

# Compilation of OCaml memory model to Power

Egor Namakonov, Anton Podkopaev



NATIONAL RESEARCH  
UNIVERSITY



MAX PLANCK INSTITUTE  
FOR SOFTWARE SYSTEMS

# An execution result is explained by alternating threads

$[x] = [y] = 0$		
$[x] := 1$	$a := [x]$	$b := [y]$
	$[y] := 1$	$c := [x]$

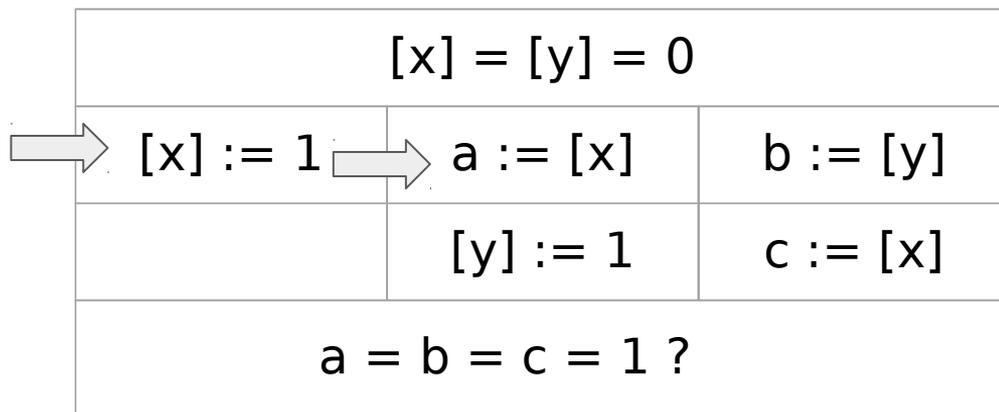
An execution result is explained by alternating threads

$[x] = [y] = 0$		
$[x] := 1$	$a := [x]$	$b := [y]$
	$[y] := 1$	$c := [x]$
$a = b = c = 1 ?$		

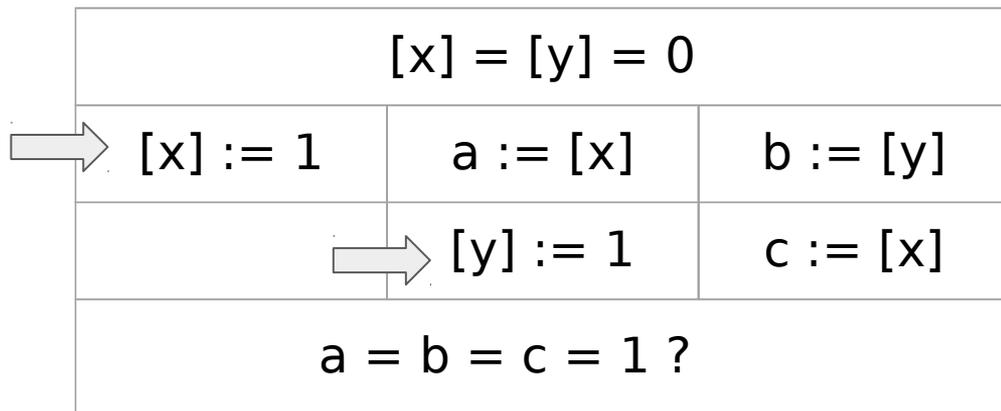
# An execution result is explained by alternating threads

$[x] = [y] = 0$		
 $[x] := 1$	$a := [x]$	$b := [y]$
	$[y] := 1$	$c := [x]$
$a = b = c = 1 ?$		

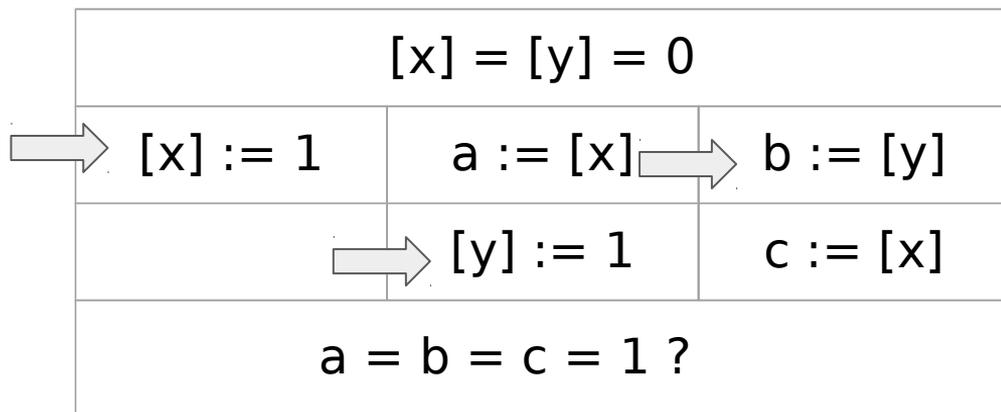
# An execution result is explained by alternating threads



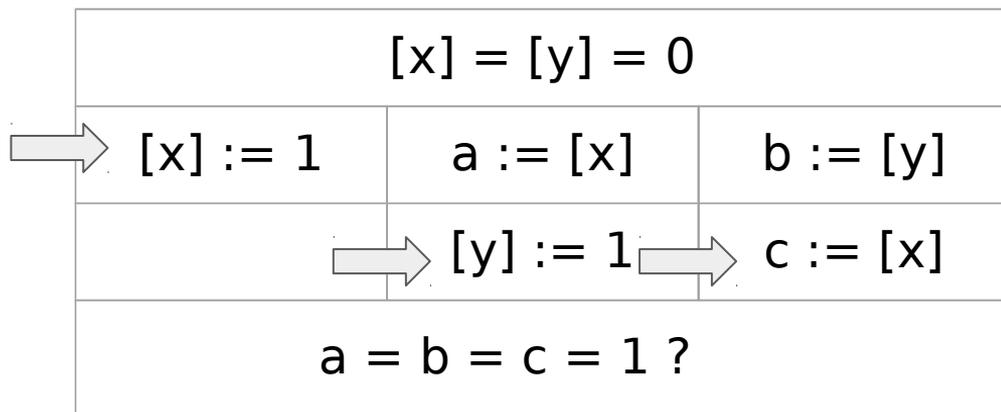
# An execution result is explained by alternating threads



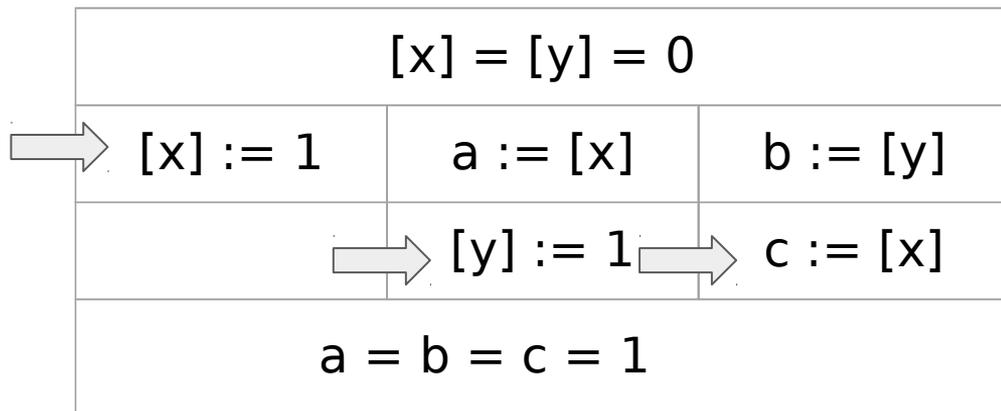
# An execution result is explained by alternating threads



# An execution result is explained by alternating threads



# An execution result is explained by alternating threads



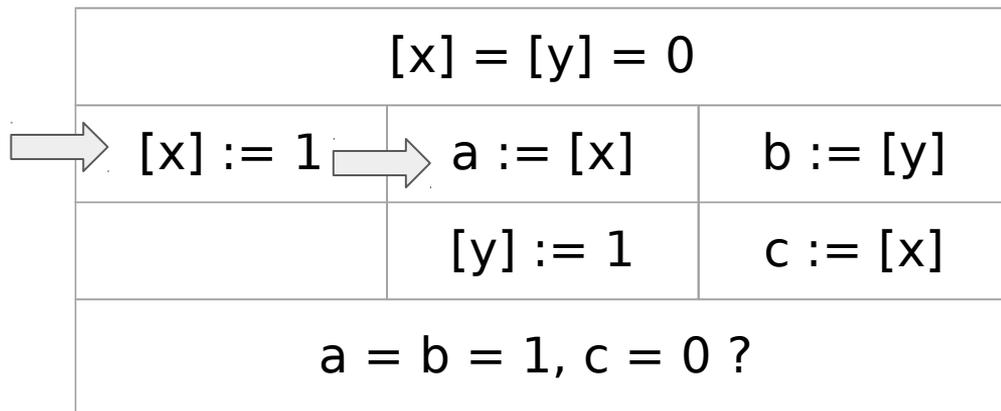
An execution result is explained by alternating threads

$[x] = [y] = 0$		
$[x] := 1$	$a := [x]$	$b := [y]$
	$[y] := 1$	$c := [x]$
$a = b = 1, c = 0 ?$		

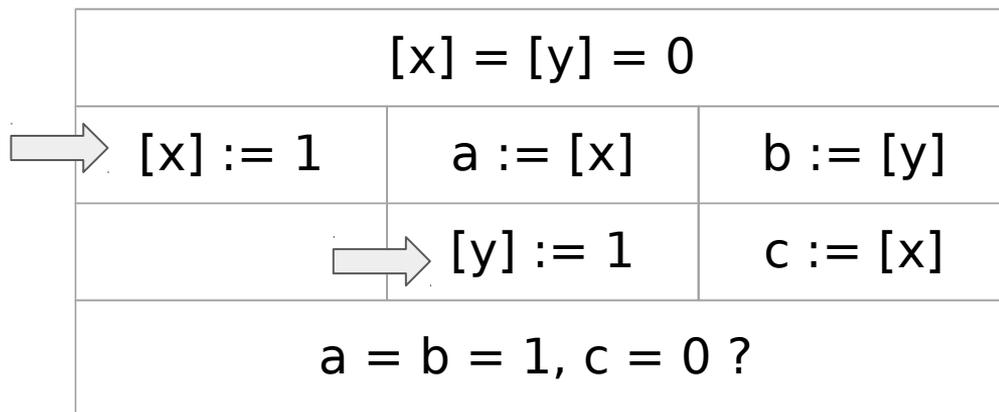
# An execution result is explained by alternating threads

$[x] = [y] = 0$		
 $[x] := 1$	$a := [x]$	$b := [y]$
	$[y] := 1$	$c := [x]$
$a = b = 1, c = 0 ?$		

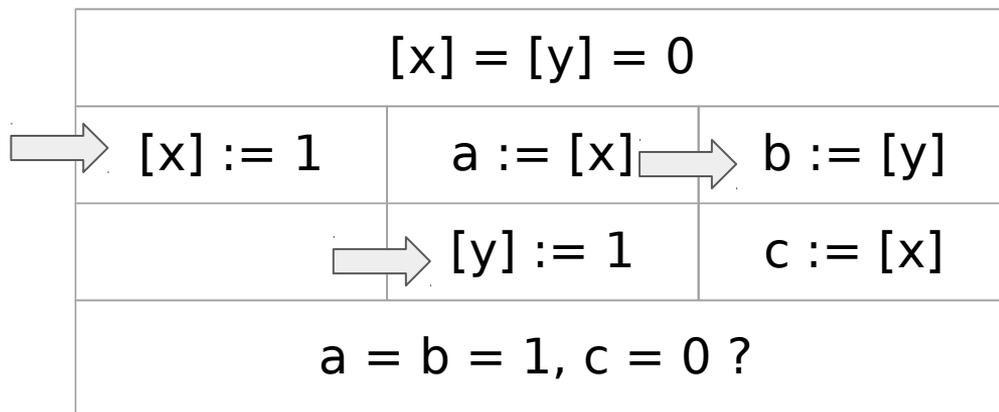
# An execution result is explained by alternating threads



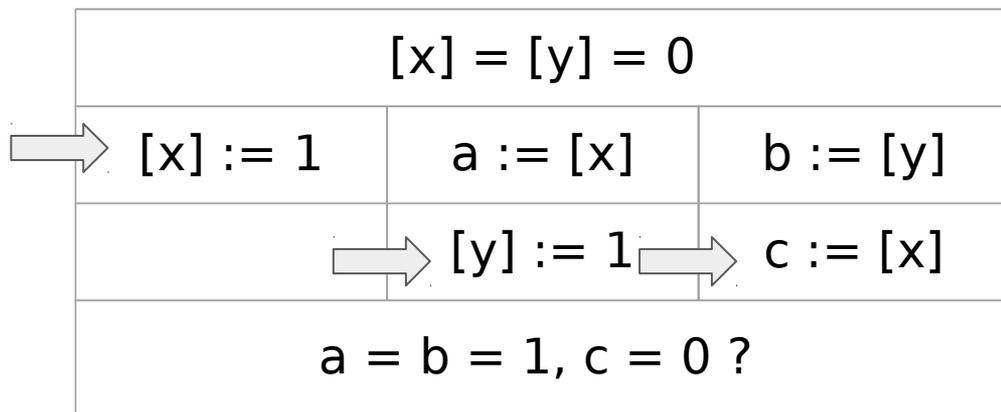
# An execution result is explained by alternating threads



