

# Synthesizing Memory Models from Framework Sketches and Litmus Tests

James Bornholt

Emina Torlak

University of Washington

**Memory consistency models define memory reordering behaviors on multiprocessors**

# Memory consistency models define memory reordering behaviors on multiprocessors

...correctness of  
my compiler...

Compiler  
writers



# Memory consistency models define memory reordering behaviors on multiprocessors

...correctness of  
my compiler...

Compiler  
writers



...rules to verify  
against...

Verification  
tools



# Memory consistency models define memory reordering behaviors on multiprocessors

...correctness of my compiler...

Compiler writers



...rules to verify against...

Verification tools



...possible low-level behaviors...

Kernel/library developers



# Memory consistency models define memory reordering behaviors on multiprocessors

...correctness of my compiler...

Compiler writers



...rules to verify against...

Verification tools



...possible low-level behaviors...

Kernel/library developers



Litmus tests and prose

# Memory consistency models define memory reordering behaviors on multiprocessors


...correctness of my compiler...

Compiler writers



...rules to verify against...

Verification tools




...possible low-level behaviors...

Kernel/library developers

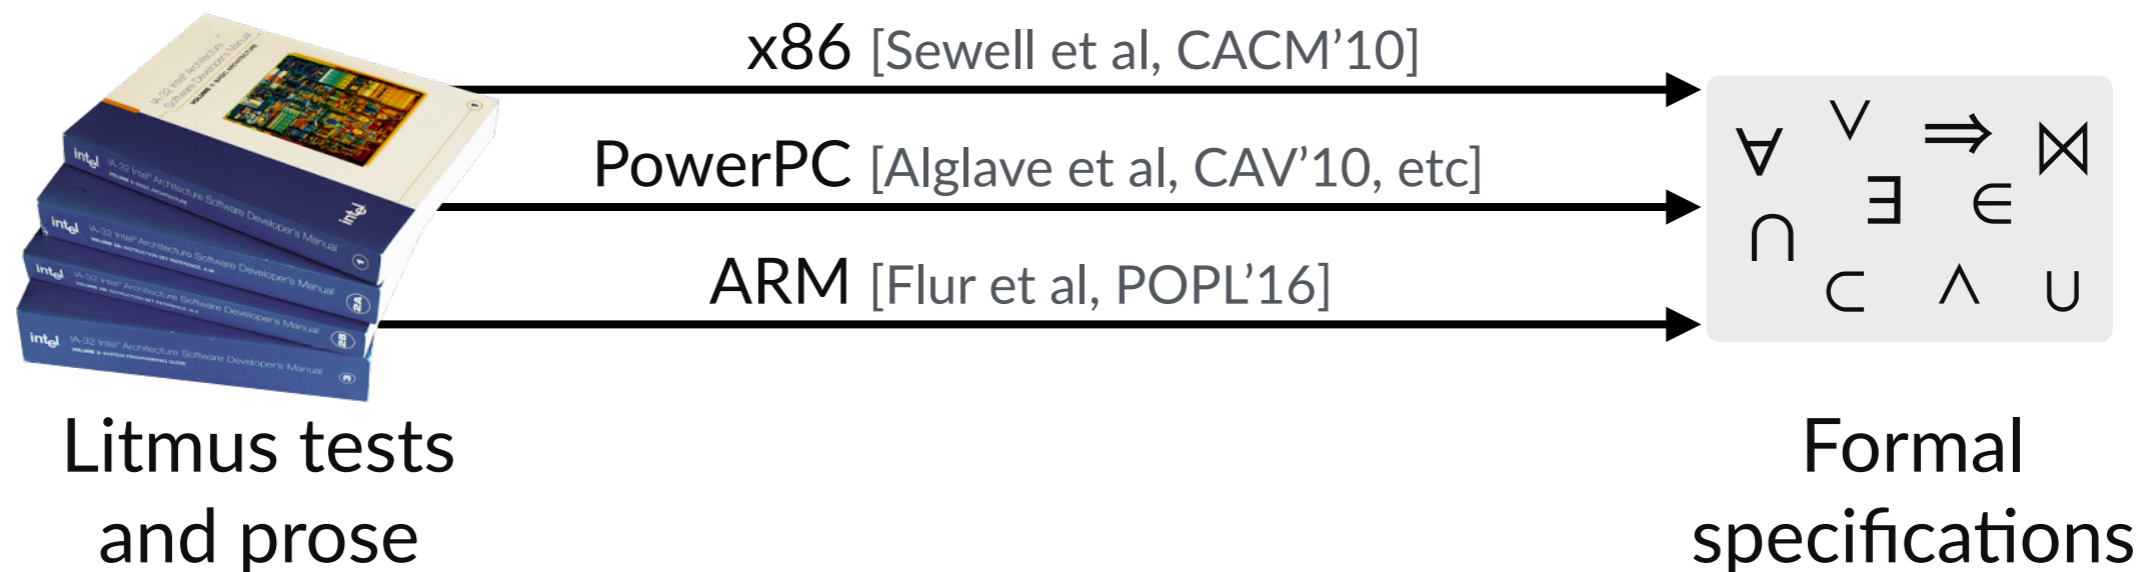
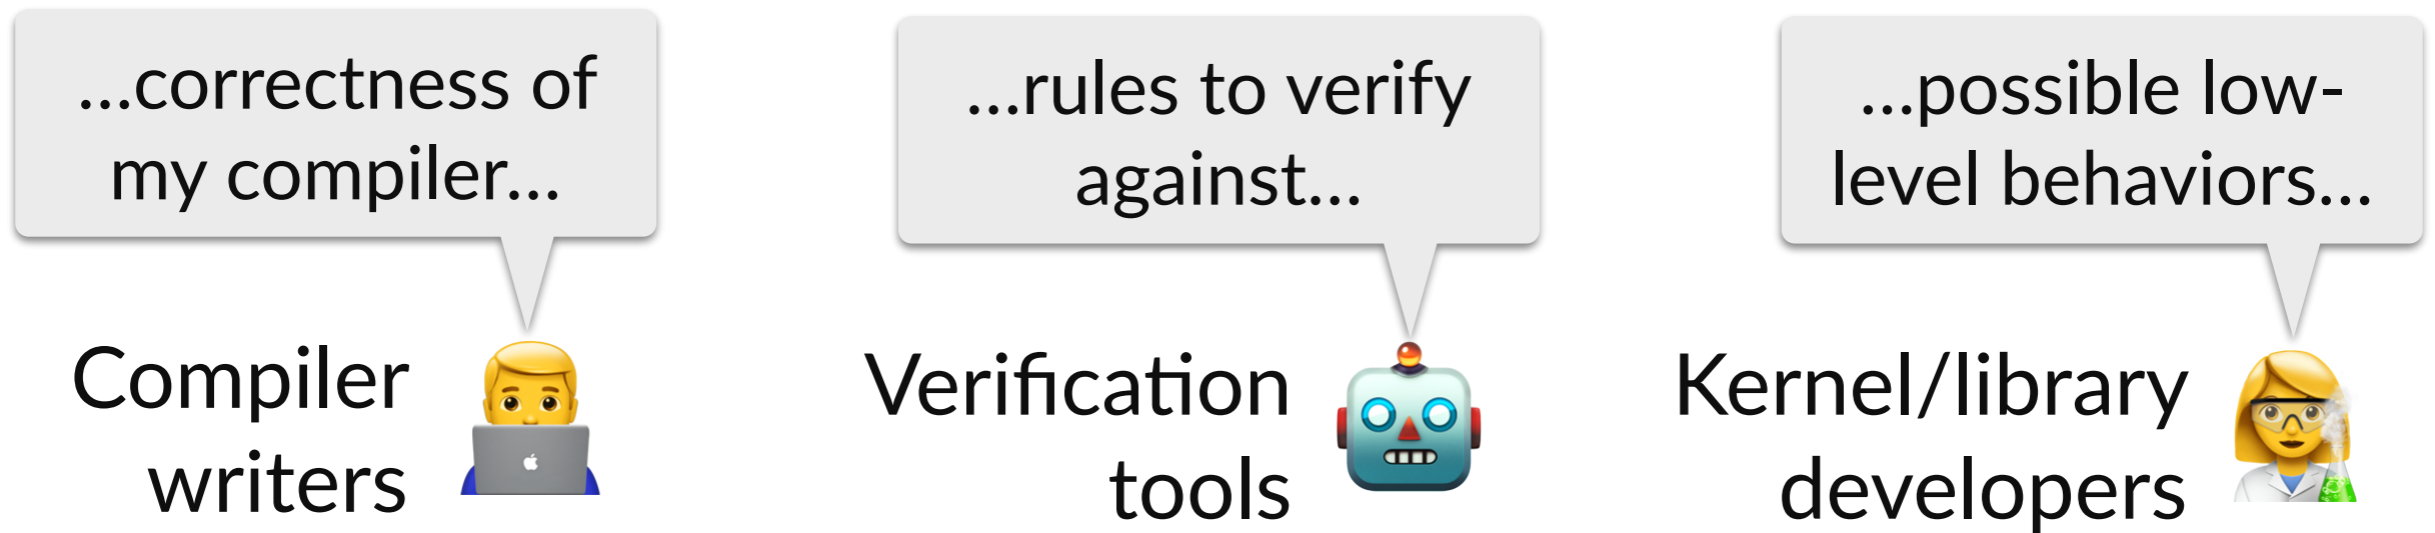


