


Automatically Comparing Memory Consistency Models



John Wickerson
Imperial College



Mark Batty
University of Kent




Tyler Sorensen
Imperial College



George Constantinides
Imperial College

Memory Models

```
x = 1;   ||   y = 1;  
r0 = y;  ||   r1 = x;
```



Summary

- **The problem is** that memory models are complex and counterintuitive. Misunderstanding them has led to lots of bugs.
- **Our contribution is** a technique for automatically generating litmus tests that distinguish two models.
- **This is useful because** we can compare two variants of a model, and we can debug compilers.

Three tasks

1. To show that **M** is **stronger** than **N**:
 - find P, σ where $\sigma \notin \llbracket P \rrbracket_M$ and $\sigma \in \llbracket P \rrbracket_N$.
2. To show that a **compiler optimisation** on **M** is unsound:
 - find P, Q, σ where $\sigma \notin \llbracket P \rrbracket_M$, P optimises to Q , and $\sigma \in \llbracket Q \rrbracket_M$.
3. To show that a **compiler mapping** from **M** to **N** is unsound:
 - find P, Q, σ where $\sigma \notin \llbracket P \rrbracket_M$, P maps to Q , and $\sigma \in \llbracket Q \rrbracket_N$.

Related work


	Manual c'examples	Manual proofs	Automatically
Comparing memory models	[Batty et al. <i>POPL'16</i>] [Nienhuis et al. <i>OOPSLA'16</i>] [Lahav et al. <i>POPL'16</i>]	[Owens et al. <i>TPHOLs'09</i>] [Mador-Haim et al. <i>CAV'12</i>]	[Mador-Haim et al. <i>CAV'10</i>] THIS WORK
Checking compiler optimisations	[Vafeiadis et al. <i>POPL'15</i>]	[Sevcik <i>PLDI'11</i>]	[Burckhardt et al. <i>CC'10</i>] [Chakraborty et al. <i>CGO'16</i>] THIS WORK
Checking compiler mappings	[Wickerson et al. <i>OOPSLA'15</i>]	[Batty et al. <i>POPL'11</i>]	[Lustig et al. <i>MICRO'14</i>] THIS WORK

Results

	Previous work	Our tool (automatic)	Time taken
Comparison: Batty et al.'s C++ vs. original	7	5	<1 sec
Comparison: Lahav et al.'s C++ vs. original	10	6	1 sec
Comparison: Nienhuis et al.'s C++ vs. original	12	12	4 min
Optimisation: strengthening memory order in C++	8	7	1 min
Optimisation: thread linearisation in C++	6	6	5 sec
Compiler mapping: from OpenCL to AMD	12	11	22 min
Compiler mapping: from C++ to Power	14	14	2 min