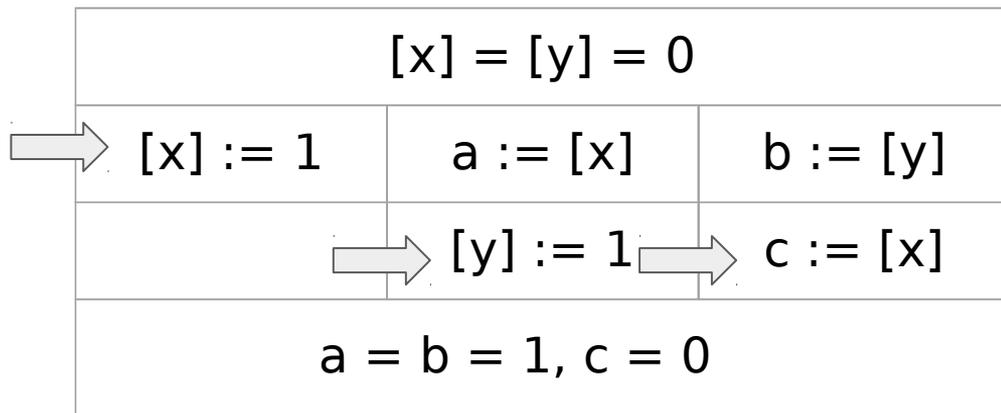
# An execution result is explained by alternating threads



# An execution result is explained by alternating threads



An execution result is explained by alternating threads... usually



C++ allows it due to a (non-atomic) data race

$[x] = [y] = 0$		
$[x] := 1$	$a := [x]$	$b := [y]$
	$[y] := 1$	$c := [x]$
$a = b = 1, c = 0$		

C++ allows it due to a (non-atomic) data race

$[x] = [y] = 0$		
$[x] := 1$	$a := [x]$	$b := [y]$
	$[y] := 1$	$c := [x]$
$a = b = 1, c = 0$		

# C++ allows it due to a (non-atomic) data race

Non-atomic  
accesses

$[x] = [y] = 0$		
$[x] := 1$	$a := [x]$	$b := [y]$
	$[y] := 1$	$c := [x]$
$a = b = 1, c = 0$		

# C++ allows it due to a (non-atomic) data race

Non-atomic  
accesses

$[x] = [y] = 0$		
$[x] := 1$	$a := [x]$	$b := [y]$
	$[y] := 1$	$c := [x]$
$a = b = 1, c = 0$		

**Standard for Programming Language C++, 6.8.2.1.20:**  
“Any such data race results in undefined behavior.”

# No races on atomics

$[x] = [y] = 0$		
$[x]^{\text{rlx}} := 1$	$a := [x]^{\text{rlx}}$	$b := [y]^{\text{rlx}}$
	$[y]^{\text{rlx}} := 1$	$c := [x]^{\text{rlx}}$
$a = b = 1, c = 0 ?$		

# No races on atomics

$[x] = [y] = 0$		
$[x]^{\text{rlx}} := 1$	$a := [x]^{\text{rlx}}$	$b := [y]^{\text{rlx}}$
<b>Constants</b> Defined in header <code>&lt;atomic&gt;</code>	$[y]^{\text{rlx}} := 1$	$c := [x]^{\text{rlx}}$
<b>Value</b>	$a = b = 1, c = 0 ?$	
<code>memory_order_relaxed</code>		

No races on atomics but the outcome is still allowed

$[x] = [y] = 0$		
$[x]^{\text{rlx}} := 1$	$a := [x]^{\text{rlx}}$	$b := [y]^{\text{rlx}}$
	$[y]^{\text{rlx}} := 1$	$c := [x]^{\text{rlx}}$
$a = b = 1, c = 0$		

# C++ memory model

$[x] = [y] = 0$		
$[x]^{\text{rlx}} := 1$	$a := [x]^{\text{rlx}}$	$b := [y]^{\text{rlx}}$
	$[y]^{\text{rlx}} := 1$	$c := [x]^{\text{rlx}}$

# C++ memory model

$[x] = [y] = 0$		
$[x]^{\text{rlx}} := 1$	$a := [x]^{\text{rlx}}$	$b := [y]^{\text{rlx}}$
	$[y]^{\text{rlx}} := 1$	$c := [x]^{\text{rlx}}$

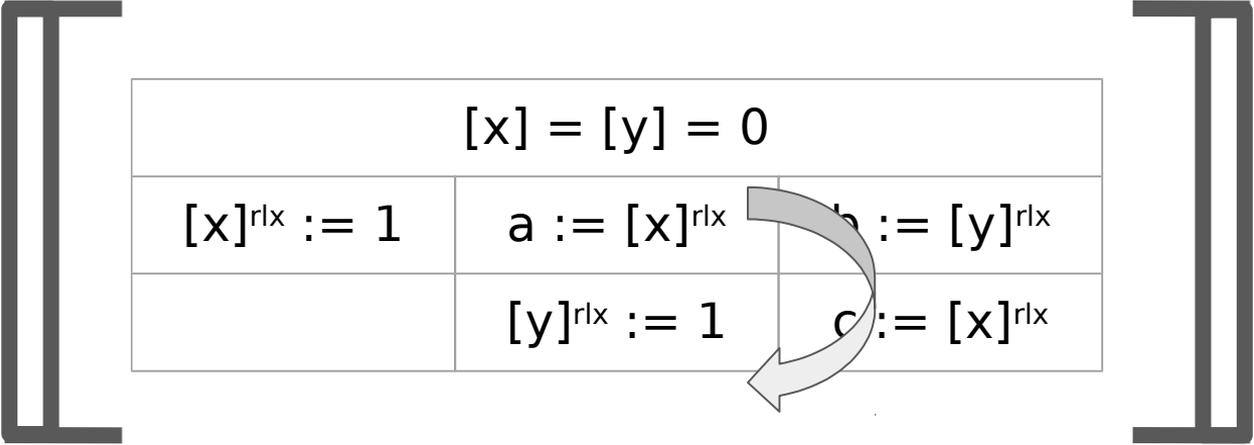
$= \{ \dots, (a=b=c=1), \dots (a=b=1, c=0), \dots \}$

# C++ memory model is **weak**

$[x] = [y] = 0$		
$[x]^{\text{rlx}} := 1$	$a := [x]^{\text{rlx}}$	$b := [y]^{\text{rlx}}$
	$[y]^{\text{rlx}} := 1$	$c := [x]^{\text{rlx}}$

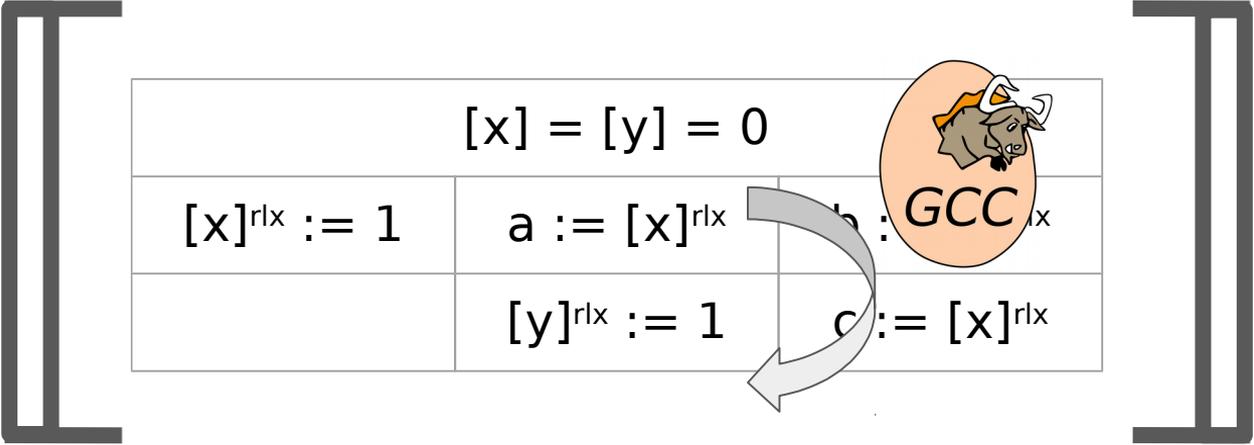
= { ..., (a=b=c=1), ... (a=b=1, c=0), ... }

C++ memory model is **weak** as it allows optimizations



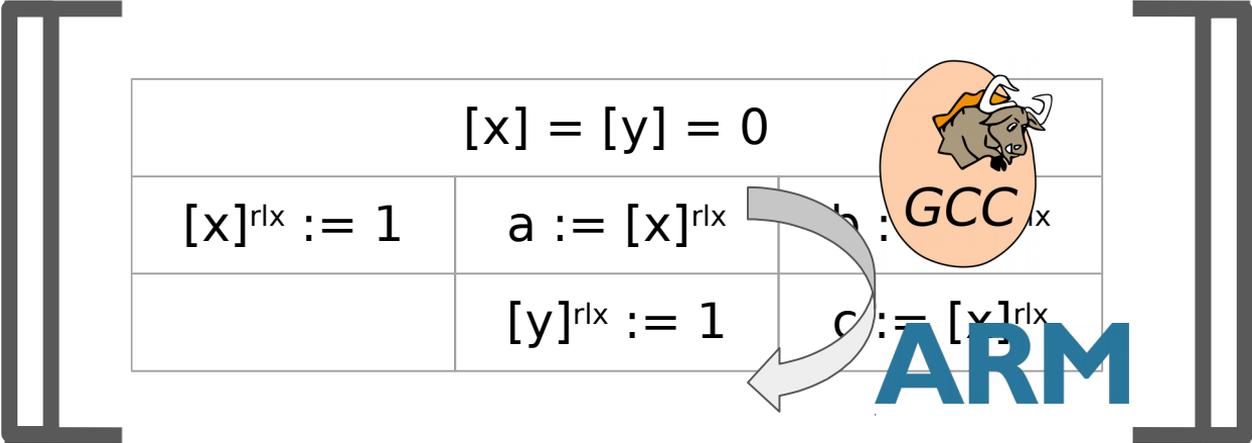
= { ..., (a=b=c=1), ... (a=b=1, c=0), ... }

# C++ memory model is **weak** as it allows optimizations



= { ..., (a=b=c=1), ... (a=b=1, c=0), ... }

# C++ memory model is **weak** as it allows optimizations

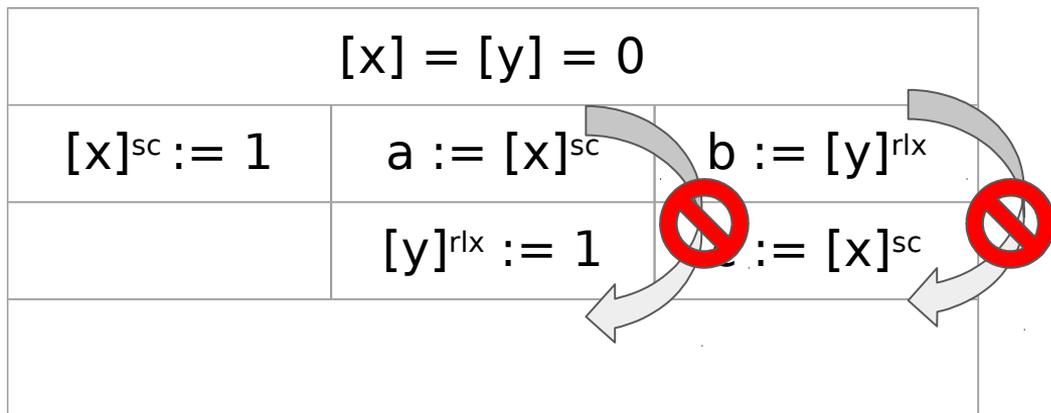


= { ..., (a=b=c=1), ... (a=b=1, c=0), ... }

# Weak behavior can be controlled with access modes

$[x] = [y] = 0$		
$[x]^{sc} := 1$	$a := [x]^{sc}$	$b := [y]^{rlx}$
	$[y]^{rlx} := 1$	$c := [x]^{sc}$

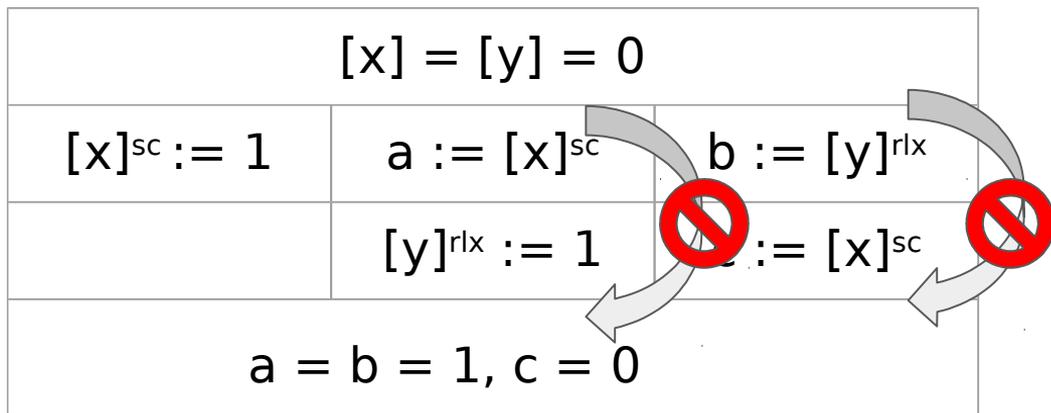
# Weak behavior can be controlled with access modes