Litmus tests and prose



Formal specifications

# Memory consistency models define memory reordering behaviors on multiprocessors






# MemSynth



Litmus tests




Formal specifications

# MemSynth



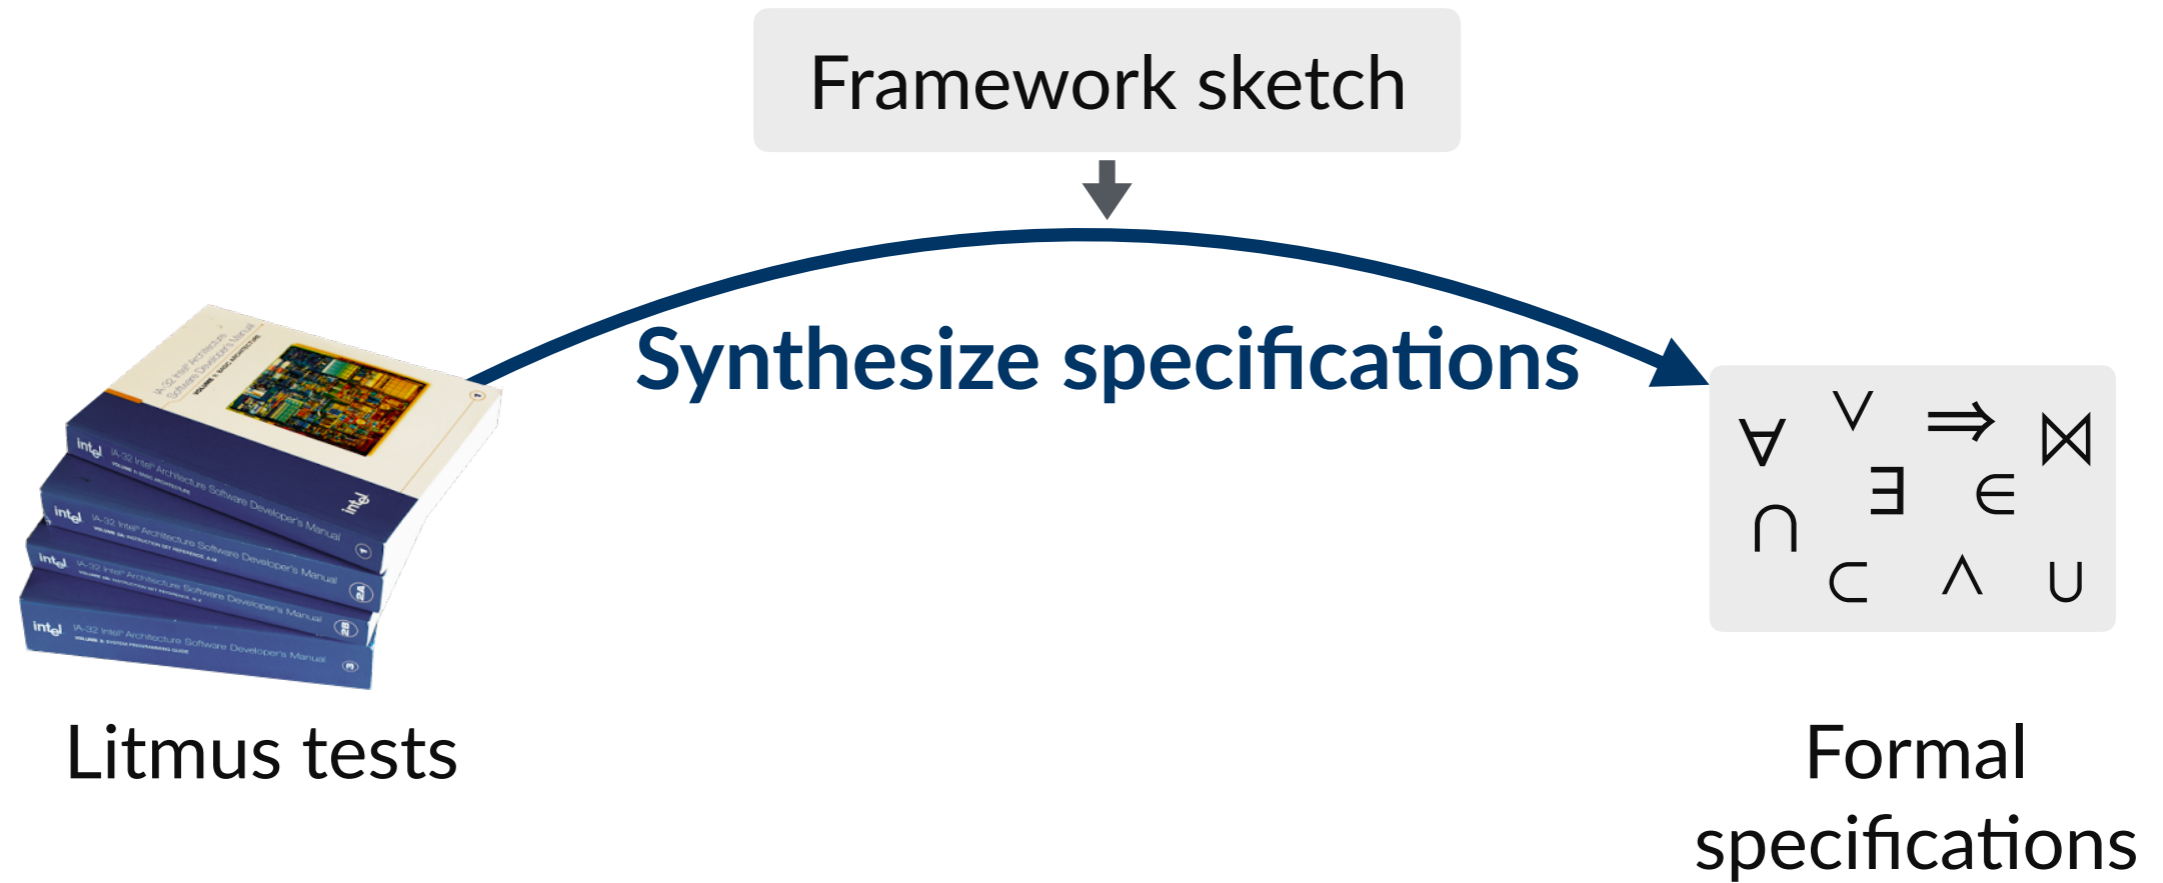
Litmus tests

Synthesize specifications

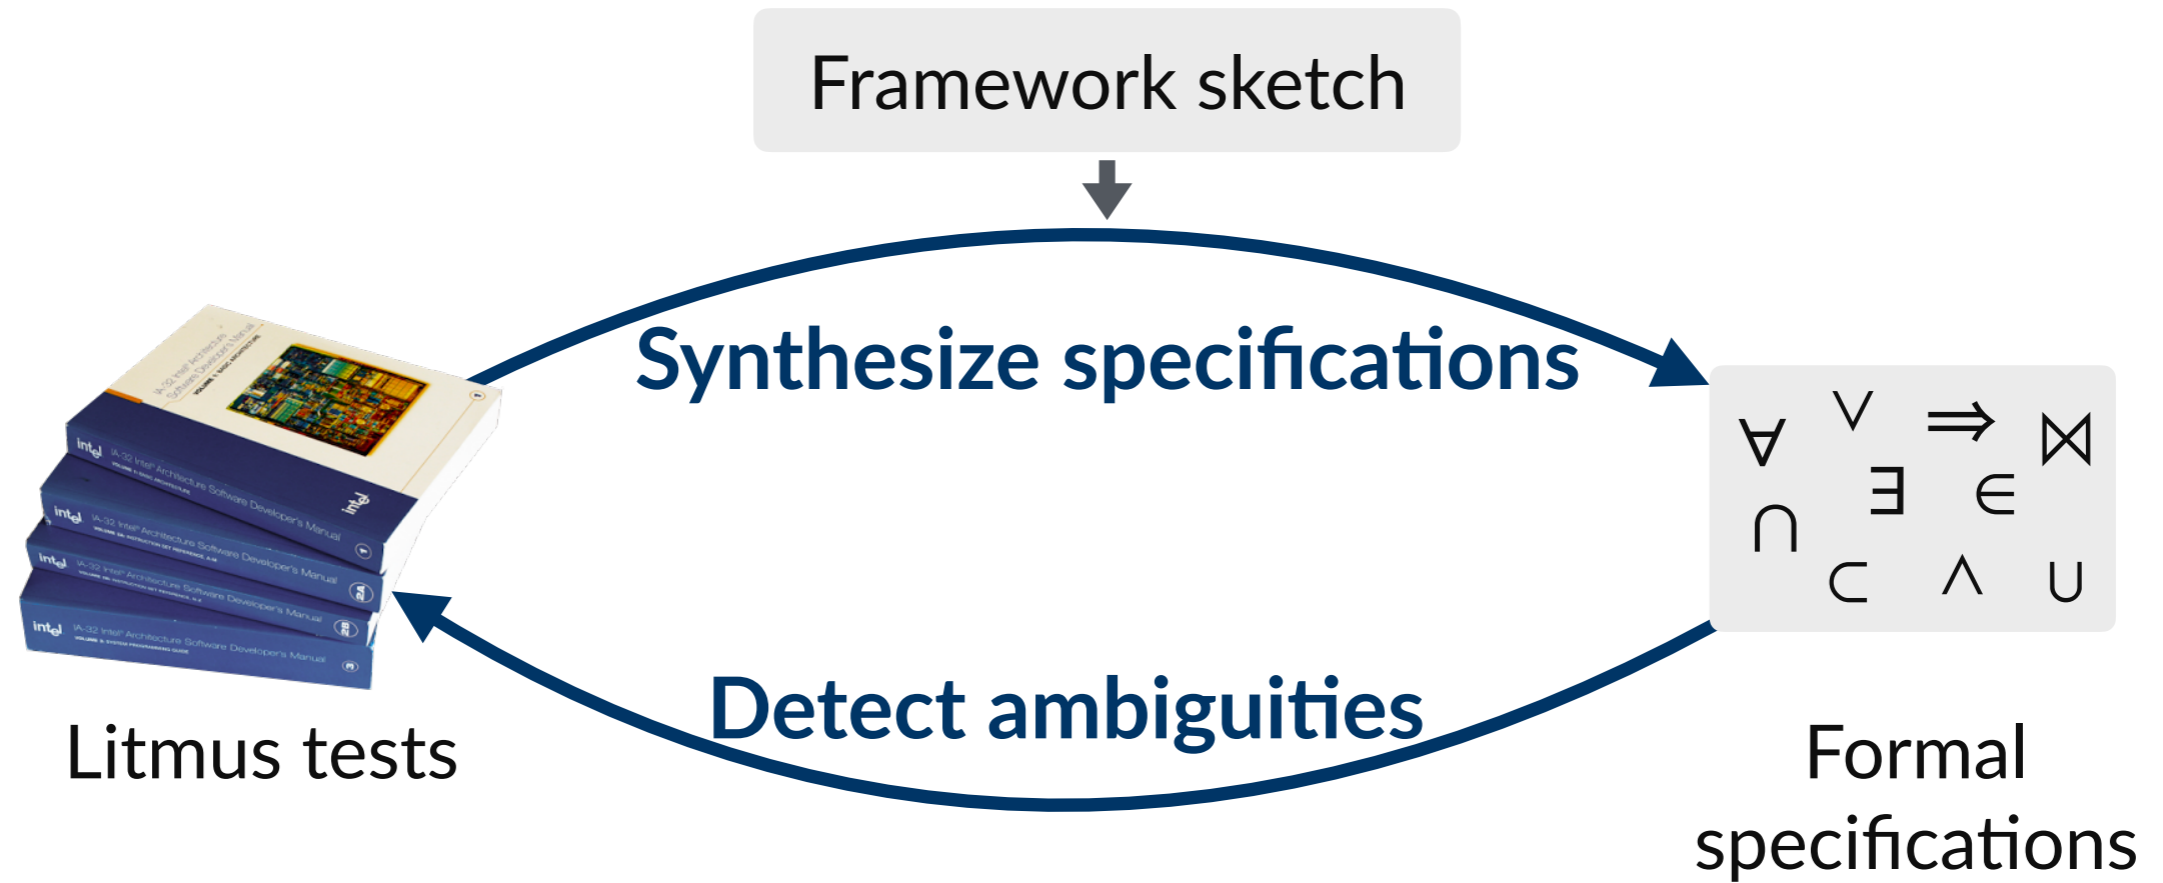


Formal specifications

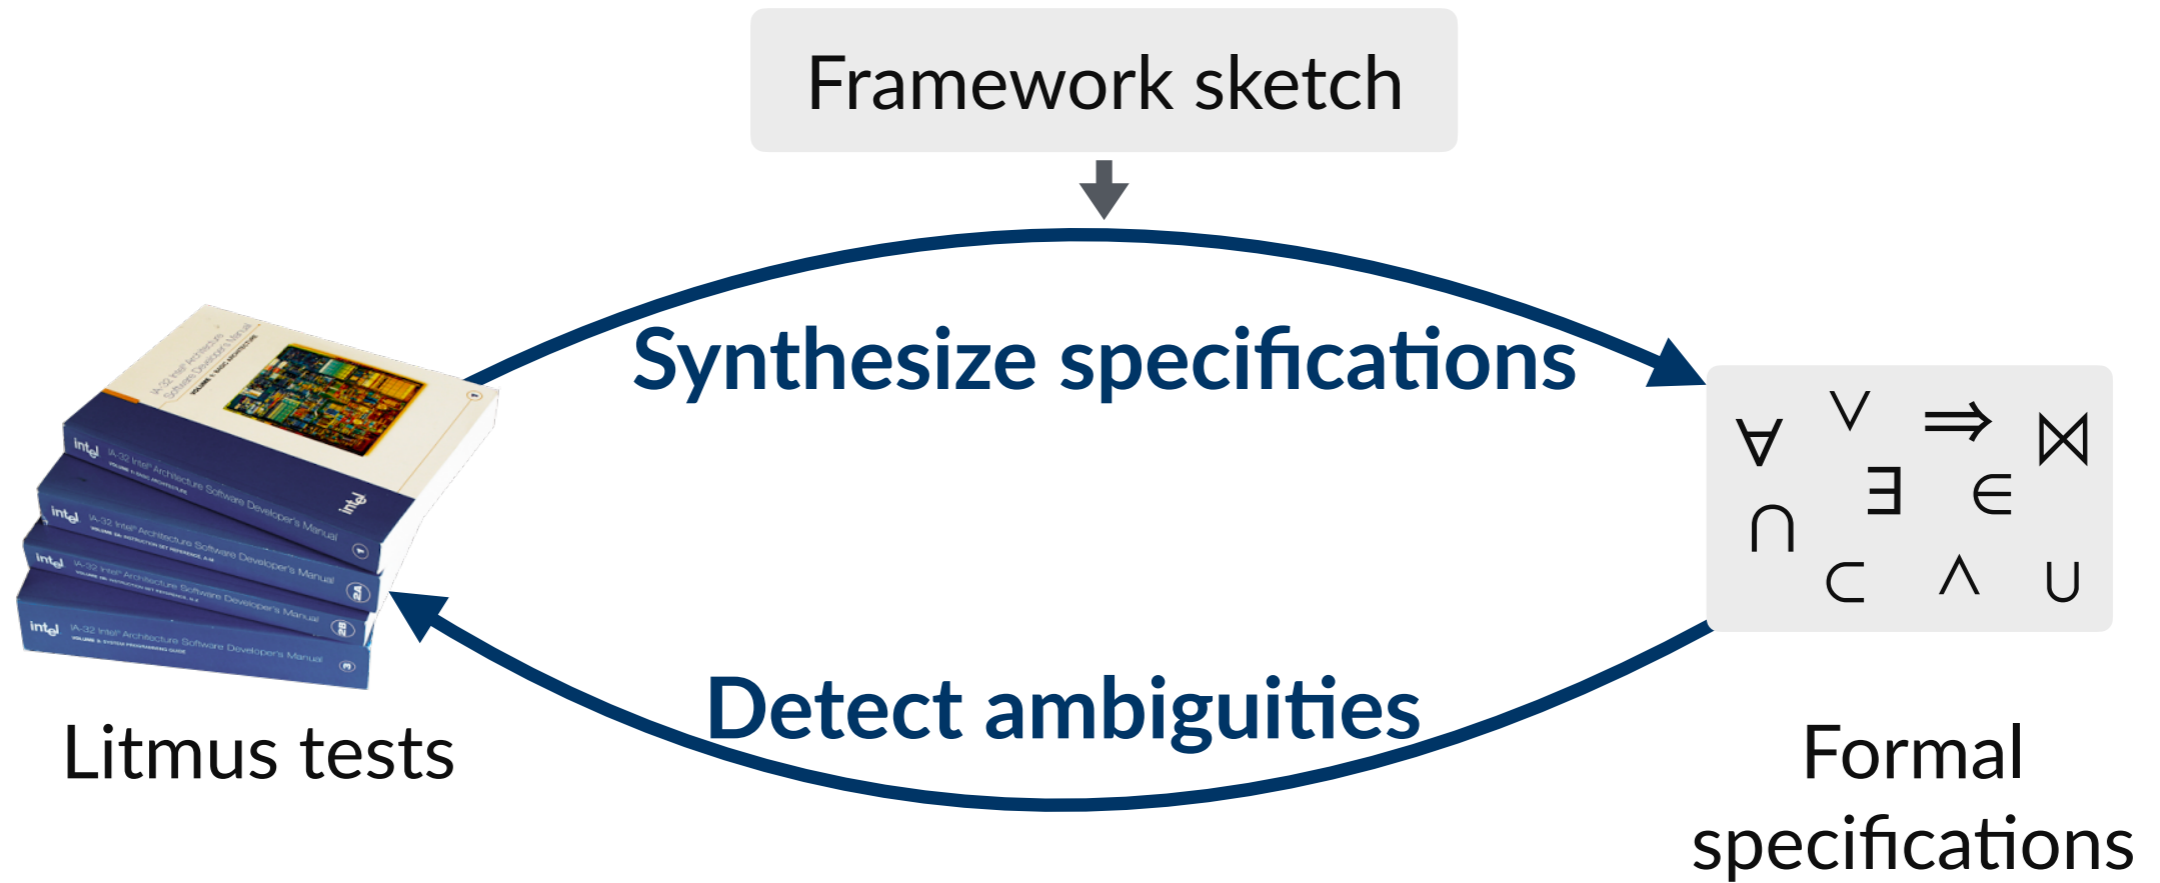
# MemSynth




# MemSynth




# MemSynth



# MemSynth



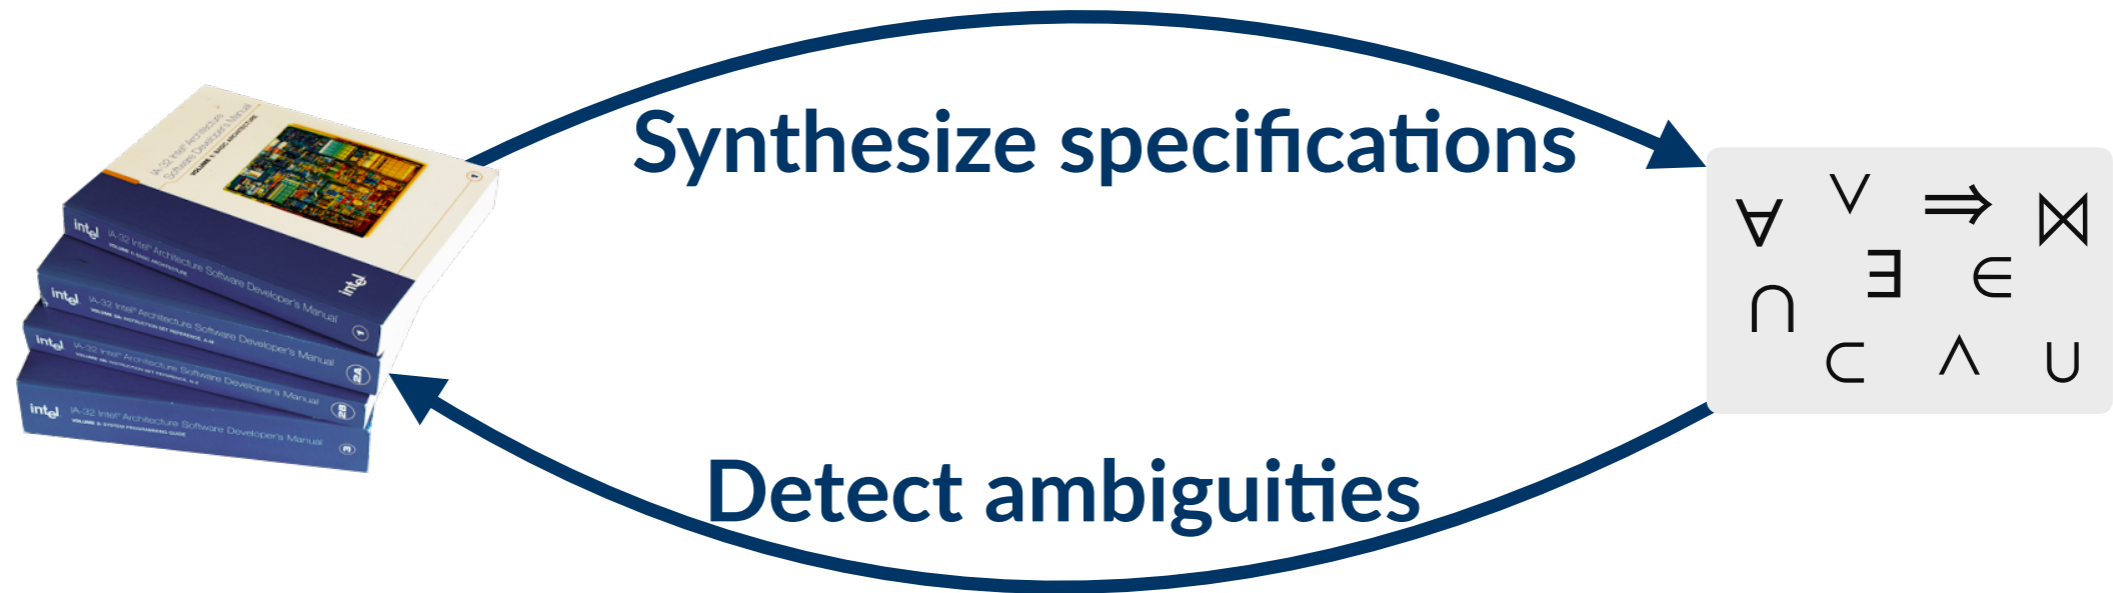
# MemSynth



## Framework sketches

define a class of memory models

# MemSynth



## Framework sketches

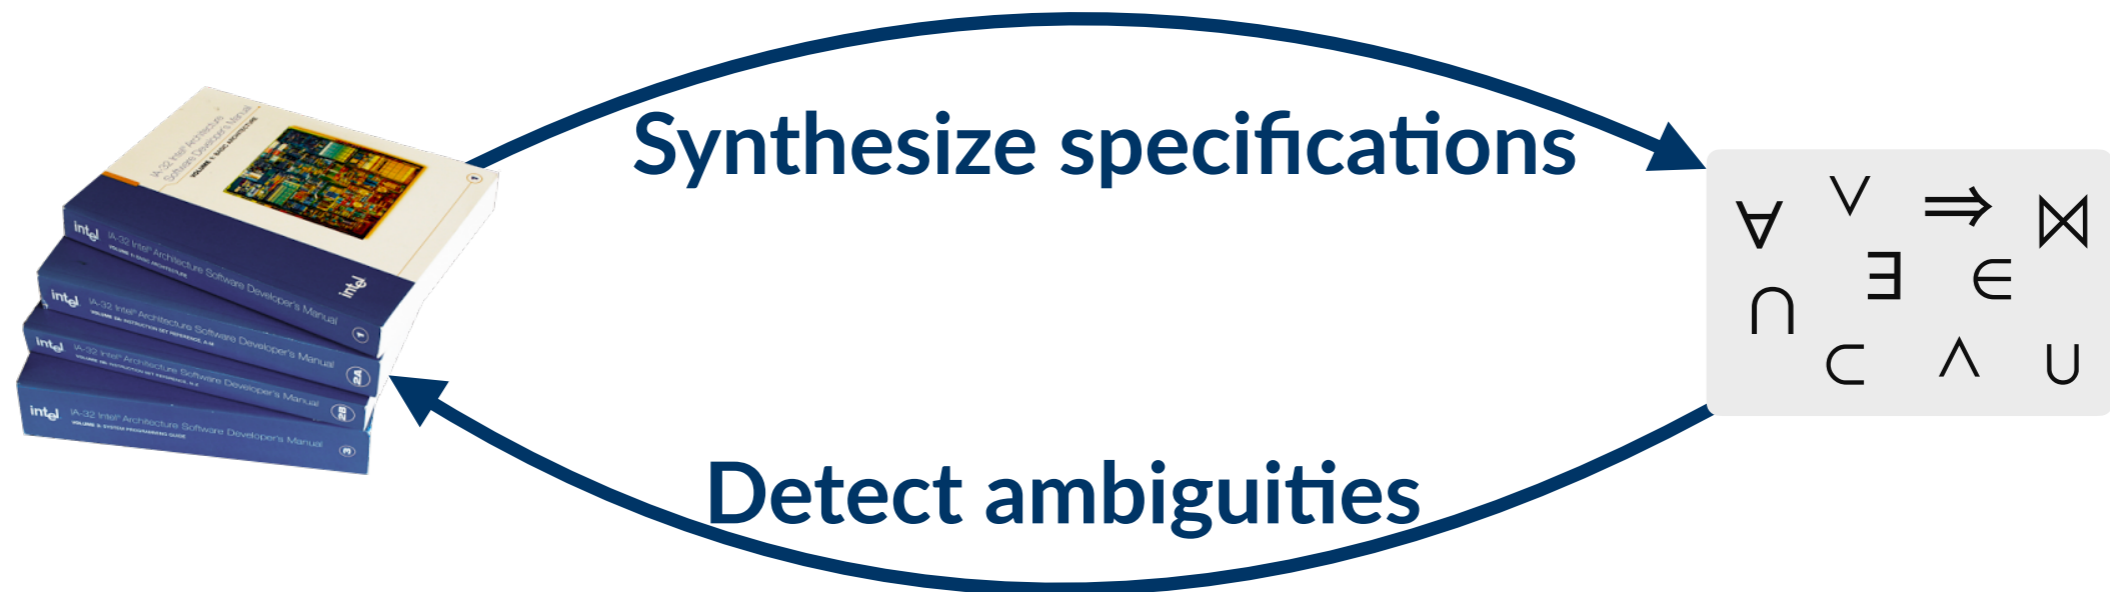
define a class of memory models

## MemSynth engine

verification, equivalence, synthesis, ambiguity



# MemSynth



## Framework sketches

define a class of memory models

## MemSynth engine

verification, equivalence, synthesis, ambiguity

## Results

synthesize real-world memory model specs

# Memory models and framework sketches

# Litmus tests illustrate memory model behavior

**Thread 1**

**Thread 2**

---

①  $X = 1$

③  $Y = 1$

②  $r1 = Y$

④  $r2 = X$

---

Can  $r1 = 0 \wedge r2 = 0$ ?

# Litmus tests illustrate memory model behavior

**Thread 1**

**Thread 2**

---

①  $X = 1$

③  $Y = 1$

②  $r1 = Y$

④  $r2 = X$

---

Can  $r1 = 0 \wedge r2 = 0$ ?

**Sequential consistency: no**

# Litmus tests illustrate memory model behavior

**Thread 1**

**Thread 2**

---

①  $X = 1$

③  $Y = 1$

②  $r1 = Y$

④  $r2 = X$

---

