i> ⟩⟩; if <i>wt</i> then <i>d</i> := <i>w</i> ; ⟨⟨ <i>ok</i> := false⟩⟩ else skip fi ni		... read() $\hat{=}$ local <i>y</i> , <i>rt</i> in ⟨⟨ <i>ok</i> := true⟩⟩; ⟨⟨ <i>rt</i> := <i>l</i> ⟩⟩; <i>y</i> := <i>c</i> [<i>rt</i>]; ⟨⟨ <i>rt</i> = <i>ok</i> ⟩⟩; if \neg <i>rt</i> then <i>y</i> := <i>d</i> else ⟨⟨ <i>ok</i> := false⟩⟩ fi; ⟨⟨ <i>rs</i> := <i>rs.y</i> ⟩⟩; return <i>y</i> ni)
---	--	--	---	---

– the reader tells the writer when there's no need to use the side channel.

A programmer's instinct demands ...

```
local c[2] = (null, null), d = null, l = 0, ok, ws = ⟨⟩.null, rs = ⟨⟩ in
```

(... write(w) ≐ local wt in ⟨⟨wt := !l⟩⟩; c[wt] := w; ⟨⟨l := wt; ws := ws.w⟩⟩; ⟨⟨wt := ok⟩⟩; if wt then d := w; ⟨⟨ok := false⟩⟩ else skip fi ni		... read() ≐ local y, rt in ⟨⟨ok := true⟩⟩; ⟨⟨rt := l⟩⟩; y := c[rt]; ⟨⟨rt = ok⟩⟩; if ¬rt then y := d else ⟨⟨ok := false⟩⟩ fi; ⟨⟨rs := rs.y⟩⟩; return y ni)
---	---	--	--	---

– the reader tells the writer when there's no need to use the side channel. It's no longer true that the writer only writes when *ok*.

A programmer's instinct demands ...

```
local c[2] = (null, null), d = null, l = 0, ok, ws = ⟨⟩, null, rs = ⟨⟩ in
```

(... write(w) ≐ local wt in ⟨⟨wt := !l⟩⟩; c[wt] := w; ⟨⟨l := wt; ws := ws.w⟩⟩; ⟨⟨wt := ok⟩⟩; if wt then d := w; ⟨⟨ok := false⟩⟩ else skip fi ni		... read() ≐ local y, rt in ⟨⟨ok := true⟩⟩; ⟨⟨rt := l⟩⟩; y := c[rt]; ⟨⟨rt = ok⟩⟩; if ¬rt then y := d else ⟨⟨ok := false⟩⟩ fi; ⟨⟨rs := rs.y⟩⟩; return y ni)
---	---	--	--	---

– the reader tells the writer when there's no need to use the side channel. It's no longer true that the writer only writes when *ok*. But from the point of view of the reader, nothing has changed!

A programmer's instinct demands ...

```
local c[2] = (null, null), d = null, l = 0, ok, ws = ⟨⟩, null, rs = ⟨⟩ in
```

(... write(w) ≐ local wt in ⟨⟨wt := !l⟩⟩; c[wt] := w; ⟨⟨l := wt; ws := ws.w⟩⟩; ⟨⟨wt := ok⟩⟩; if wt then d := w; ⟨⟨ok := false⟩⟩ else skip fi ni)	... read() ≐ local y, rt in ⟨⟨ok := true⟩⟩; ⟨⟨rt := l⟩⟩; y := c[rt]; ⟨⟨rt = ok⟩⟩; if ¬rt then y := d else ⟨⟨ok := false⟩⟩ fi; ⟨⟨rs := rs.y⟩⟩; return y ni
---	---	--	---	--

– the reader tells the writer when there's no need to use the side channel. It's no longer true that the writer only writes when *ok*. But from the point of view of the reader, nothing has changed!

In fact the writer writes when *ok* or the reader is asleep.