Weak behavior can be controlled with access modes  
but the effect is not obvious



POWER



# C++ solution: strengthen access mode everywhere

$[x] = [y] = 0$		
$[x]^{sc} := 1$	$a := [x]^{sc}$	$b := [y]^{sc}$
	$[y]^{sc} := 1$	$c := [x]^{sc}$
<del><math>a = b = 1, c = 0</math></del>		

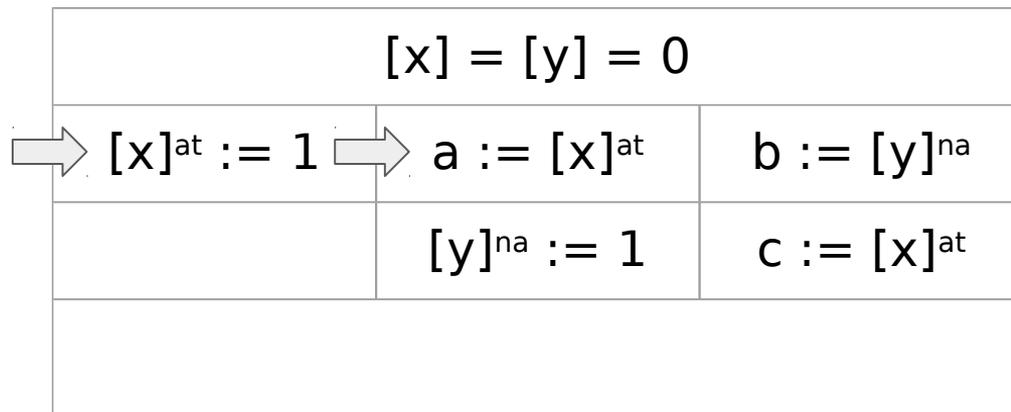
# OCaml MM: reasonable rules for racy programs

$[x] = [y] = 0$		
$[x]^{\text{at}} := 1$	$a := [x]^{\text{at}}$	$b := [y]^{\text{na}}$
	$[y]^{\text{na}} := 1$	$c := [x]^{\text{at}}$

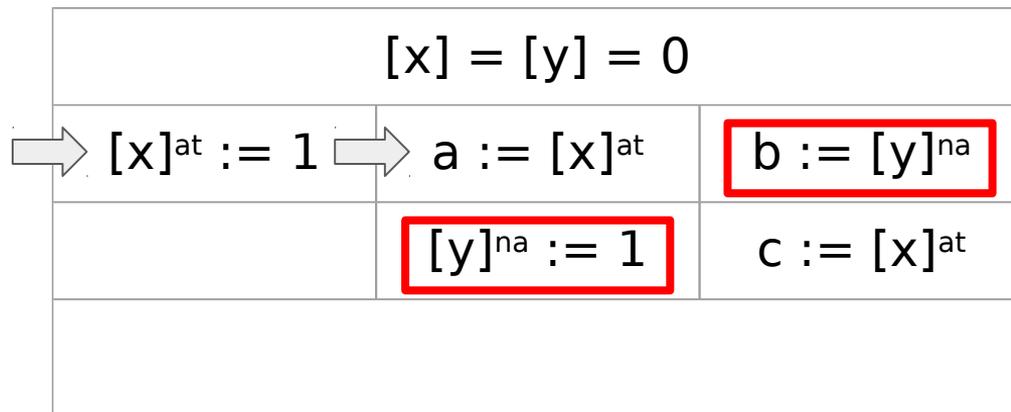
# OCaml MM: reasonable rules for racy programs

$[x] = [y] = 0$		
 $[x]^{\text{at}} := 1$	$a := [x]^{\text{at}}$	$b := [y]^{\text{na}}$
	$[y]^{\text{na}} := 1$	$c := [x]^{\text{at}}$

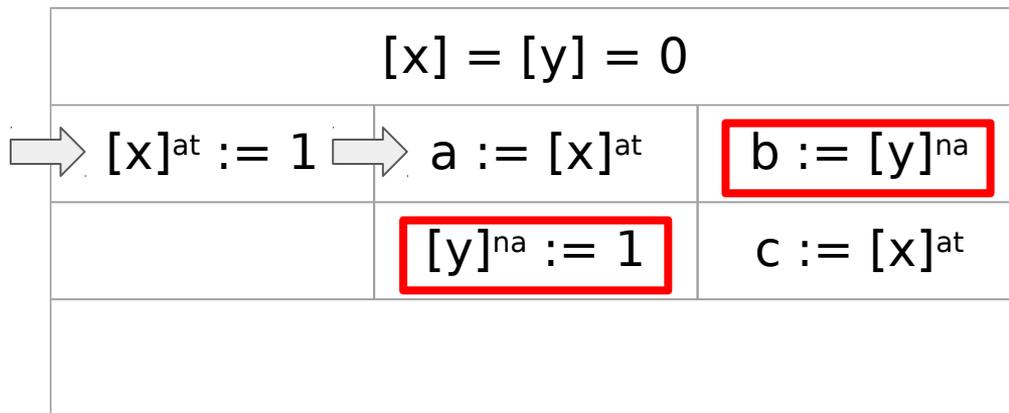
# OCaml MM: reasonable rules for racy programs



# OCaml MM: reasonable rules for racy programs

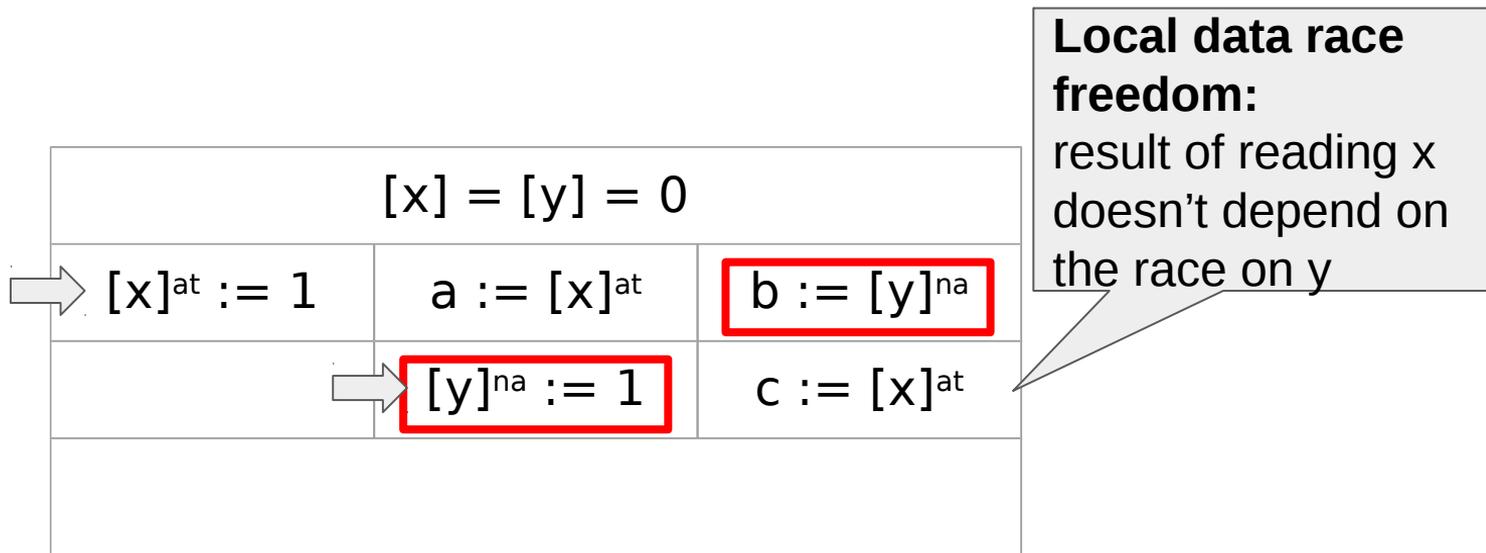


# OCaml MM: reasonable rules for racy programs

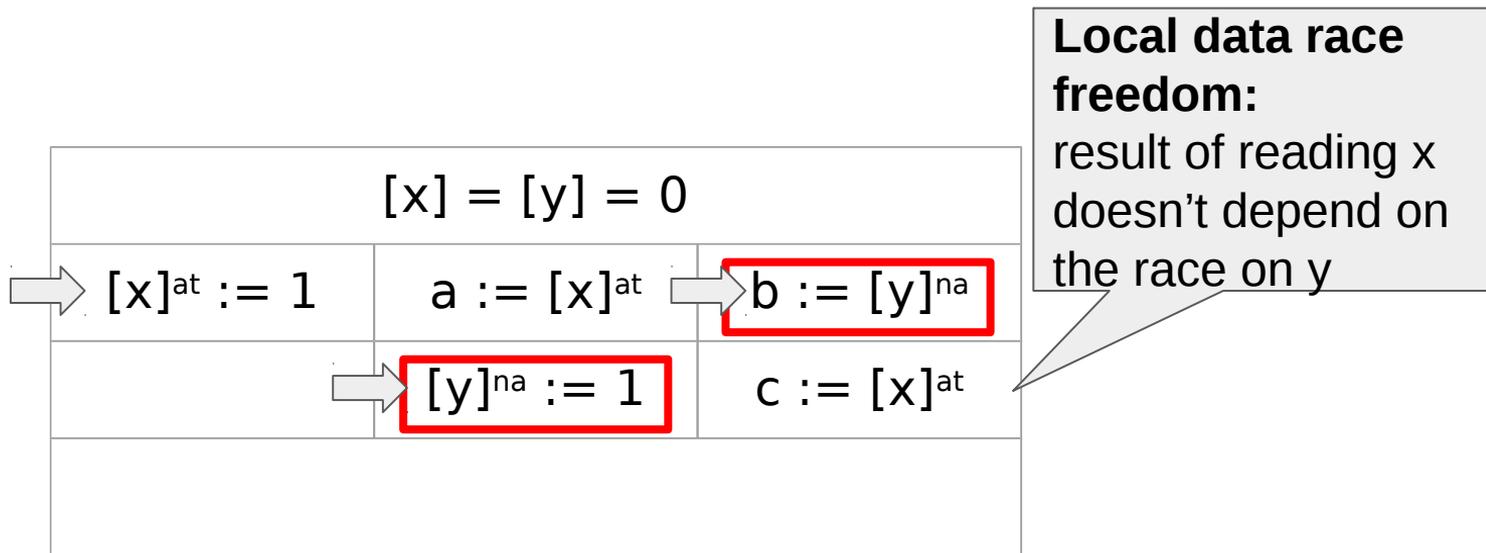


**Local data race freedom:**  
result of reading  $x$   
doesn't depend on  
the race on  $y$

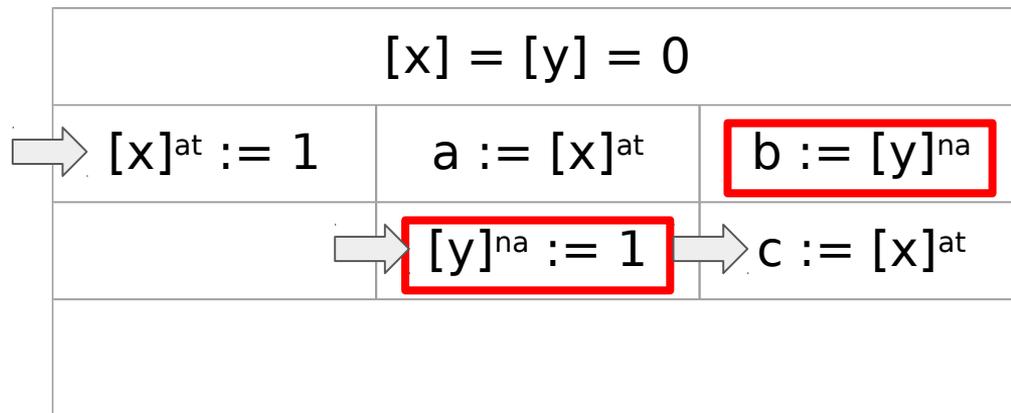
# OCaml MM: reasonable rules for racy programs