Can  $r1 = 0 \wedge r2 = 0$ ?

**Sequential consistency: no**

**x86: yes!**

# Litmus tests illustrate memory model behavior

Thread 1	Thread 2
① $X = 1$	③ $Y = 1$
② $r1 = Y$	④ $r2 = X$

Can  $r1 = 0 \wedge r2 = 0$ ?

Sequential consistency: no

x86: yes!

A memory model  $M$  is a set of constraints that define the possible executions (outcomes) of a program.

# Litmus tests illustrate memory model behavior

**Thread 1**

**Thread 2**

---

①  $X = 1$

③  $Y = 1$

②  $r1 = Y$

④  $r2 = X$

---

Can  $r1 = 0 \wedge r2 = 0$ ?

**Sequential consistency: no**

**x86: yes!**

A **memory model M** is a set of constraints that define the possible executions (outcomes) of a program.

Memory model M **allows** litmus test T if there exists an execution that satisfies M's constraints.

# Litmus tests illustrate memory model behavior

**Thread 1**

①  $X = 1$

②  $r1 = Y$

**Thread 2**

③  $Y = 1$

④  $r2 = X$

Can  $r1 = 0 \wedge r2 = 0$ ?

**Sequential consistency: no**

**x86: yes!**

Memory model  
M allows test T:  
 $\exists E. M(T,E)$

A memory model M is a set of constraints that define the possible executions (outcomes) of a program.



# Memory models, formally

Common formalizations based on **relational logic**

Memory model  
M allows test T:  
 $\exists E. M(T, E)$

Example for **sequential consistency**:

**no**  $\wedge$  (ws + fr + po + rf + fences) & **iden**

# Memory models, formally

Common formalizations based on **relational logic**

Memory model  