Memory Models

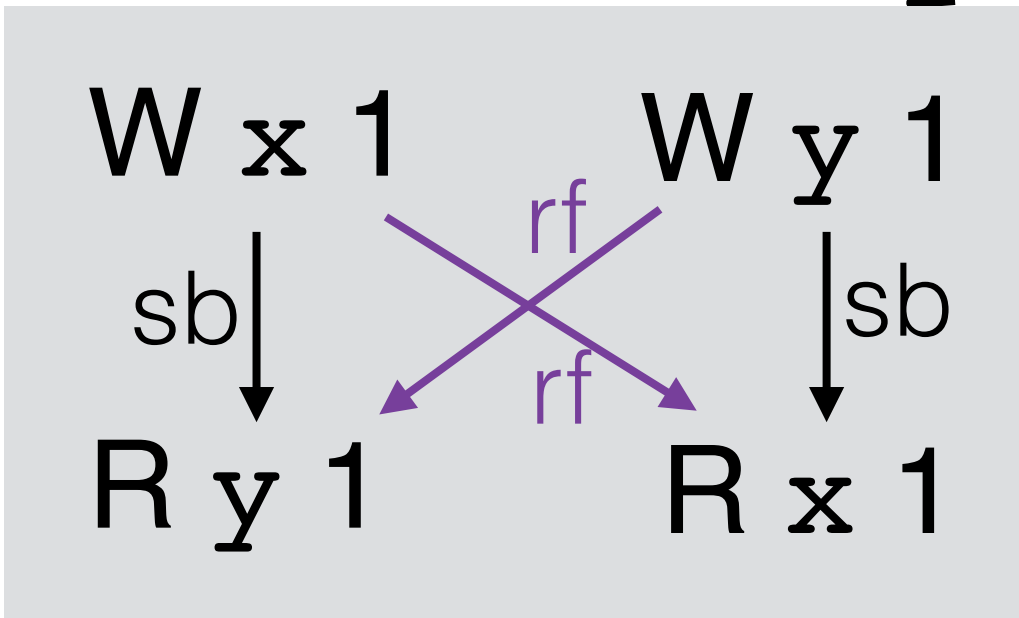
```
x = 1;   ||   y = 1;  
r0 = y;  ||   r1 = x;
```



```

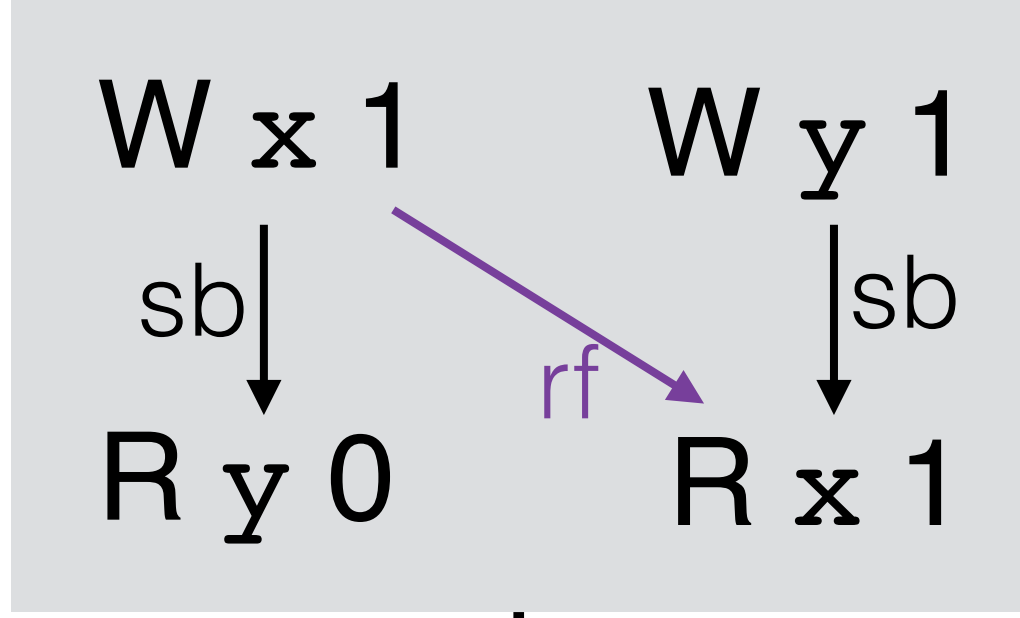
x = 1;   ||   y = 1;
r0 = y;  ||   r1 = x;