# OCaml MM: reasonable rules for racy programs

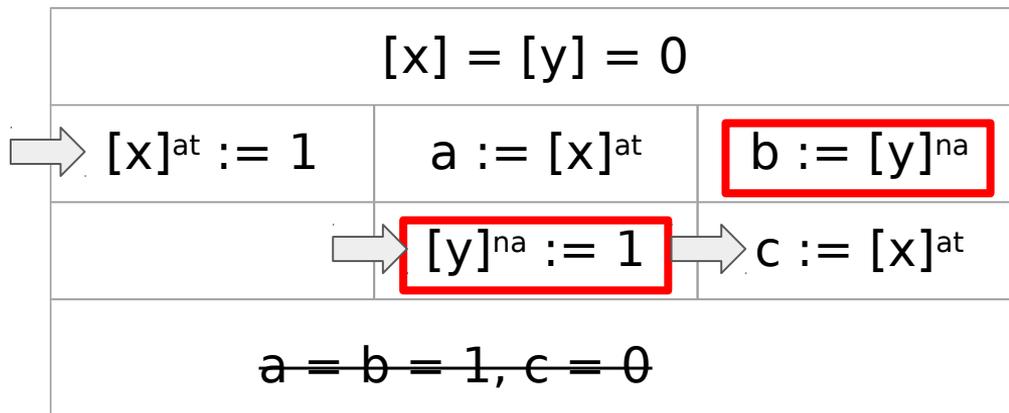


# OCaml MM: reasonable rules for racy programs



**Local data race freedom:**  
result of reading  $x$   
doesn't depend on  
the race on  $y$

# OCaml MM: reasonable rules for racy programs



**Local data race freedom:**  
result of reading  $x$   
doesn't depend on  
the race on  $y$

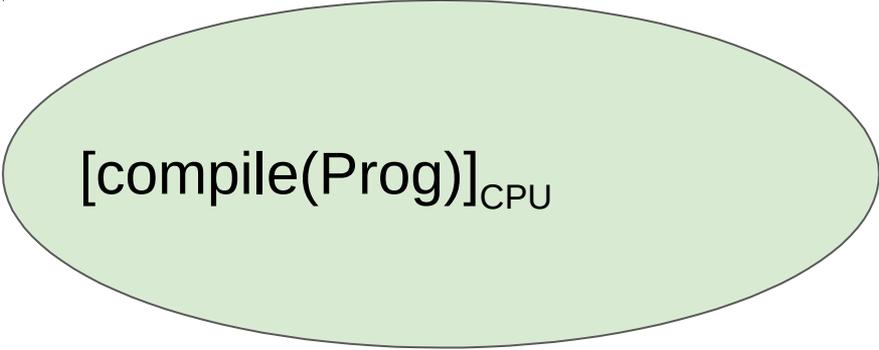
# OCaml MM guarantees should be implemented



POWER

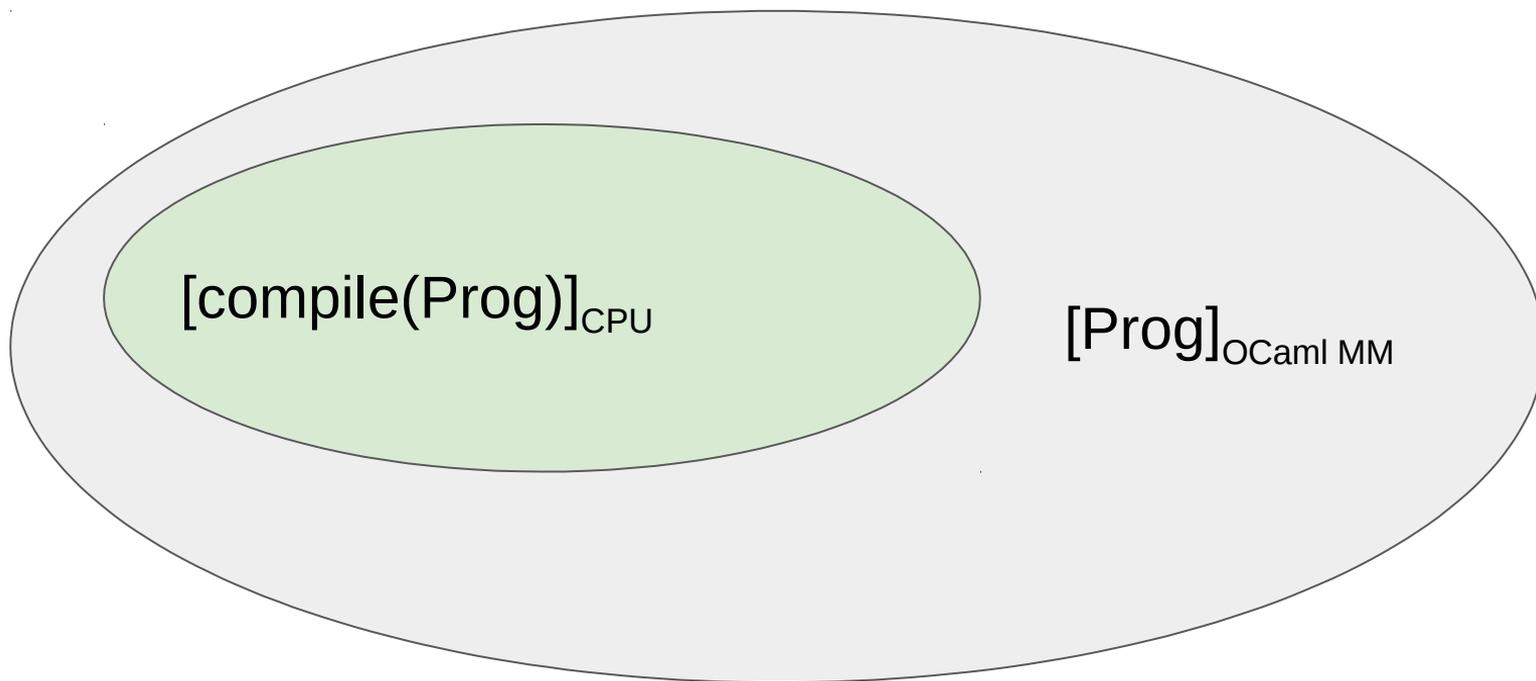
$[x] = [y] = 0$		
$[x]^{sc} := 1$	$a := [x]^{sc}$	$b := [y]^{rlx}$
	$[y]^{rlx} := 1$	$c := [x]^{sc}$
$a = b = 1, c = 0$		

OCaml MM guarantees should be implemented

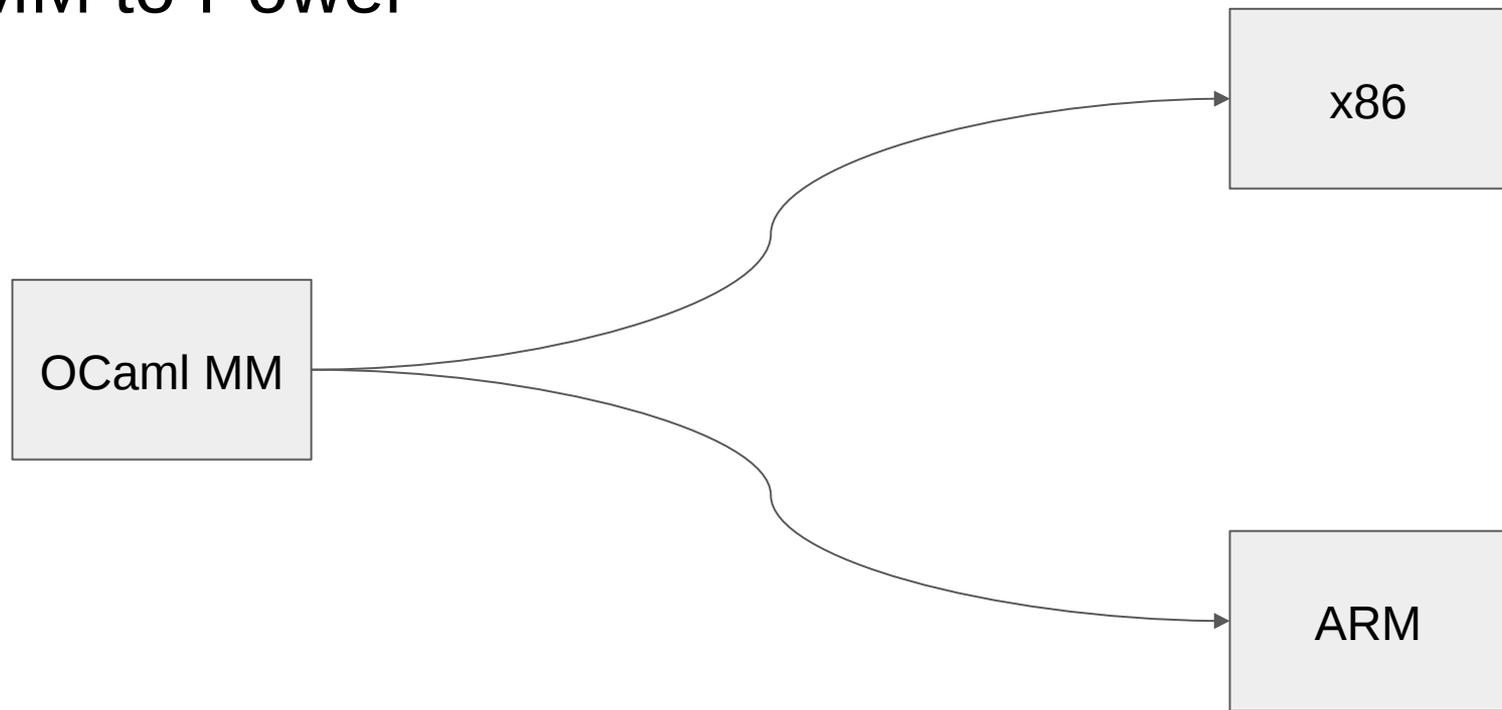


$[\text{compile}(\text{Prog})]_{\text{CPU}}$

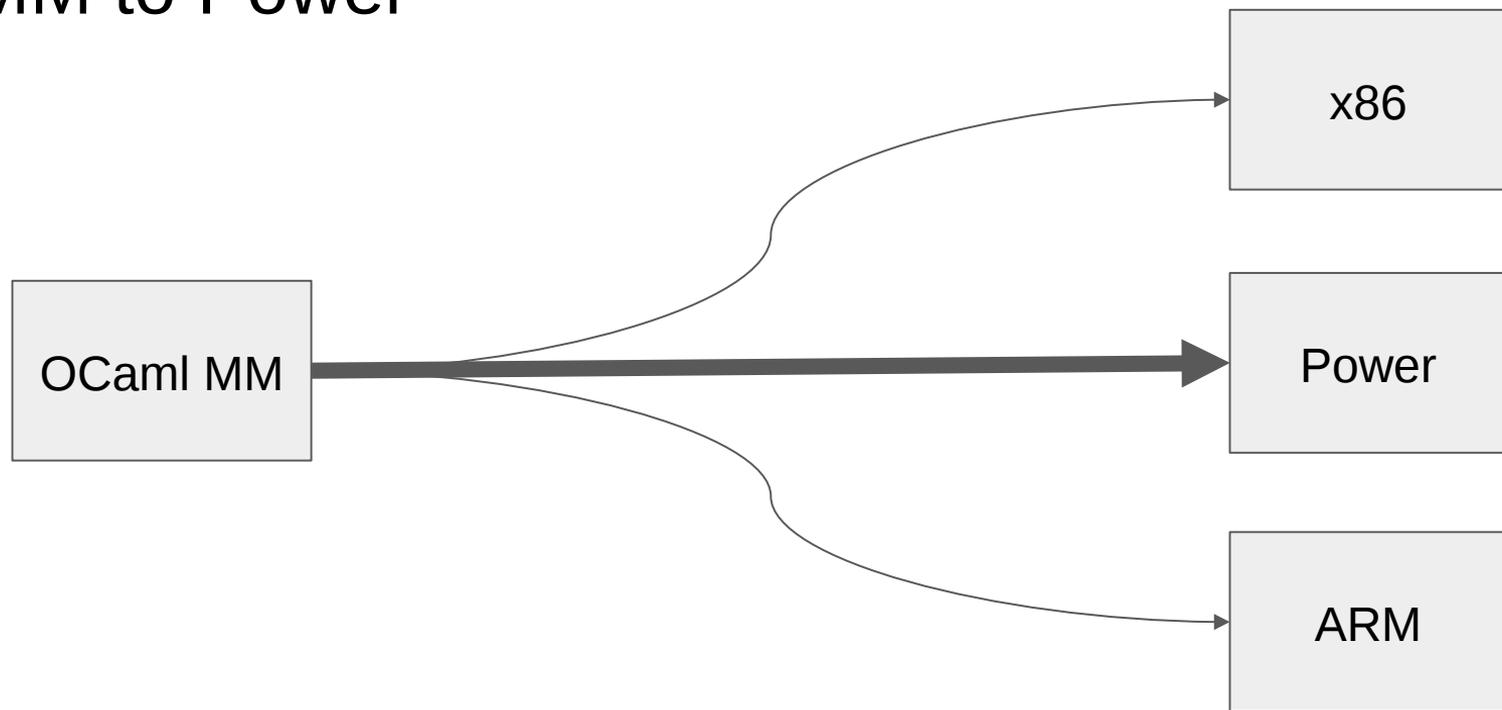
OCaml MM guarantees should be implemented by providing a correct compilation scheme