M allows test T:  
 $\exists E. M(T, E)$

Example for **sequential consistency**:

**no**  $\wedge$  (ws + fr + po + rf + fences) & **iden**

Binary relations over  
program instructions

# Memory models, formally

Common formalizations based on **relational logic**

Memory model  
M allows test T:  
 $\exists E. M(T, E)$

Example for **sequential consistency**:

**no** <sup>happens-before order</sup>  $\wedge (ws + fr + po + rf + fences) \& \mathbf{idem}$

Binary relations over  
program instructions

# Memory models, formally

Common formalizations based on **relational logic**

Memory model  
M allows test T:  
 $\exists E. M(T, E)$

Example for **sequential consistency**:

**no**  $\wedge$  (ws + fr + po + rf + fences) **& iden**

happens-before order

is acyclic

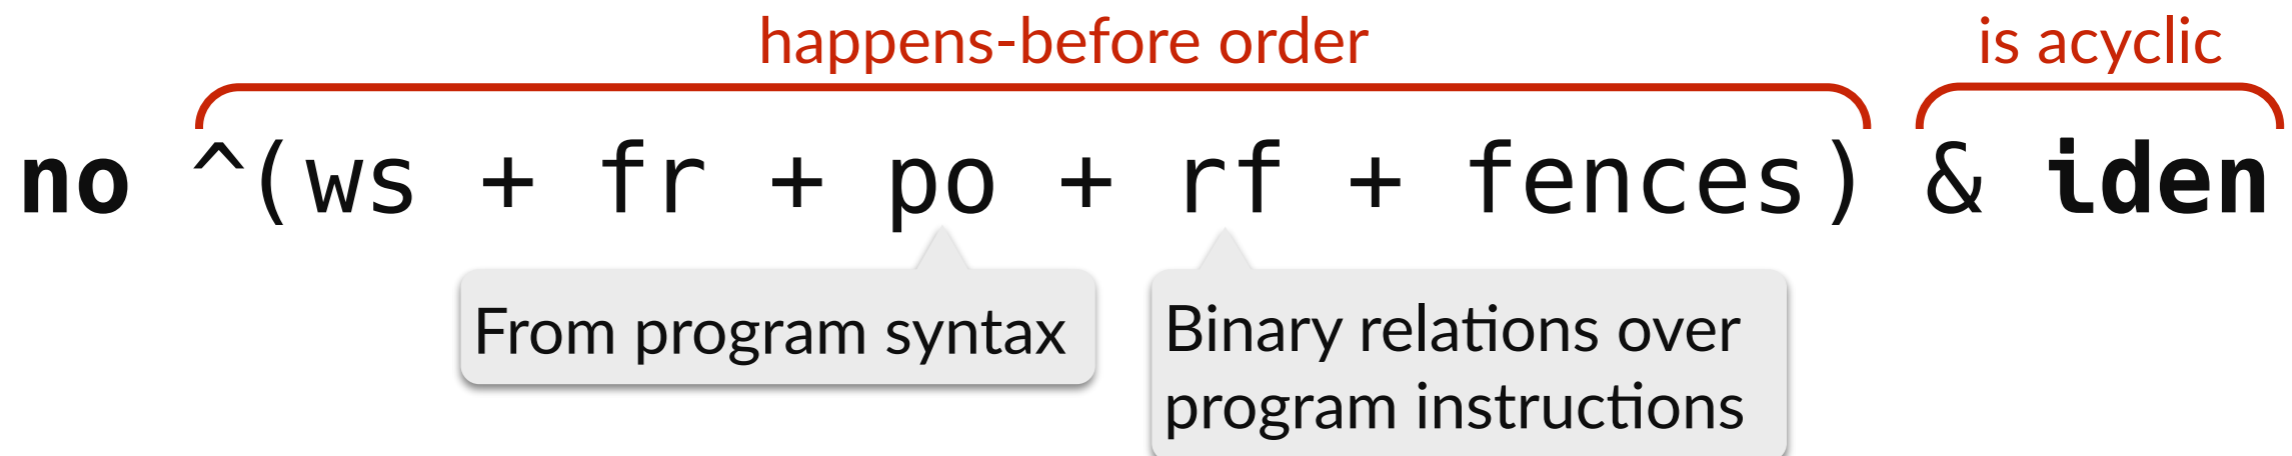
Binary relations over  
program instructions

# Memory models, formally

Common formalizations based on **relational logic**

Memory model  
M allows test T:  
 $\exists E. M(T, E)$

Example for **sequential consistency**:



# Memory models, formally

Common formalizations based on **relational logic**

Memory model  $M$  allows test  $T$ :  
 $\exists E. M(T, E)$

Example for **sequential consistency**:

**no**  $\wedge$  (ws + fr + po + rf + fences) **& iden**

happens-before order

is acyclic

From program syntax

Binary relations over program instructions

**Thread 1**

**Thread 2**

1  $X = 1$

3  $Y = 1$

2  $r1 = Y$

4  $r2 = X$

Can  $r1 = 0 \wedge r2 = 0$ ?