```



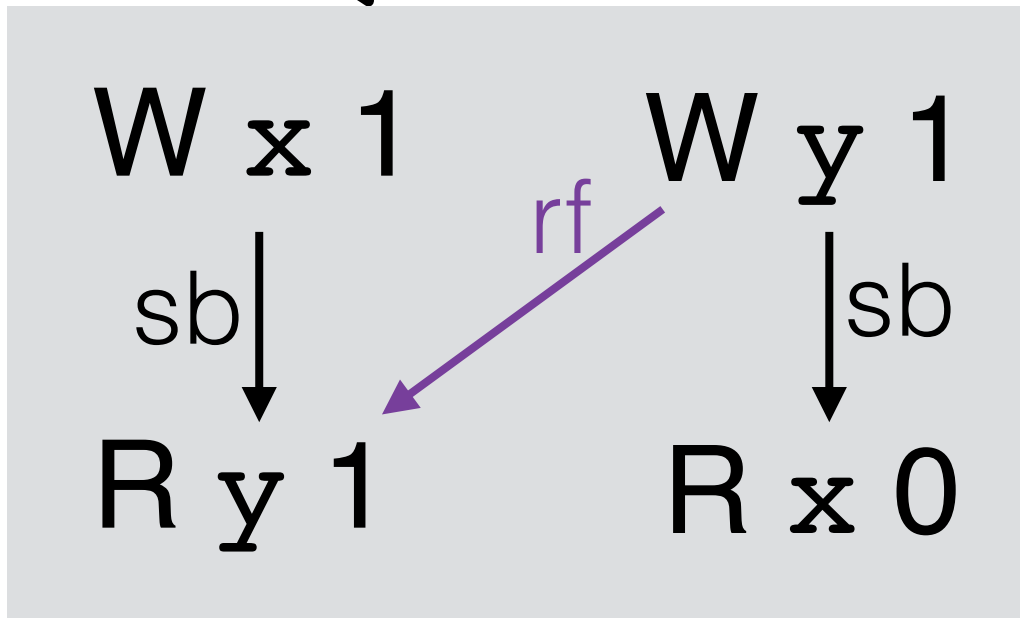
SC: ✓
x86: ✓

r0==1
r1==1



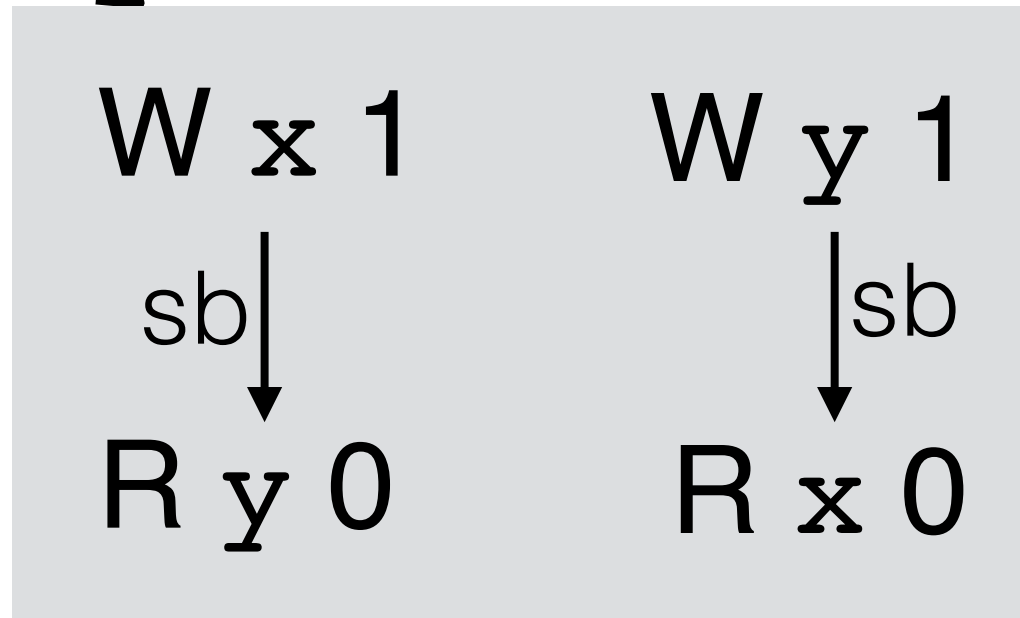
SC: ✓
x86: ✓

r0==0
r1==1



SC: ✓
x86: ✓

r0==1
r1==0

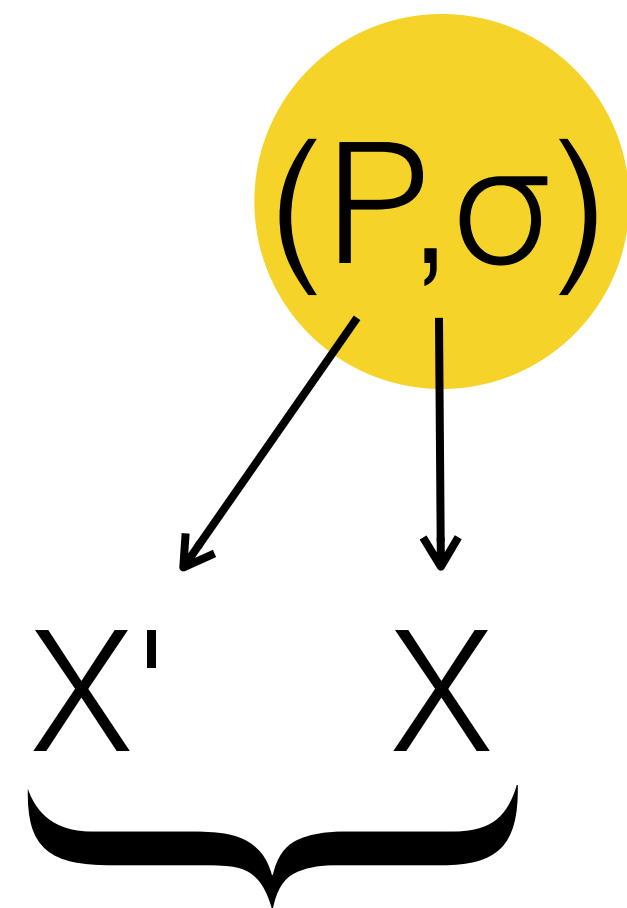


SC: ✗
x86: ✓

r0==0
r1==0

Why is this tricky?