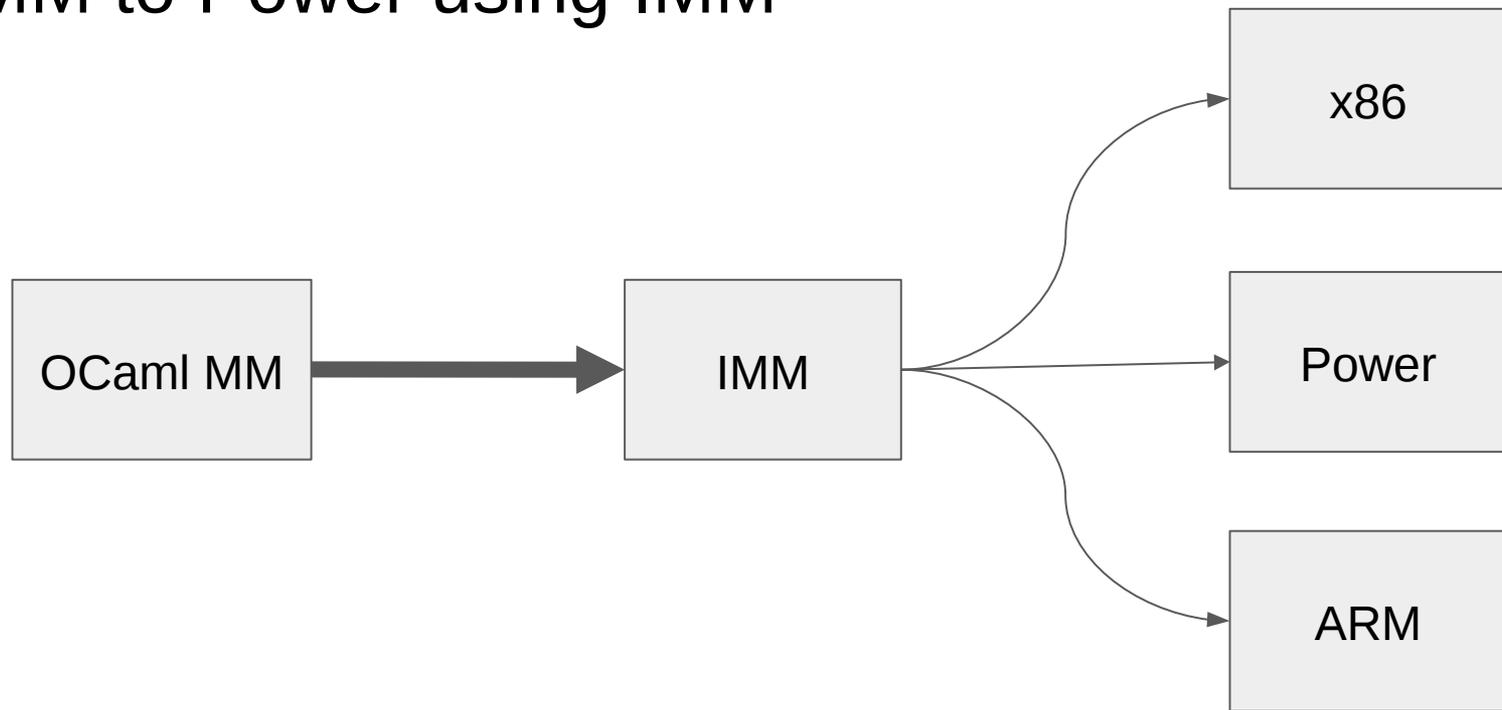
# We've proved compilation correctness from OCaml MM to Power



# We've proved compilation correctness from OCaml MM to Power



# We've proved compilation correctness from OCaml MM to Power using IMM

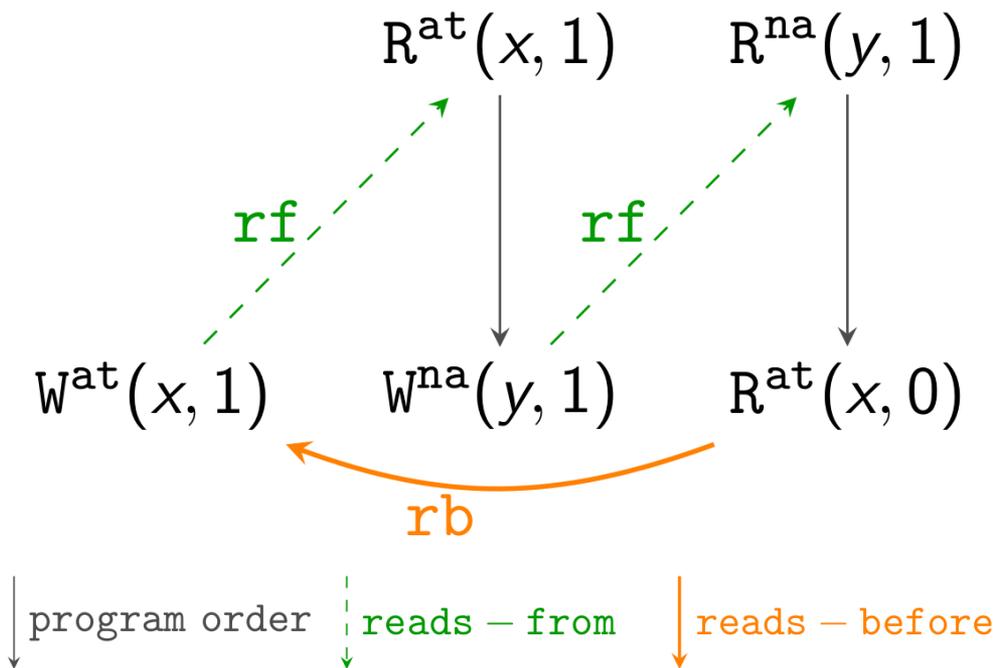


# Another execution representation is needed

$[x] = [y] = 0$		
$[x] := 1$	$a := [x]$	$b := [y]$
	$[y] := 1$	$c := [x]$
$a = b = 1, c = 0$		

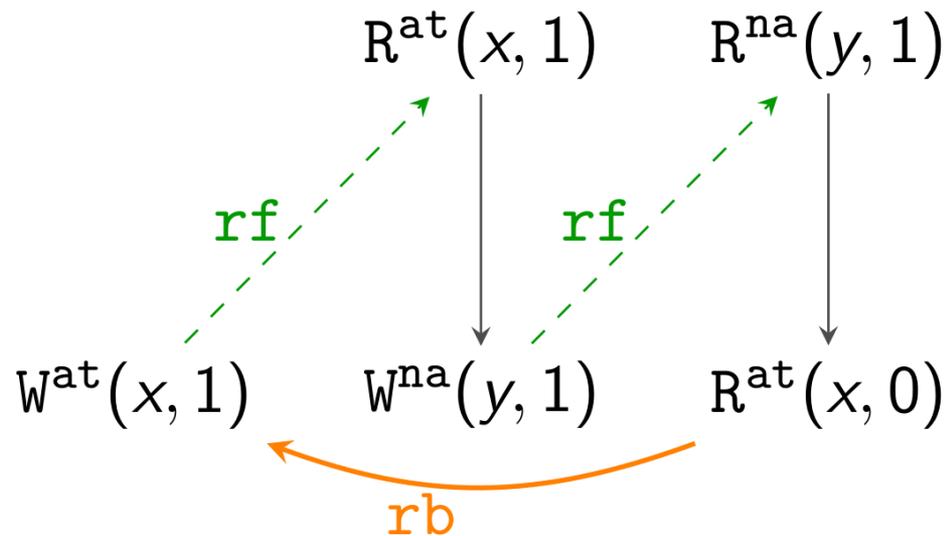
# Consider the execution as a graph

$[x] = [y] = 0$		
$[x]^{\text{at}} := 1$	$a := [x]^{\text{at}}$	$b := [y]^{\text{na}}$
	$[y]^{\text{na}} := 1$	$c := [x]^{\text{at}}$
$a = b = 1, c = 0$		



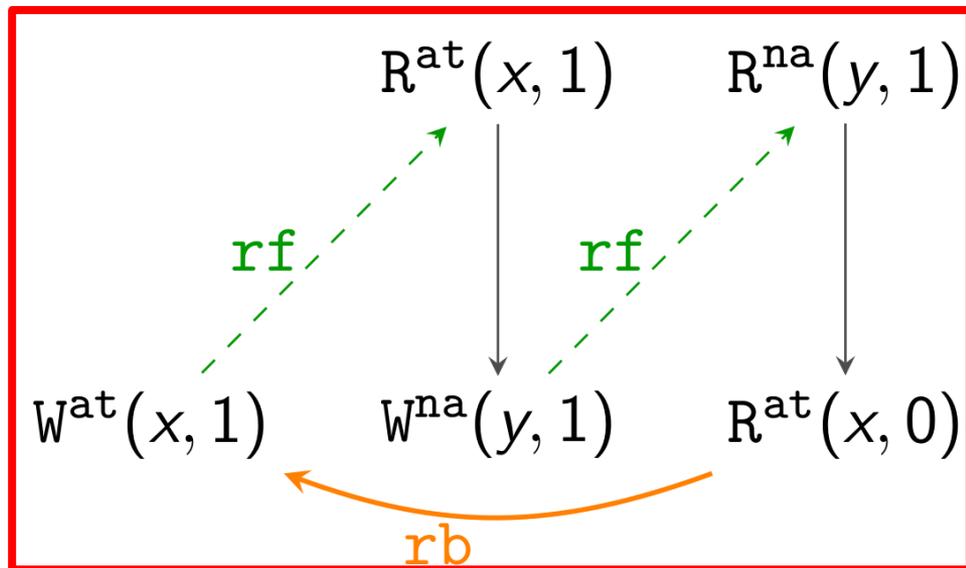
# A permission of execution is determined by its graph

$[x] = [y] = 0$		
$[x]^{\text{at}} := 1$	$a := [x]^{\text{at}}$	$b := [y]^{\text{na}}$
	$[y]^{\text{na}} := 1$	$c := [x]^{\text{at}}$



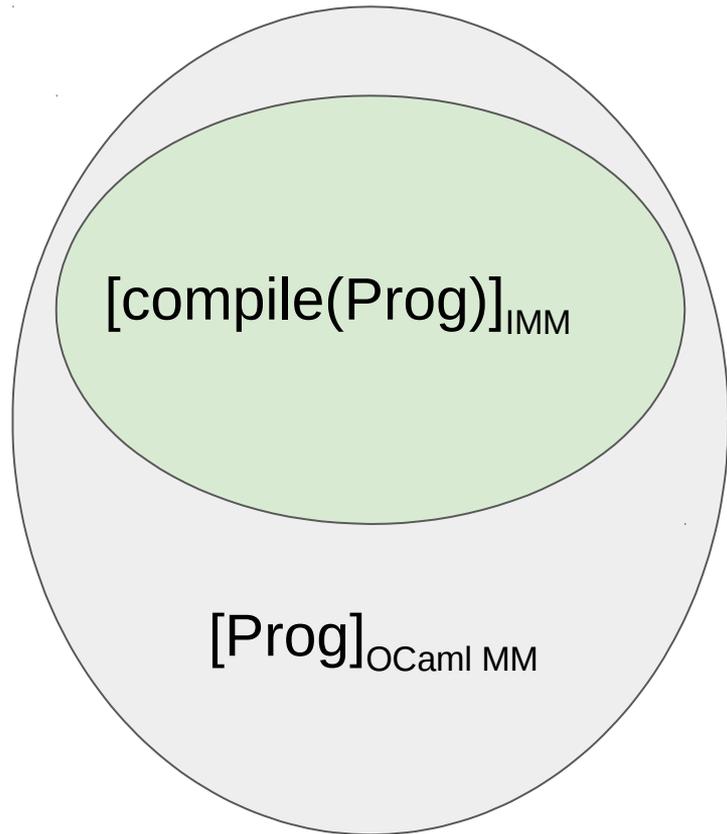
A permission of execution is determined by its graph

$[x] = [y] = 0$		
$[x]^{\text{at}} := 1$	$a := [x]^{\text{at}}$	$b := [y]^{\text{na}}$
	$[y]^{\text{na}} := 1$	$c := [x]^{\text{at}}$
OCaml MM: <del><math>a = b = 1, c = 0</math></del>		