# Memory models, formally

Common formalizations based on **relational logic**

Memory model  $M$  allows test  $T$ :  
 $\exists E. M(T, E)$

Example for **sequential consistency**:

**no**  $\wedge$  (ws + fr + po + rf + fences) **& iden**

happens-before order

is acyclic

From program syntax

Binary relations over program instructions

**Thread 1**

**Thread 2**

1  $X = 1$

3  $Y = 1$

2  $r1 = Y$

4  $r2 = X$

Program order:

$po = \{(1, 2), (3, 4)\}$

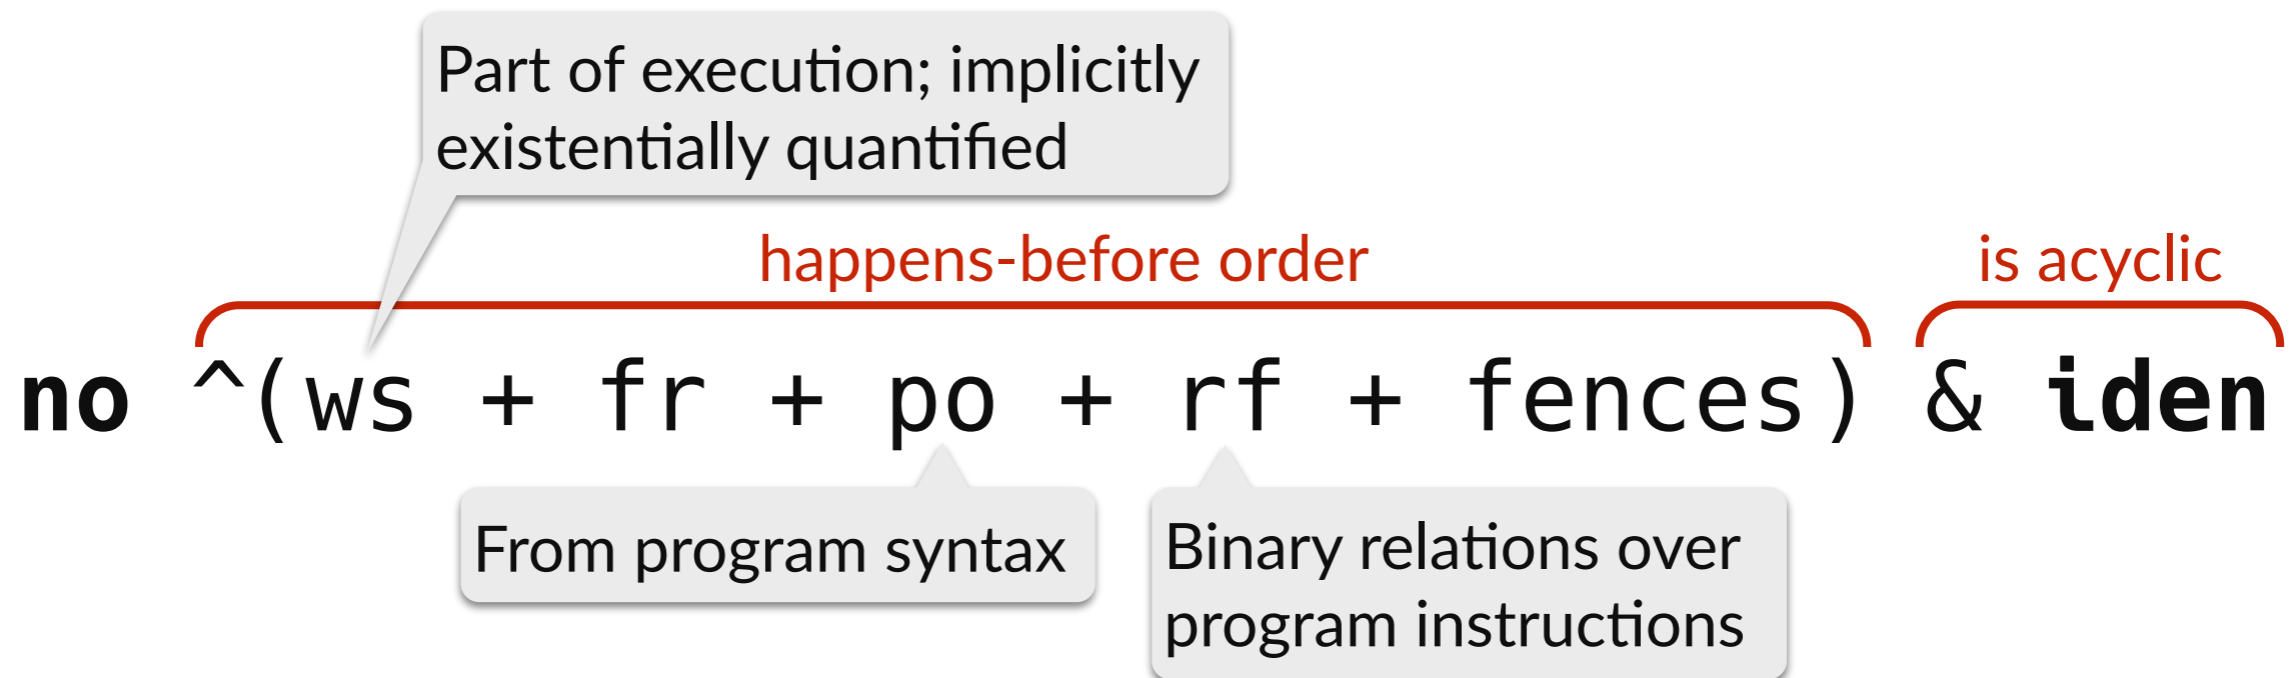
Can  $r1 = 0 \wedge r2 = 0$ ?

# Memory models, formally

Common formalizations based on **relational logic**

Memory model  $M$  allows test  $T$ :  
 $\exists E. M(T, E)$

Example for **sequential consistency**:



**Thread 1**

**Thread 2**

1  $X = 1$

3  $Y = 1$

2  $r1 = Y$

4  $r2 = X$

Can  $r1 = 0 \wedge r2 = 0$ ?

Program order:

$po = \{(1, 2), (3, 4)\}$



# Framework sketches

A **framework sketch** defines the search space for synthesizing a memory model  $M$  by including *holes* in constraints

**no**  $\wedge$  (ws + fr + po + rf + fences) & **iden**

# Framework sketches

A **framework sketch** defines the search space for synthesizing a memory model  $M$  by including *holes* in constraints

Expression holes  
for a synthesizer  
to complete

**no**  $\wedge$  (ws + fr + ?? + ?? + ?? ) & **iden**

# Framework sketches

A **framework sketch** defines the search space for synthesizing a memory model  $M$  by including *holes* in constraints

Expression holes  
for a synthesizer  
to complete

**no**  $\wedge$  (ws + fr + ?? + ?? + ?? ) & **iden**

Framework sketches are the **key design tool** for synthesizing memory model specifications – they define the “interesting” candidate models

# Memory model frameworks

**no**  $\wedge$  (ws + fr + ?? + ?? + ?? ) & **iden**

# Memory model frameworks

**no**  $\wedge$  (*ws* + *fr* + *ppo* + *grf* + *fences*) & **iden**

Preserved program order (same-thread reorderings)

Global reads from (inter-thread order)

Fence cumulativity (for Power, ARM, etc)

# Memory model frameworks

**no**  $\wedge$  (*ws* + *fr* + *ppo* + *grf* + *fences*) & **iden**

Preserved program order (same-thread reorderings)

Global reads from (inter-thread order)

Fence cumulativity (for Power, ARM, etc)

po

rf

∅

**Sequential consistency**

# Memory model frameworks

**no**  $\wedge$  (*ws* + *fr* + *ppo* + *grf* + *fences*) & **iden**

Preserved program order (same-thread reorderings)

Global reads from (inter-thread order)

Fence cumulativity (for Power, ARM, etc)

**Sequential consistency**

*po*

*rf*

$\emptyset$

**Total store order (x86)**

*po* - (*Wr*→*Rd*) *rf* & SameThd

$\emptyset$

# Memory model frameworks are common

Global time  
relational model  
[Alglave et al, CAV'10]

Axiomatic “must-  
not-reorder”  
functions  
[Mador-Haim et al,  
DAC'11]

Exexecutable  
distributed  
consistency models  
[Yang et al, IPDPS'04]

...



# Ocelot: relational logic with holes



A **relational logic DSL** with synthesis support

Built on the Rosette solver-aided language [Torlak & Bodik, PLDI'14]

Expression holes  
for a synthesizer  
to complete

```
no ^ ( ws + fr + ?? + ?? + ?? ) & iden
```

 Available as a Racket package: `raco pkg install ocelot`

# Ocelot: relational logic with holes



A **relational logic DSL** with synthesis support

Built on the Rosette solver-aided language [Torlak & Bodik, PLDI'14]

Expression holes  
for a synthesizer  
to complete

**no**  $\wedge$  ( **ws** + **fr** + ?? + ?? + ?? ) & **iden**

Completions are expressions in  
relational logic with chosen  
operators, terminals, and depth.

 Available as a Racket package: `raco pkg install ocelot`

# Ocelot: relational logic with holes



A **relational logic DSL** with synthesis support

Built on the Rosette solver-aided language [Torlak & Bodik, PLDI'14]

Expression holes  
for a synthesizer  
to complete

**no**  $\wedge$  ( **ws** + **fr** + ?? + ?? + ?? ) & **iden**

Completions are expressions in  
relational logic with chosen  
operators, terminals, and depth.

operators = {+, &}

terminals = {po, ws}

depth = 1

 Available as a Racket package: `raco pkg install ocelot`

# Ocelot: relational logic with holes



A **relational logic DSL** with synthesis support

Built on the Rosette solver-aided language [Torlak & Bodik, PLDI'14]

Expression holes  
for a synthesizer  
to complete

**no**  $\wedge$  (ws + fr + ?? + ?? + ?? ) & **iden**

Completions are expressions in  
relational logic with chosen  
operators, terminals, and depth.

operators = {+, &}  
terminals = {po, ws}  
depth = 1

po  
ws  
po + ws  
po & ws

 Available as a Racket package: `raco pkg install ocelot`

# Queries

- ▶ Verification
- ▶ Equivalence
- ▶ Synthesis
- ▶ Ambiguity

# Verification and equivalence

Common queries for automated memory model reasoning tools

Memory model  
M allows test T:  
 $\exists E. M(T, E)$

Herd [Alglave et al, CAV'10]; MemAlloy [Wickerson et al, POPL'17]; etc.

# Verification and equivalence

Common queries for automated memory model reasoning tools

Memory model  
M allows test T:  
 $\exists E. M(T,E)$

Herd [Alglave et al, CAV'10]; MemAlloy [Wickerson et al, POPL'17]; etc.

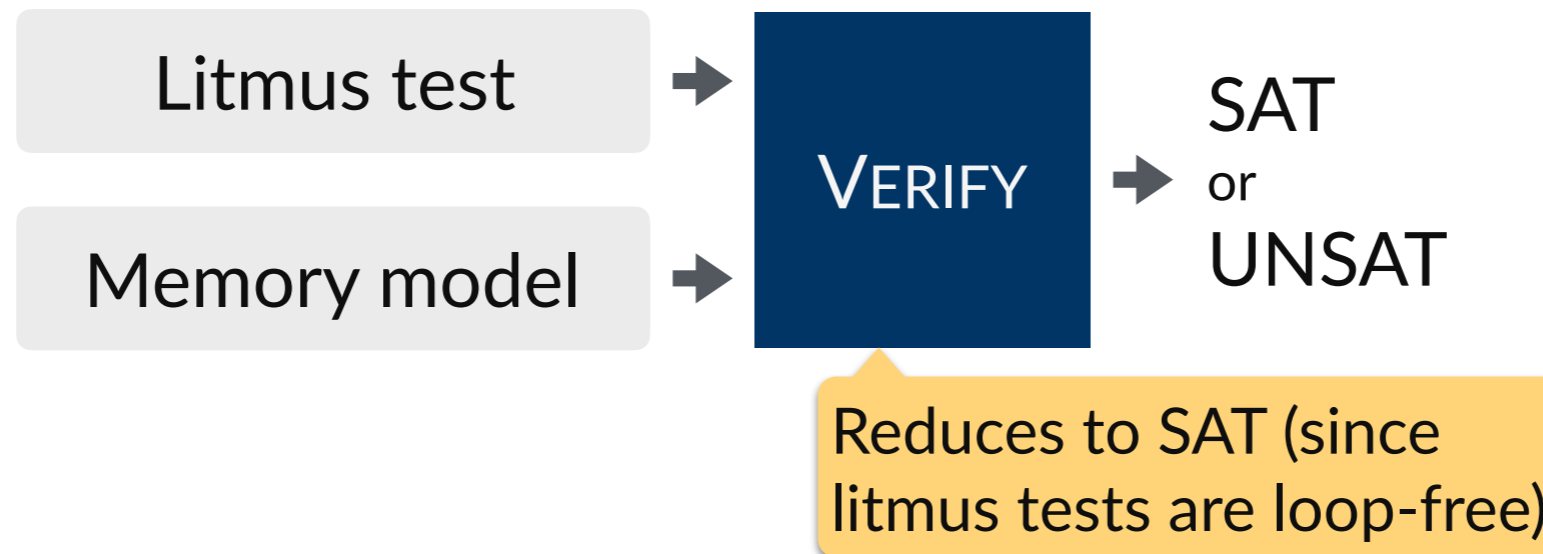


# Verification and equivalence

Common queries for automated memory model reasoning tools

Memory model  $M$  allows test  $T$ :  
 $\exists E. M(T, E)$

Herd [Alglave et al, CAV'10]; MemAlloy [Wickerson et al, POPL'17]; etc.