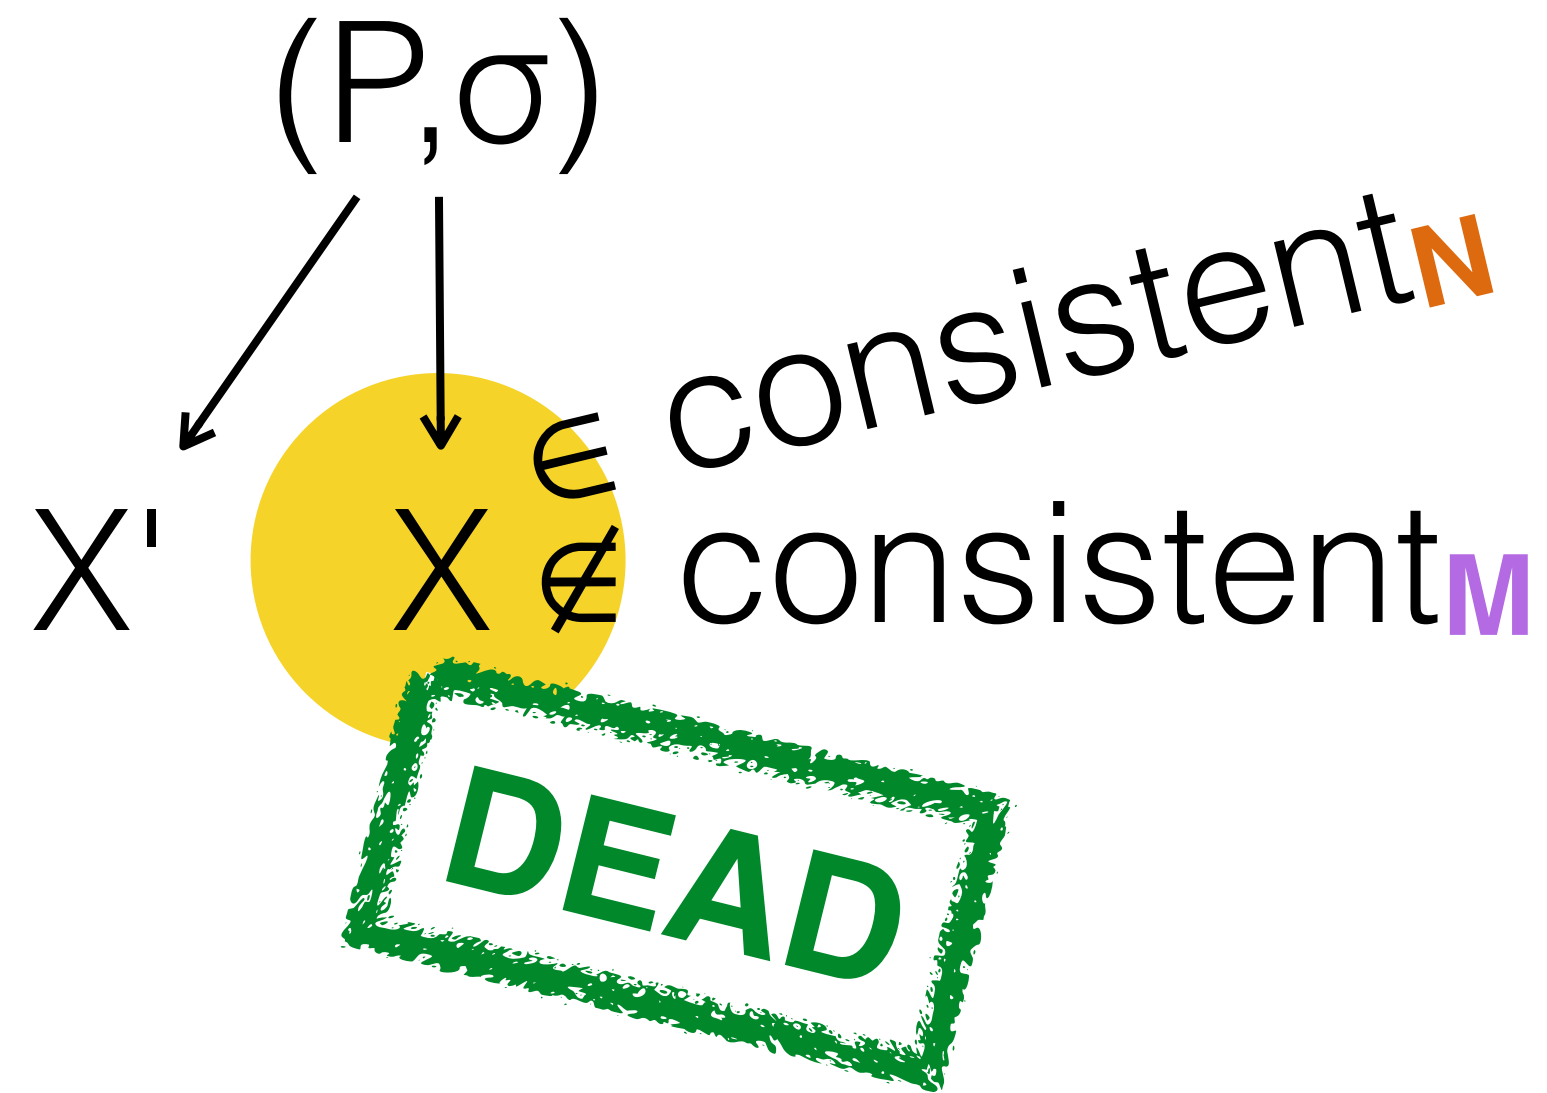
$$\{(P, \sigma) \mid \sigma \notin \llbracket P \rrbracket_M \wedge \sigma \in \llbracket P \rrbracket_N\}$$



all are inconsistent under **M**
some are consistent under **N**

Our solution

$$\{(P, \sigma) \mid \sigma \notin \llbracket P \rrbracket_M \wedge \sigma \in \llbracket P \rrbracket_N\}$$




Generating useful tests

$W_{na} \ x \ 1$ $R_{acq} \ y \ 1$
 $sb \downarrow$ $rf \nearrow$ $sb \downarrow$
 $W_{rel} \ y \ 1$ $R_{na} \ x \ 0$

C++: ✘

NOT DEAD




```
x=1;
st(y, 1, REL);
||
r0=ld(y, ACQ);
r1=x;
```

$W_{na} \ x \ 1$ $R_{acq} \ y \ 1$
 $sb \downarrow$ $rf \nearrow$ $sb \downarrow \downarrow cd$
 $W_{rel} \ y \ 1$ $R_{na} \ x \ 0$

C++: ✘

DEAD



```
x=1;
st(y, 1, REL);
||
r0=ld(y, ACQ);
if(r0) r1=x;
```

Conclusion

- **In summary:** we use Alloy to compare memory consistency models, and check for compiler bugs.
- **Future work:** how to extend our work to handle more recent styles of model like Promising model.