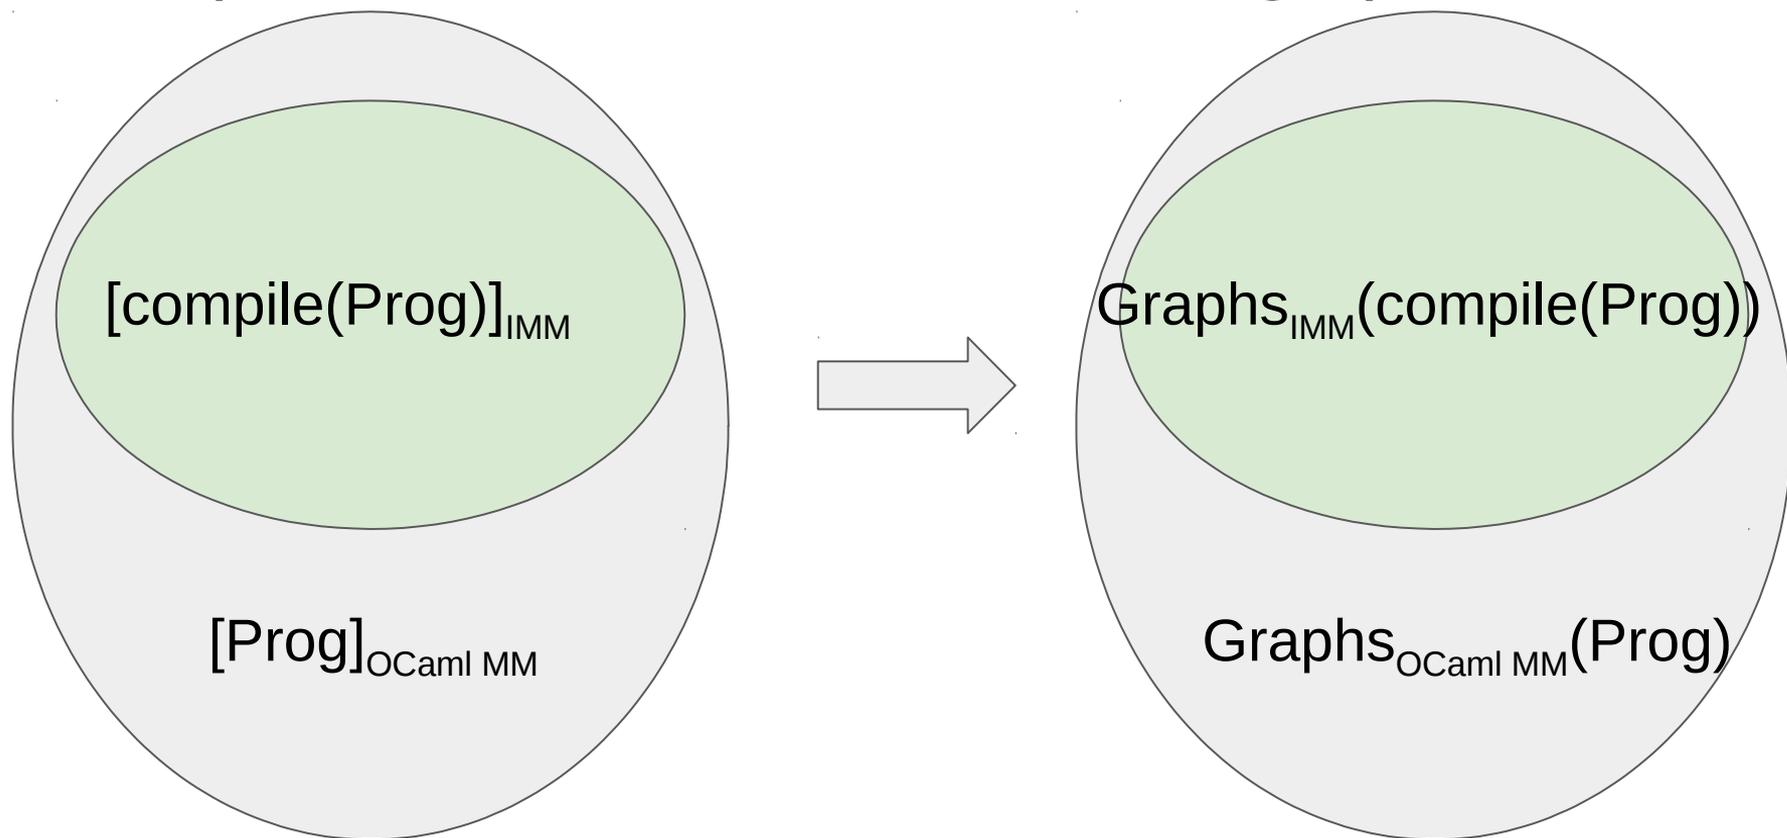


OMM: no cycles made of **po**, **rf** and **rb**

# Compilation correctness in terms of graphs



# Compilation correctness in terms of graphs



# The identity compilation scheme won't work

$\text{compile}(\text{Prog}) = [\text{na} \rightarrow \text{rlx}, \text{at} \rightarrow \text{sc}]\text{Prog}$

# The identity compilation scheme won't work

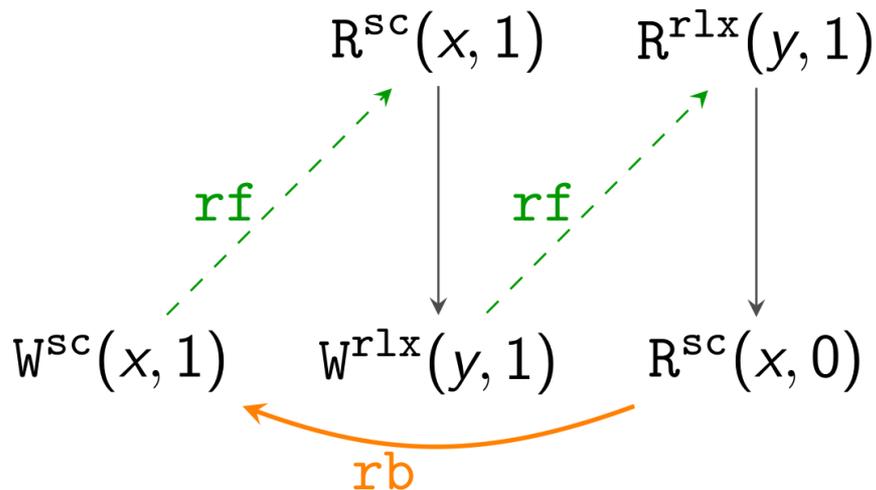
$\text{compile}(\text{Prog}) = [\text{na} \rightarrow \text{rlx}, \text{at} \rightarrow \text{sc}]\text{Prog}$

$[x] = [y] = 0$		
$[x]^{\text{sc}} := 1$	$a := [x]^{\text{sc}}$	$b := [y]^{\text{rlx}}$
	$[y]^{\text{rlx}} := 1$	$c := [x]^{\text{sc}}$

# The identity compilation scheme won't work

$\text{compile}(\text{Prog}) = [\text{na} \rightarrow \text{rlx}, \text{at} \rightarrow \text{sc}]\text{Prog}$

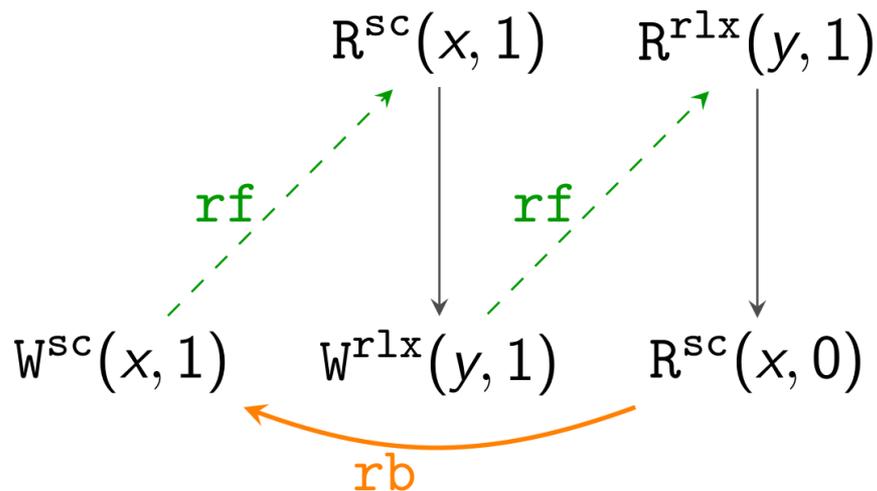
$[x] = [y] = 0$		
$[x]^{\text{sc}} := 1$	$a := [x]^{\text{sc}}$	$b := [y]^{\text{rlx}}$
	$[y]^{\text{rlx}} := 1$	$c := [x]^{\text{sc}}$



# The identity compilation scheme won't work

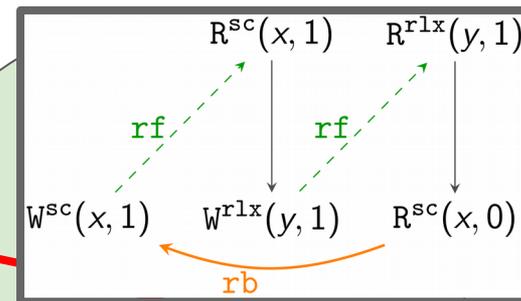
$\text{compile}(\text{Prog}) = [\text{na} \rightarrow \text{rlx}, \text{at} \rightarrow \text{sc}]\text{Prog}$

$[x] = [y] = 0$		
$[x]^{\text{sc}} := 1$	$a := [x]^{\text{sc}}$	$b := [y]^{\text{rlx}}$
	$[y]^{\text{rlx}} := 1$	$c := [x]^{\text{sc}}$
$a = b = 1, c = 0$		



IMM: can have a cycle made of **po**, **rf** and **rb**

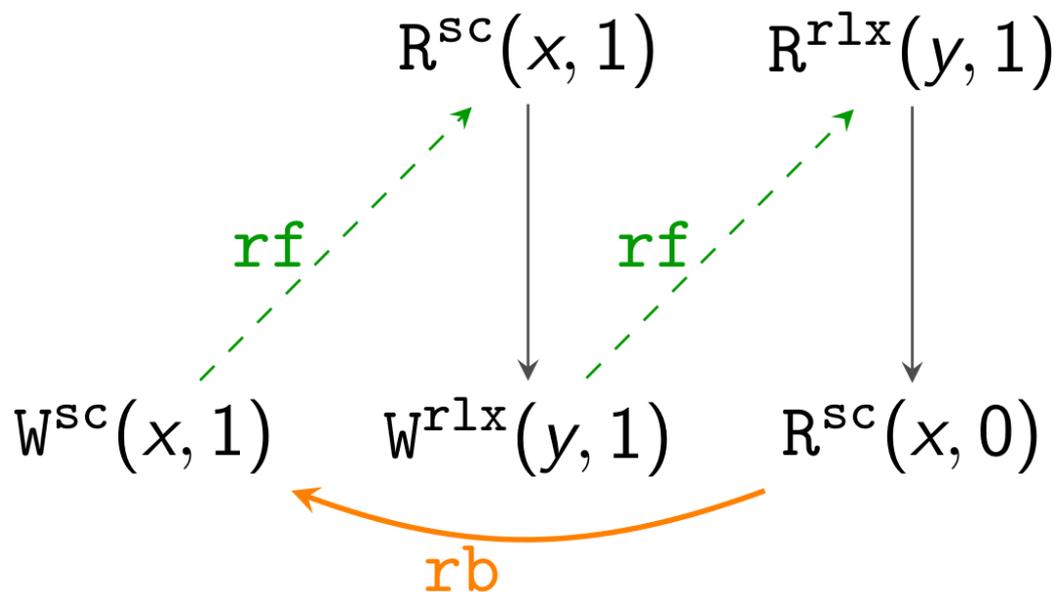
# The identity compilation scheme won't work



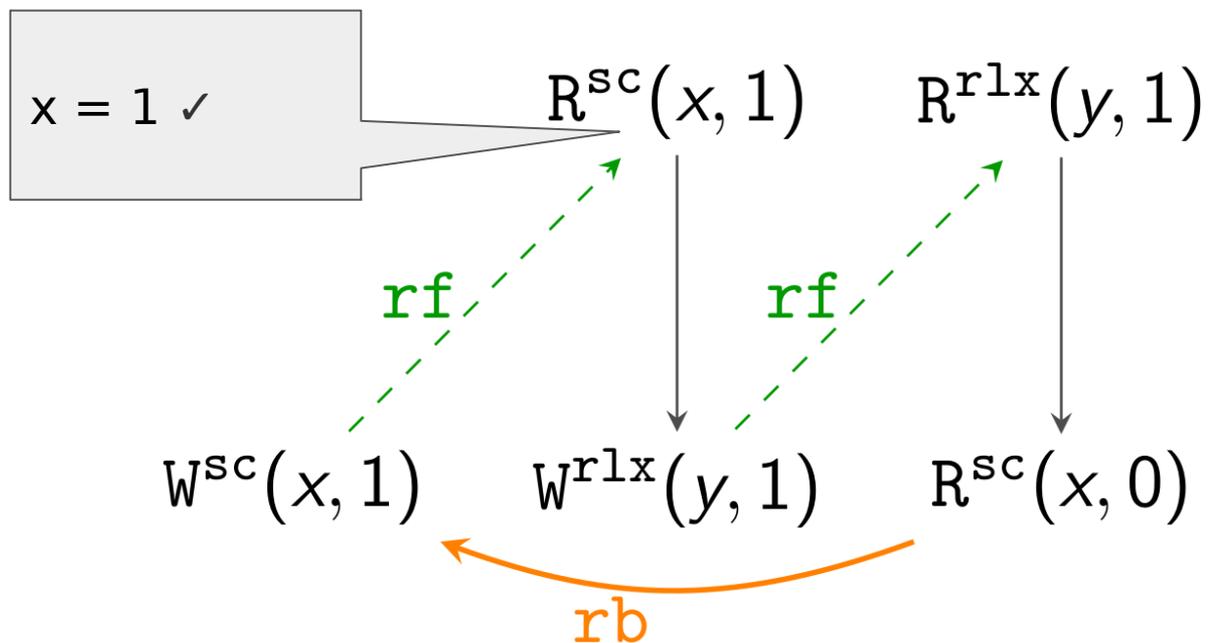
$\text{Graphs}_{\text{OCaml MM}}(\text{Prog})$