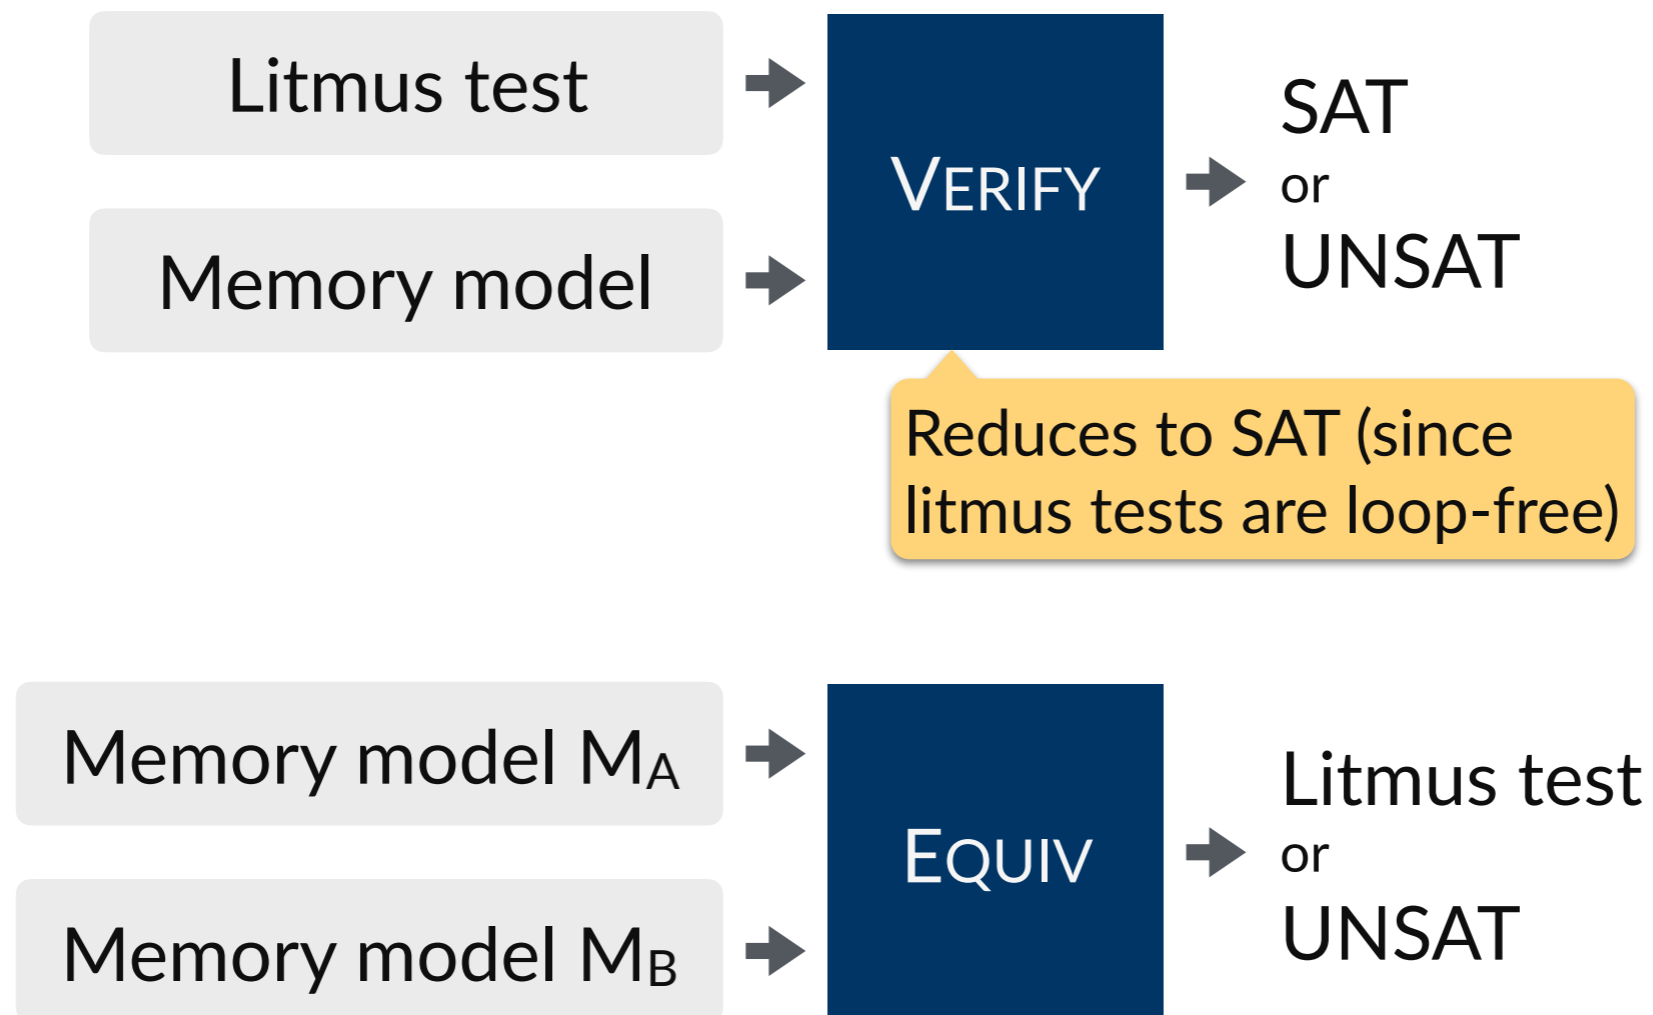


# Verification and equivalence

Common queries for automated memory model reasoning tools

Memory model  $M$  allows test  $T$ :  
 $\exists E. M(T, E)$

Herd [Alglave et al, CAV'10]; MemAlloy [Wickerson et al, POPL'17]; etc.