A programmer's instinct demands ...

```
local c[2] = (null, null), d = null, l = 0, ok, ws = ⟨⟩, null, rs = ⟨⟩ in
```

(... write(w) $\hat{=}$ local wt in ⟨⟨wt := !l⟩⟩; c[wt] := w; ⟨⟨l := wt; ws := ws.w⟩⟩; ⟨⟨wt := ok⟩⟩; if wt then d := w; ⟨⟨ok := false⟩⟩ else skip fi ni		... read() $\hat{=}$ local y, rt in ⟨⟨ok := true⟩⟩; ⟨⟨rt := l⟩⟩; y := c[rt]; ⟨⟨rt = ok⟩⟩; if \neg rt then y := d else ⟨⟨ok := false⟩⟩ fi; ⟨⟨rs := rs.y⟩⟩; return y ni)
---	---	--	--	---

– the reader tells the writer when there's no need to use the side channel. It's no longer true that the writer only writes when *ok*. But from the point of view of the reader, nothing has changed!

In fact the writer writes when *ok* or the reader is asleep. Easy to fix with another auxiliary variable, proof available on request.

Pipelines, caches, weak memory models

In my argument I relied on sequencing of actions and on serialisation of (small) memory accesses.

Pipelines, caches, weak memory models

In my argument I relied on sequencing of actions and on serialisation of (small) memory accesses.

In modern machines these assumptions don't hold.

Pipelines, caches, weak memory models

In my argument I relied on sequencing of actions and on serialisation of (small) memory accesses.

In modern machines these assumptions don't hold.

Pipelines reorder instructions if there isn't an obvious dependency – e.g. in $c[wt] := w; l := wt$ or in $d := w; ok := \text{false}$.

Pipelines, caches, weak memory models

In my argument I relied on sequencing of actions and on serialisation of (small) memory accesses.

In modern machines these assumptions don't hold.

Pipelines reorder instructions if there isn't an obvious dependency – e.g. in $c[wt] := w; l := wt$ or in $d := w; ok := \text{false}$.

Caches delay interaction with the store.

Pipelines, caches, weak memory models

In my argument I relied on sequencing of actions and on serialisation of (small) memory accesses.

In modern machines these assumptions don't hold.

Pipelines reorder instructions if there isn't an obvious dependency – e.g. in $c[wt] := w; l := wt$ or in $d := w; ok := \text{false}$.

Caches delay interaction with the store. Viewed from another machine, store actions are reordered.

Pipelines, caches, weak memory models

In my argument I relied on sequencing of actions and on serialisation of (small) memory accesses.

In modern machines these assumptions don't hold.

Pipelines reorder instructions if there isn't an obvious dependency – e.g. in $c[wt] := w; l := wt$ or in $d := w; ok := \text{false}$.

Caches delay interaction with the store. Viewed from another machine, store actions are reordered.

So-called “weak memory models” are difficult to understand.

Pipelines, caches, weak memory models

In my argument I relied on sequencing of actions and on serialisation of (small) memory accesses.

In modern machines these assumptions don't hold.

Pipelines reorder instructions if there isn't an obvious dependency – e.g. in $c[wt] := w; l := wt$ or in $d := w; ok := \text{false}$.

Caches delay interaction with the store. Viewed from another machine, store actions are reordered.

So-called “weak memory models” are difficult to understand. That is, to understand well enough to reason about programs running under them.

Pipelines, caches, weak memory models

In my argument I relied on sequencing of actions and on serialisation of (small) memory accesses.

In modern machines these assumptions don't hold.

Pipelines reorder instructions if there isn't an obvious dependency – e.g. in $c[wt] := w; l := wt$ or in $d := w; ok := \text{false}$.

Caches delay interaction with the store. Viewed from another machine, store actions are reordered.

So-called “weak memory models” are difficult to understand. That is, to understand well enough to reason about programs running under them.

Maybe an EPSRC project ...

And finally

Programming is creative, experimental.

And finally

Programming is creative, experimental.

That means we will go wrong.

And finally

Programming is creative, experimental.

That means we will go wrong.

When we go wrong we can fix our program, refine our conjecture, or both.

And finally

Programming is creative, experimental.

That means we will go wrong.

When we go wrong we can fix our program, refine our conjecture, or both.

The mechanisms of refinement can guide us, but we may (I would say must) sometimes make “false overstatements”.

And finally

Programming is creative, experimental.

That means we will go wrong.

When we go wrong we can fix our program, refine our conjecture, or both.

The mechanisms of refinement can guide us, but we may (I would say must) sometimes make “false overstatements”.

Lakatos’s dialectic lives!