$\text{Graphs}_{\text{IMM}}(\text{compile}(\text{Prog}))$

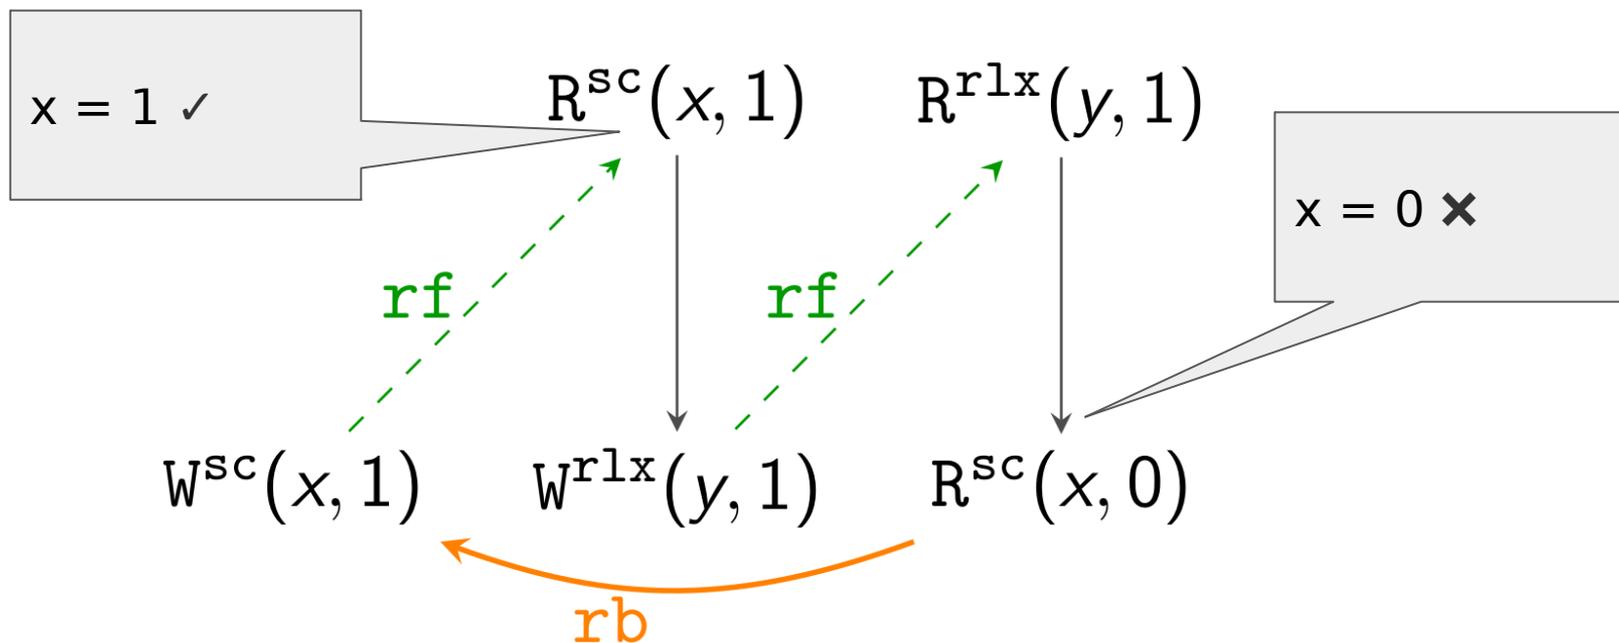
Observed writes should remain so



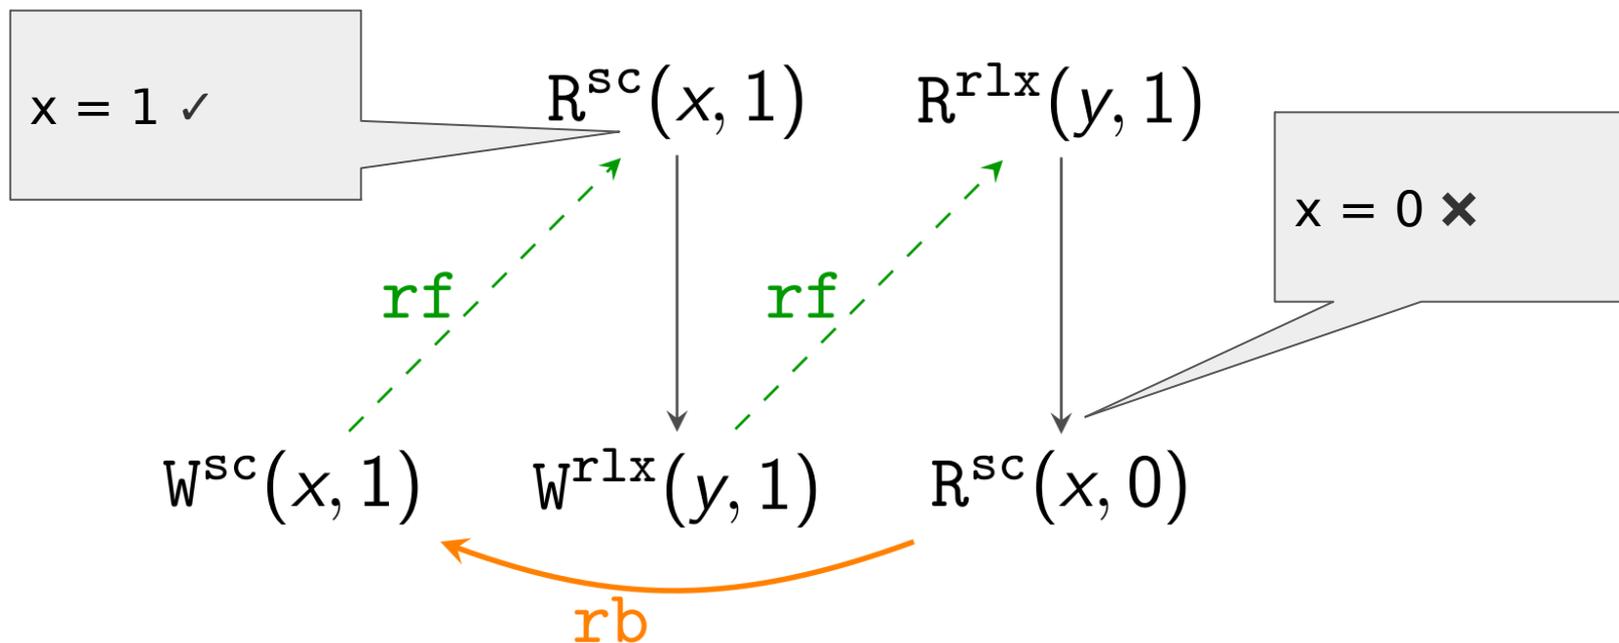
Observed writes should remain so



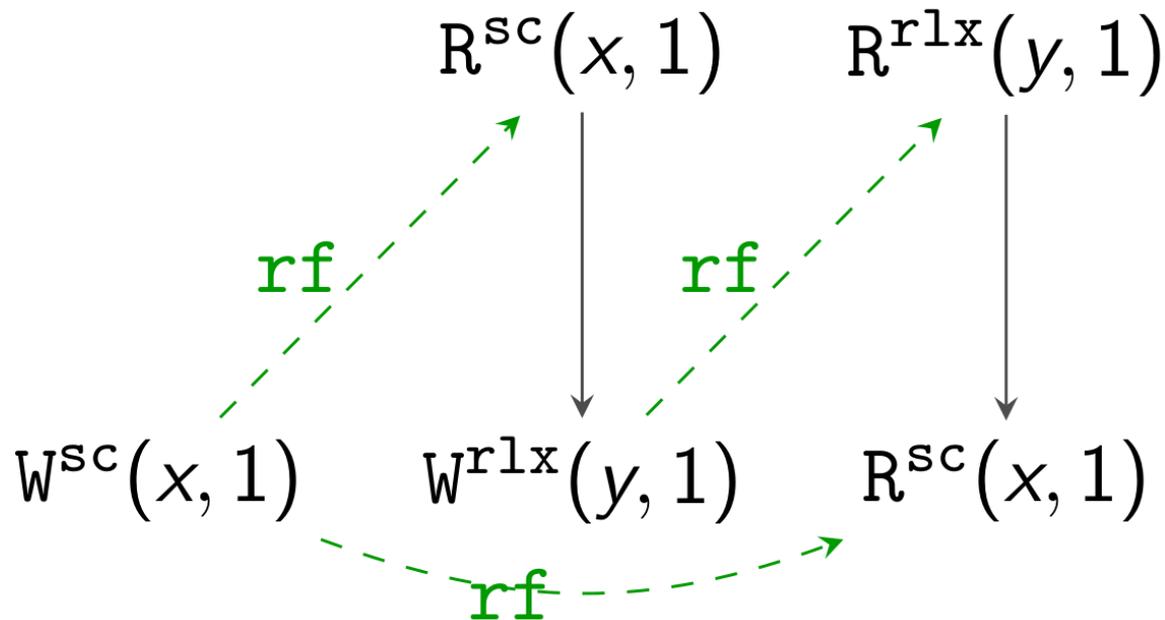
# Observed writes should remain so



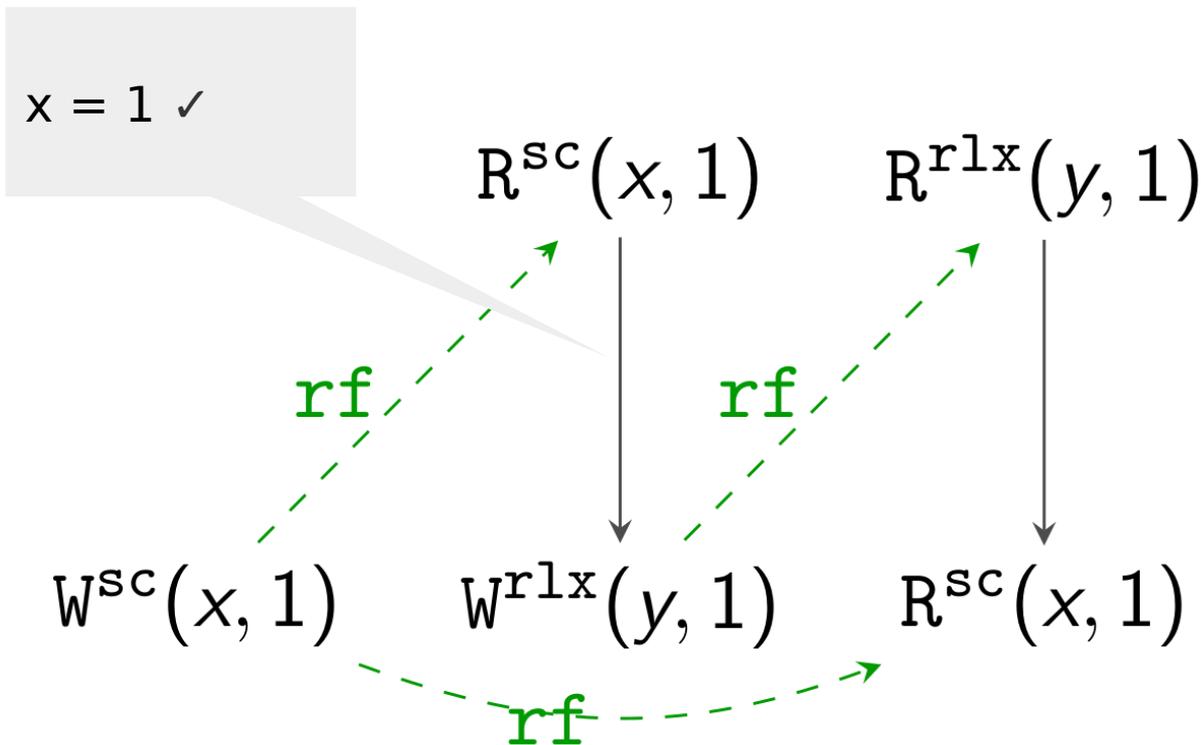
Observed writes should remain so  
= graph should have no cycles with rb