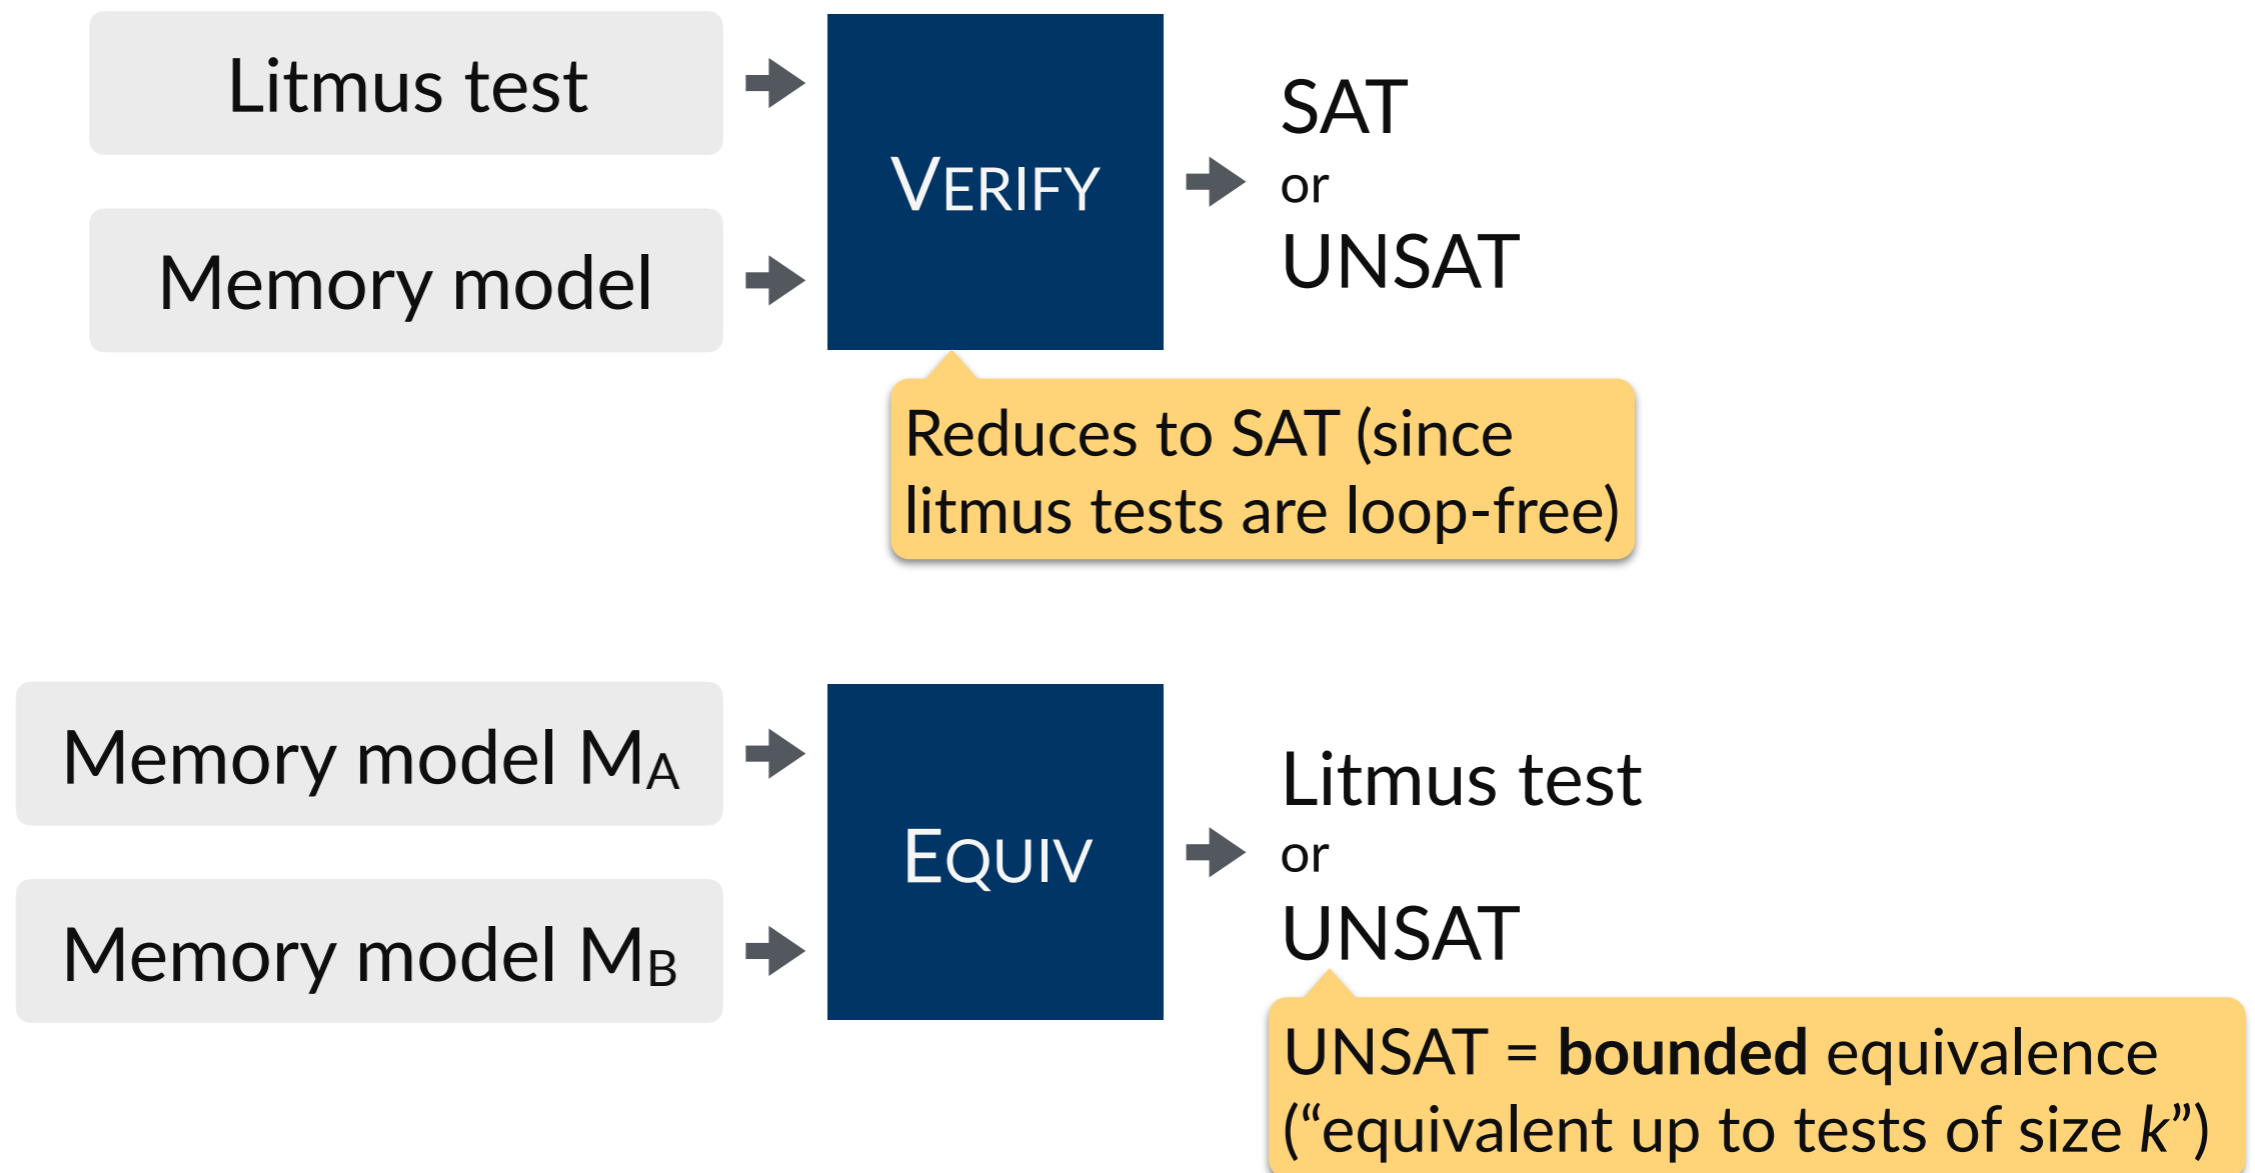


# Verification and equivalence

Common queries for automated memory model reasoning tools

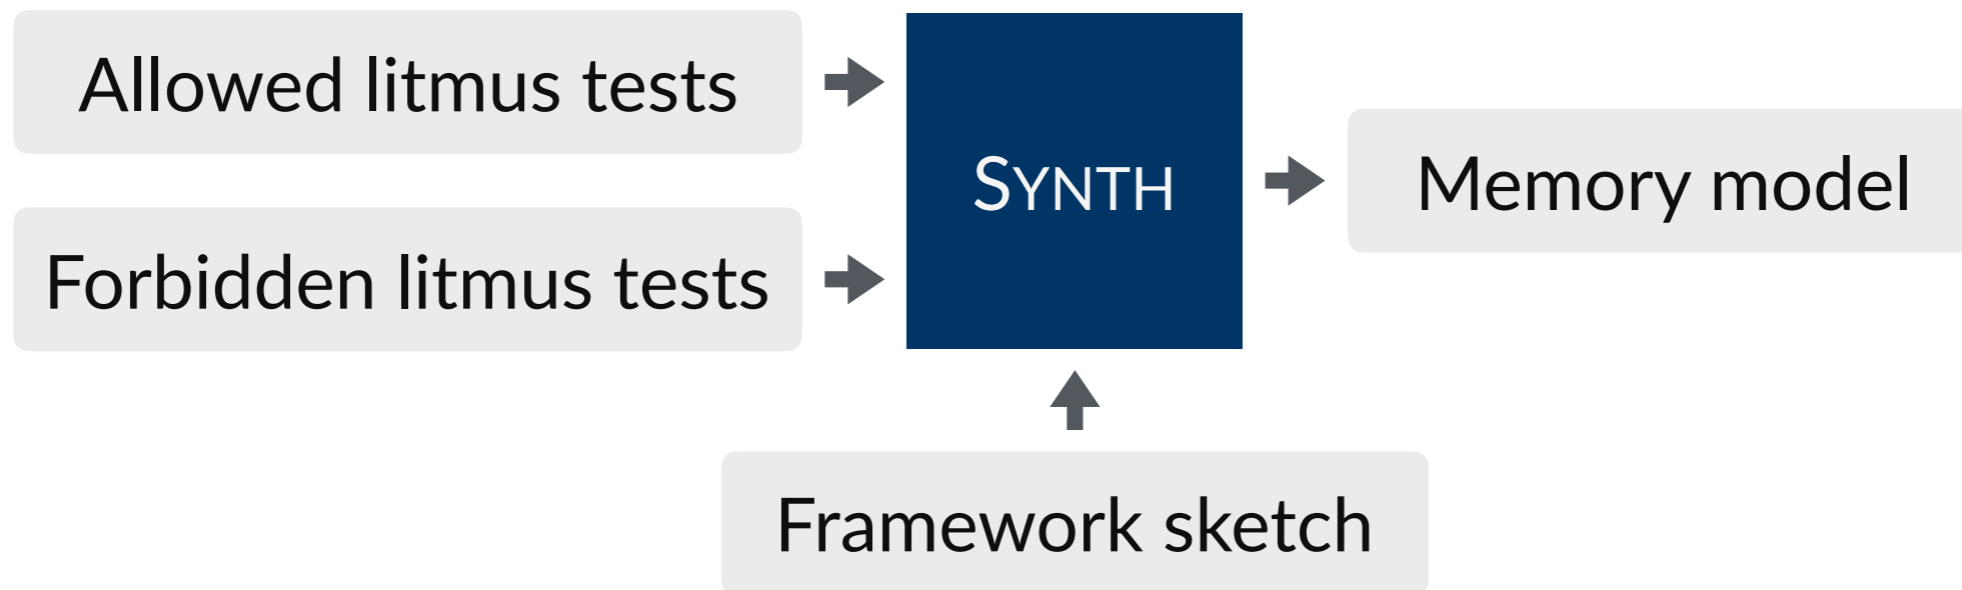
Memory model  $M$  allows test  $T$ :  
 $\exists E. M(T, E)$

Herd [Alglave et al, CAV'10]; MemAlloy [Wickerson et al, POPL'17]; etc.