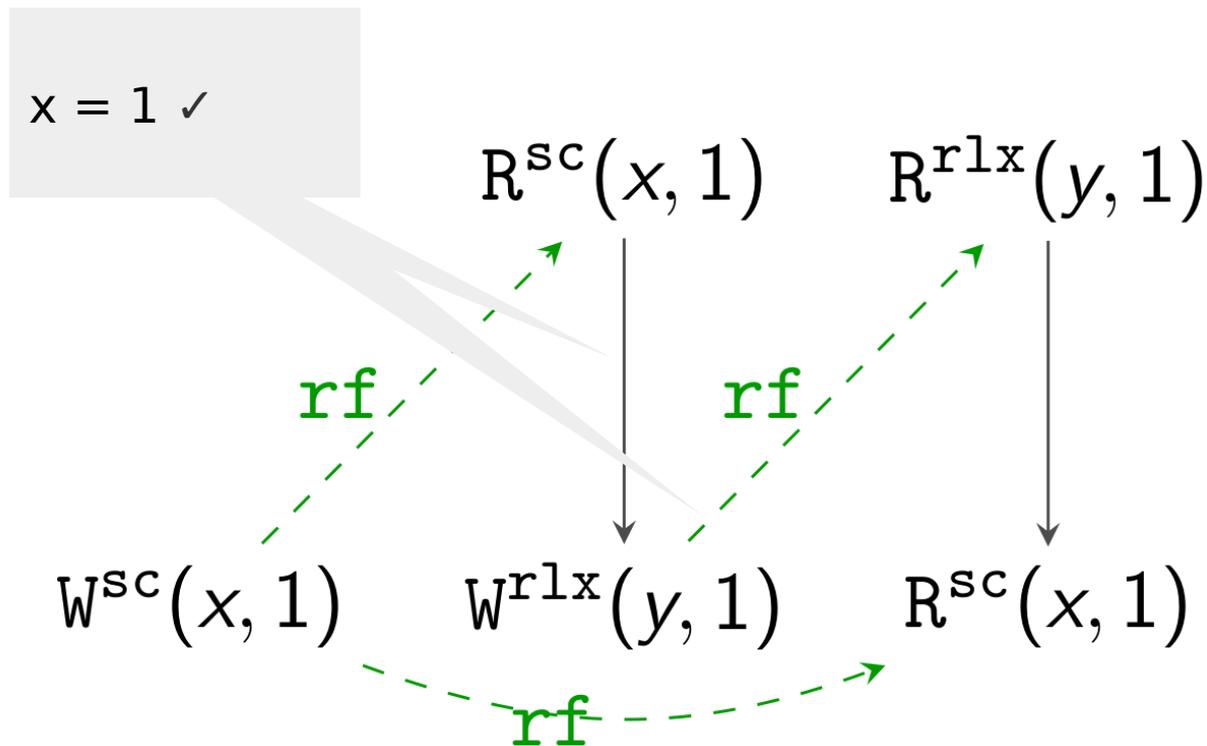
“Release” known writes and “acquire” them next



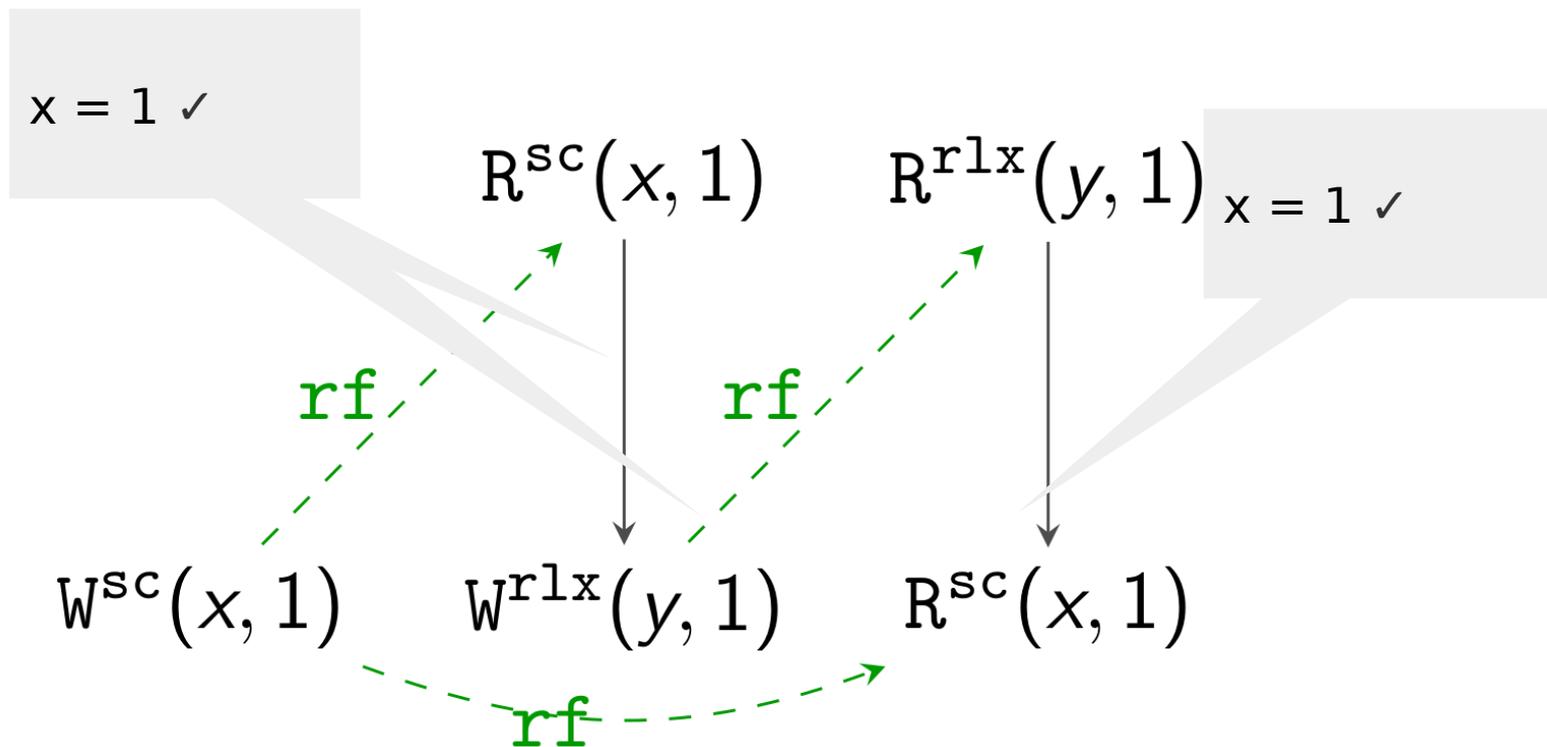
“Release” known writes and “acquire” them next



“Release” known writes and “acquire” them next



“Release” known writes and “acquire” them next



Implemented with release and acquire fences

# Implemented with release and acquire fences

$\text{compile}(\text{Prog}) = [\text{na} \rightarrow \text{rlx}, \text{at} \rightarrow \text{sc}]\text{Prog} + \text{Fences}^{\text{rel}} + \text{Fences}^{\text{acq}}$

# Implemented with release and acquire fences

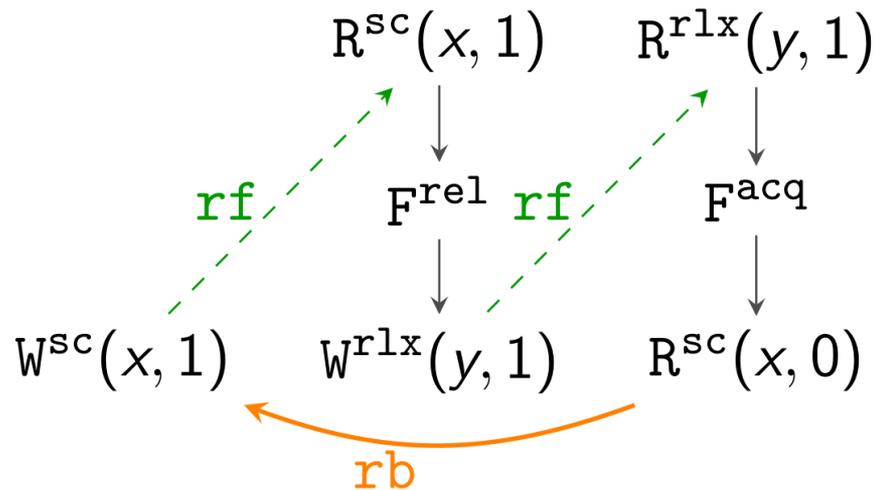
$\text{compile}(\text{Prog}) = [\text{na} \rightarrow \text{rlx}, \text{at} \rightarrow \text{sc}]\text{Prog} + \text{Fences}^{\text{rel}} + \text{Fences}^{\text{acq}}$

$[x] = [y] = 0$		
$[x]^{\text{sc}} := 1$	$a := [x]^{\text{sc}}$	$b := [y]^{\text{rlx}}$
	$\text{fence}^{\text{rel}}$	$\text{fence}^{\text{acq}}$
	$[y]^{\text{rlx}} := 1$	$c := [x]^{\text{sc}}$

# Implemented with release and acquire fences

$\text{compile}(\text{Prog}) = [\text{na} \rightarrow \text{rlx}, \text{at} \rightarrow \text{sc}]\text{Prog} + \text{Fences}^{\text{rel}} + \text{Fences}^{\text{acq}}$

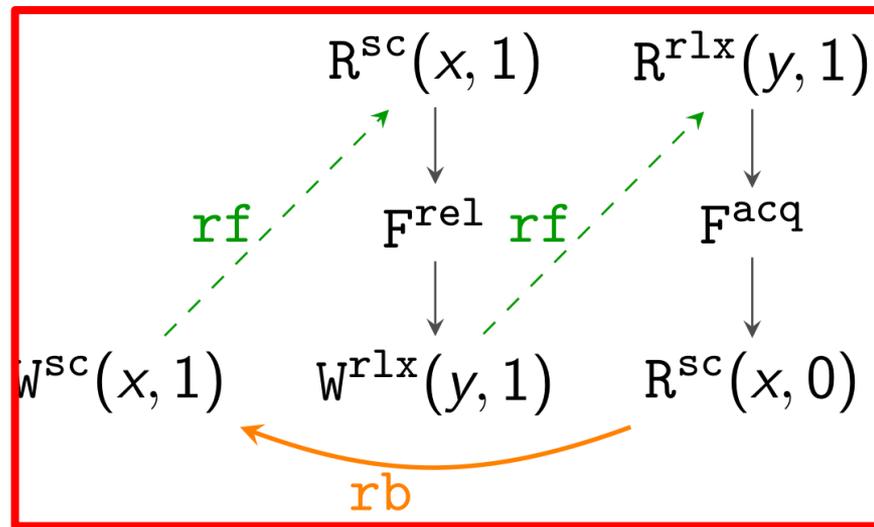
$[x] = [y] = 0$		
$[x]^{\text{sc}} := 1$	$a := [x]^{\text{sc}}$	$b := [y]^{\text{rlx}}$
	$\text{fence}^{\text{rel}}$	$\text{fence}^{\text{acq}}$
	$[y]^{\text{rlx}} := 1$	$c := [x]^{\text{sc}}$



# Implemented with release and acquire fences

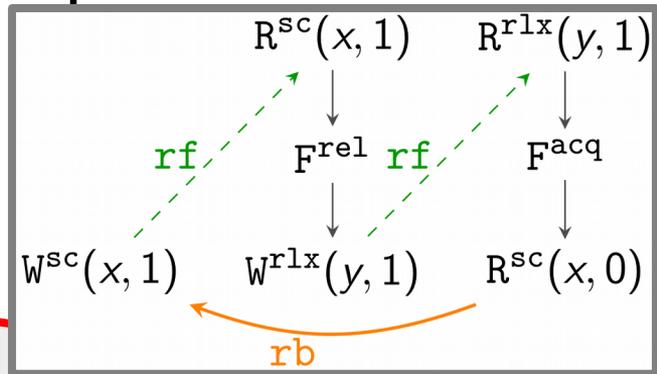
$\text{compile}(\text{Prog}) = [\text{na} \rightarrow \text{rlx}, \text{at} \rightarrow \text{sc}]\text{Prog} + \text{Fences}^{\text{rel}} + \text{Fences}^{\text{acq}}$

$[x] = [y] = 0$		
$[x]^{\text{sc}} := 1$	$a := [x]^{\text{sc}}$	$b := [y]^{\text{rlx}}$
	$\text{fence}^{\text{rel}}$	$\text{fence}^{\text{acq}}$
	$[y]^{\text{rlx}} := 1$	$c := [x]^{\text{sc}}$
<del><math>a = b = 1, c = 0</math></del>		



IMM: can have a cycle made of **po**, **rf** and **rb** if there is **rf** without sc and fences

# An IMM-inconsistent behavior is now prohibited



$\text{Graphs}_{\text{OCaml MM}}(\text{Prog})$

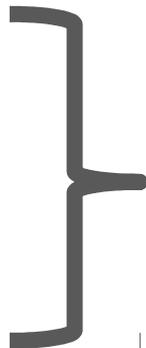
$\text{Graphs}_{\text{IMM}}(\text{compile}(\text{Prog}))$

The resulting scheme prohibits unwanted behaviors

OCaml MM	IMM
$r := [x]^{na}$	$r := [x]^{rlx}$
$[x]^{na} := v$	$fence^{acqrel}; [x]^{rlx} := v$
$r := [x]^{at}$	$fence^{acq}; r := [x]^{sc}$
$[x]^{at} := v$	$fence^{acq}; exchange^{sc}(x, v)$

# Takeaway

- Compilation scheme from OCaml MM to IMM
- Proved to be correct
- Formalized in Coq



Machine-verified  
compilation scheme from  
OCaml MM to Power

<https://github.com/weakmemory/imm>