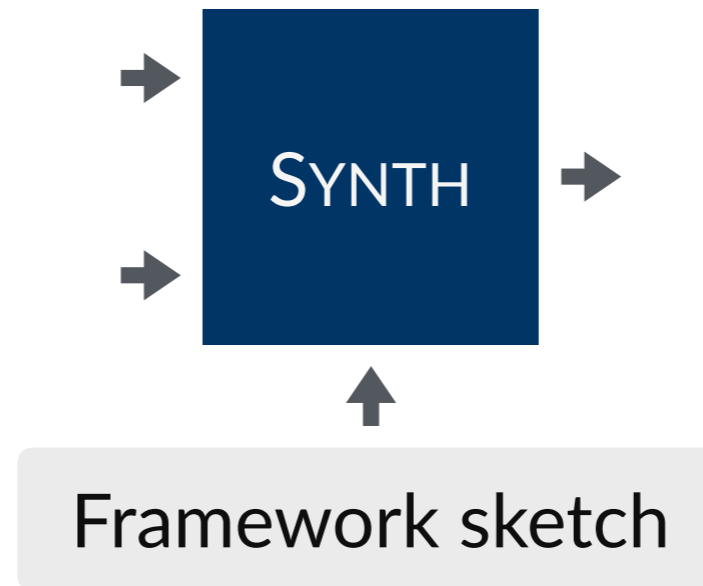
# Synthesis

Find a memory model consistent with a set of litmus tests



# Synthesis

Find a memory model consistent with a set of litmus tests




# Synthesis

Find a memory model consistent with a set of litmus tests

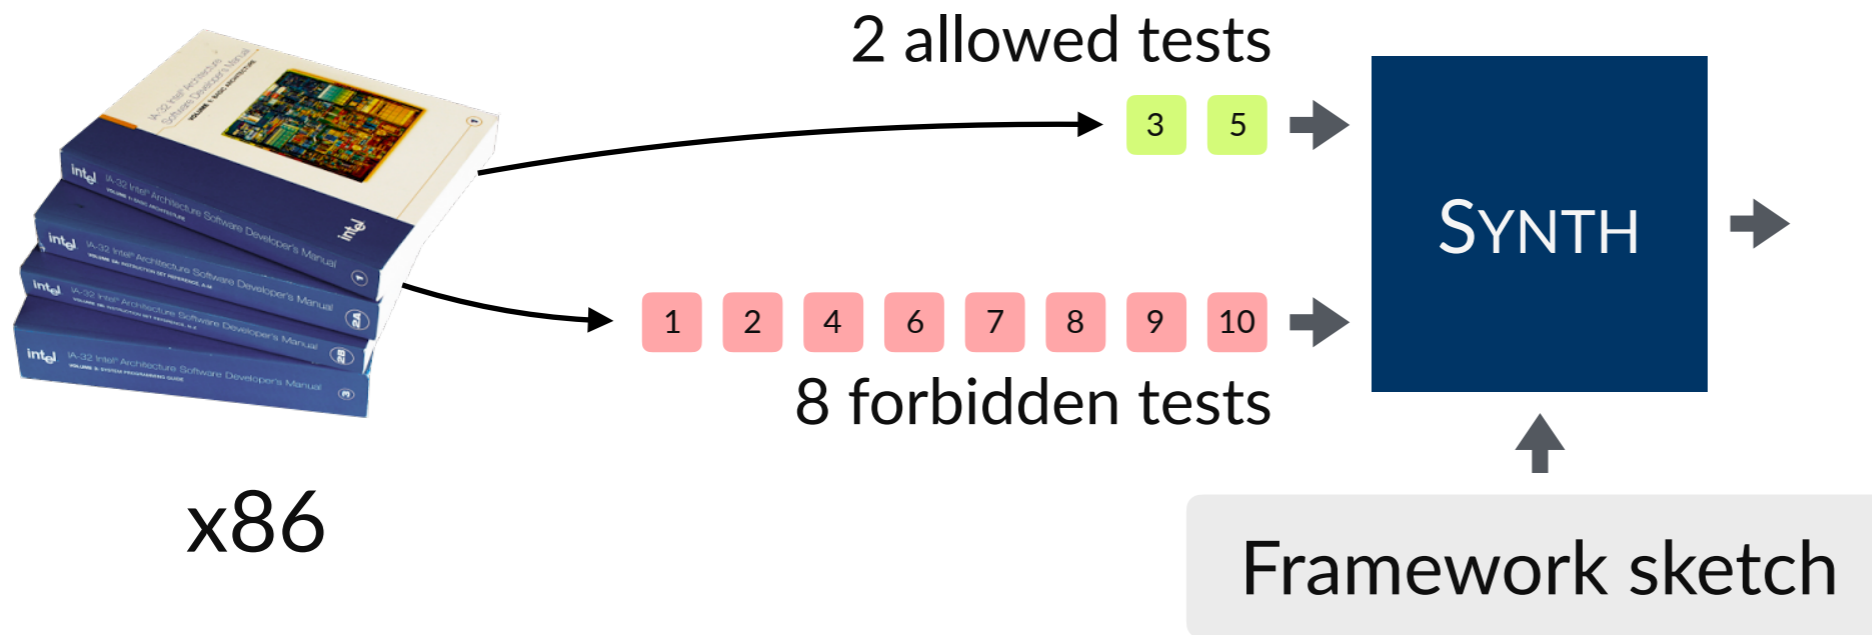


x86



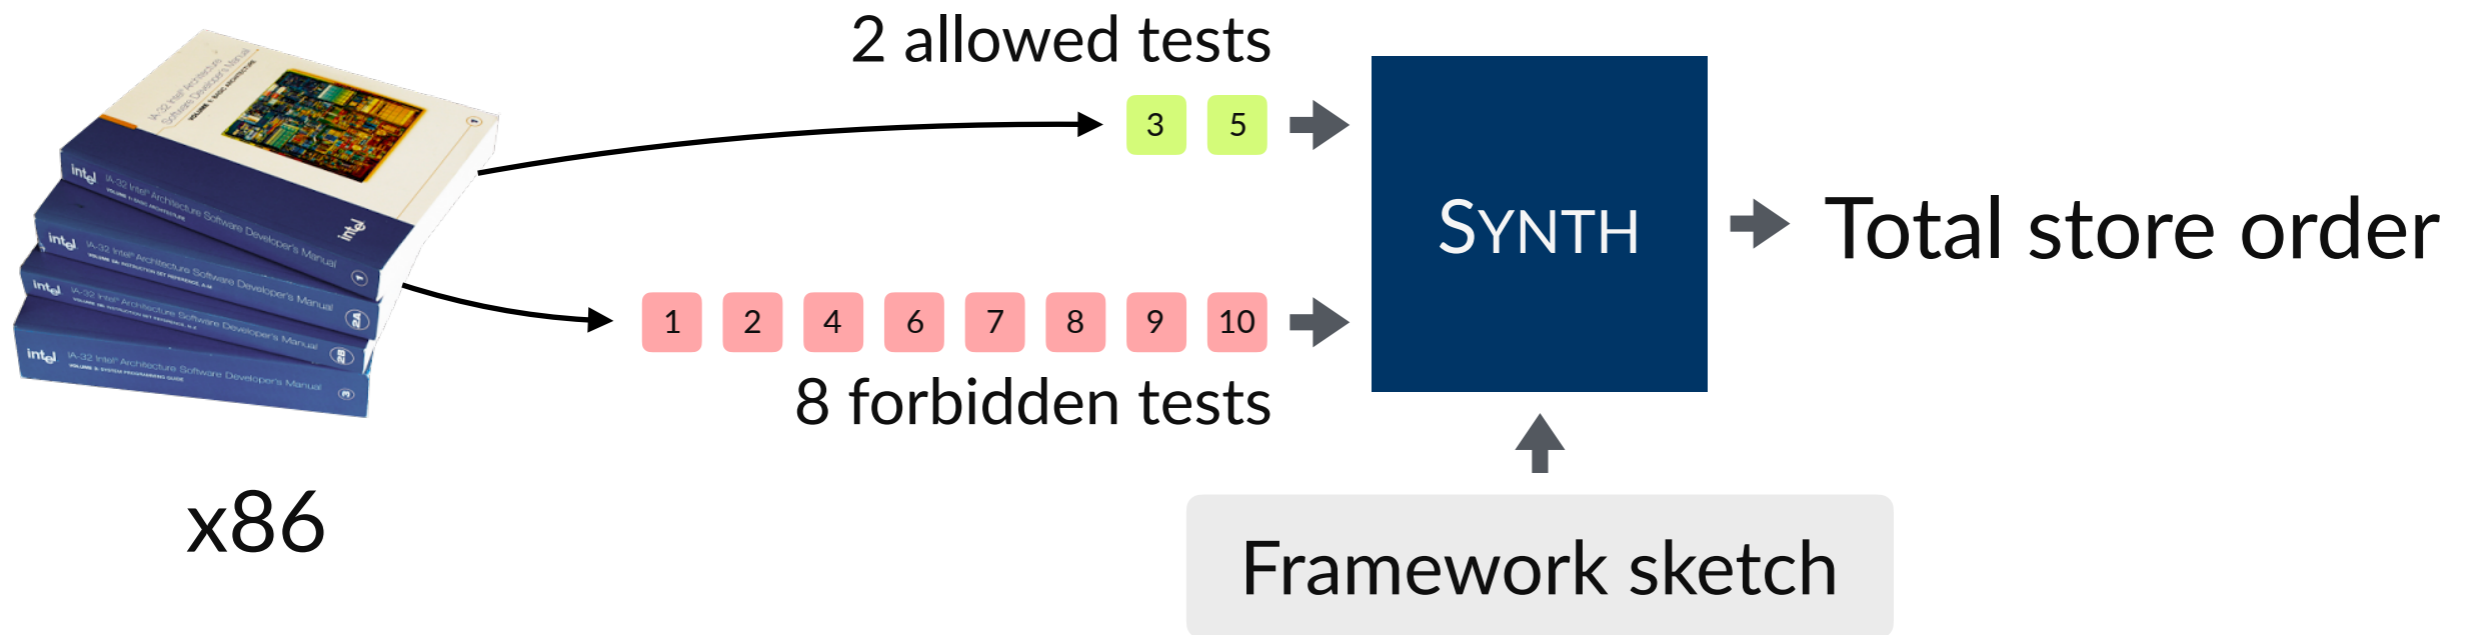
# Synthesis

Find a memory model consistent with a set of litmus tests



# Synthesis

Find a memory model consistent with a set of litmus tests



# Synthesis

Find a memory model consistent with a set of litmus tests

Allowed litmus tests

$T^+$  →

Forbidden litmus tests

$T^-$  →



Memory model

↑  
 $M$

Framework sketch

Memory model  
 $M$  allows test  $T$ :  
 $\exists E. M(T,E)$



# Synthesis

Find a memory model consistent with a set of litmus tests

Allowed litmus tests

$$T^+ \rightarrow \bigwedge_{T \in T^+} \exists E. M(T, E)$$

Forbidden litmus tests

$$T^- \rightarrow$$

Memory model

$\uparrow$   
 $M$

Framework sketch

Memory model  
M allows test T:  
 $\exists E. M(T, E)$

# Synthesis

Find a memory model consistent with a set of litmus tests

Allowed litmus tests

$$T^+ \rightarrow \bigwedge_{T \in T^+} \exists E. M(T, E)$$

Forbidden litmus tests

$$T^- \rightarrow \bigwedge_{T \in T^-} \forall E. \neg M(T, E)$$

$M$

Framework sketch



Memory model

Memory model  
M allows test T:  
 $\exists E. M(T, E)$

# Synthesis

Find a memory model consistent with a set of litmus tests

Allowed litmus tests

$$T^+ \rightarrow \bigwedge_{T \in T^+} \exists E. M(T, E)$$

Forbidden litmus tests

$$T^- \rightarrow \bigwedge_{T \in T^-} \forall E. \neg M(T, E)$$

Solved incrementally, like counterexample-guided inductive synthesis (CEGIS)

$M$

Framework sketch

Memory model

Memory model  $M$  allows test  $T$ :  
 $\exists E. M(T, E)$

# Ambiguity

Find a distinguishing litmus test that exposes an ambiguity in a model

**Key idea:** after synthesis, is there a *different* memory model that explains the tests?

AMBIG

# Ambiguity

Find a distinguishing litmus test that exposes an ambiguity in a model

**Key idea:** after synthesis, is there a *different* memory model that explains the tests?

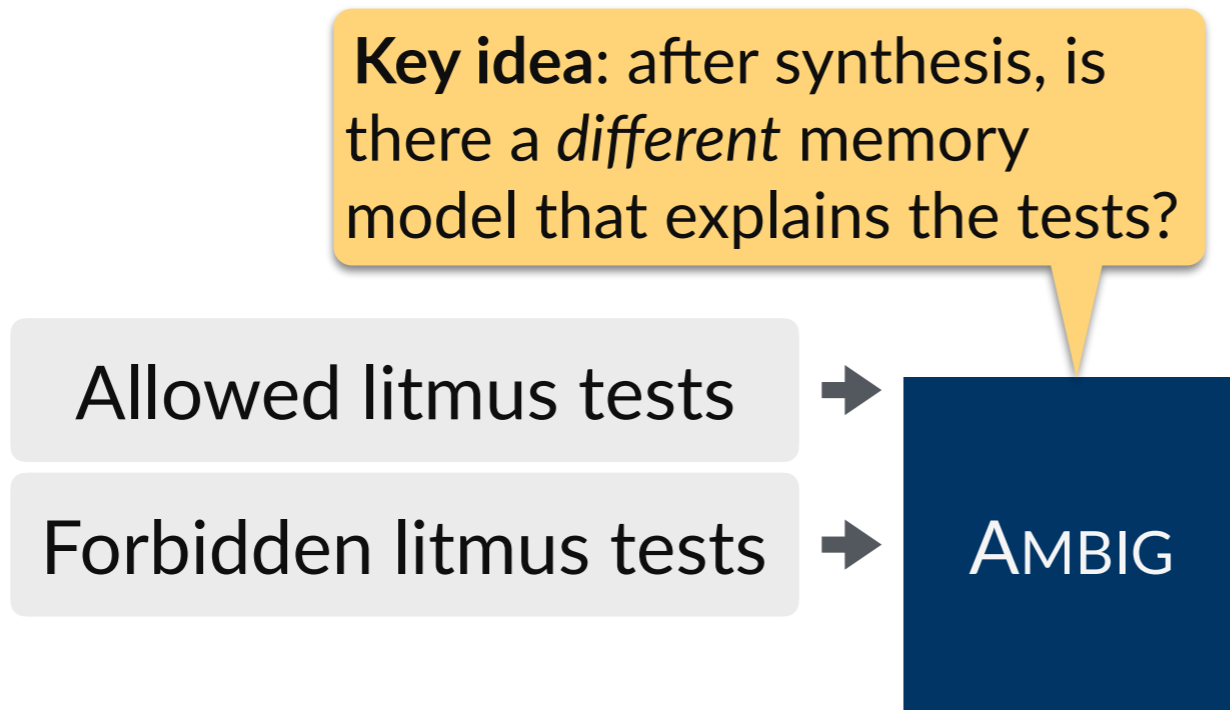
Allowed litmus tests



Forbidden litmus tests



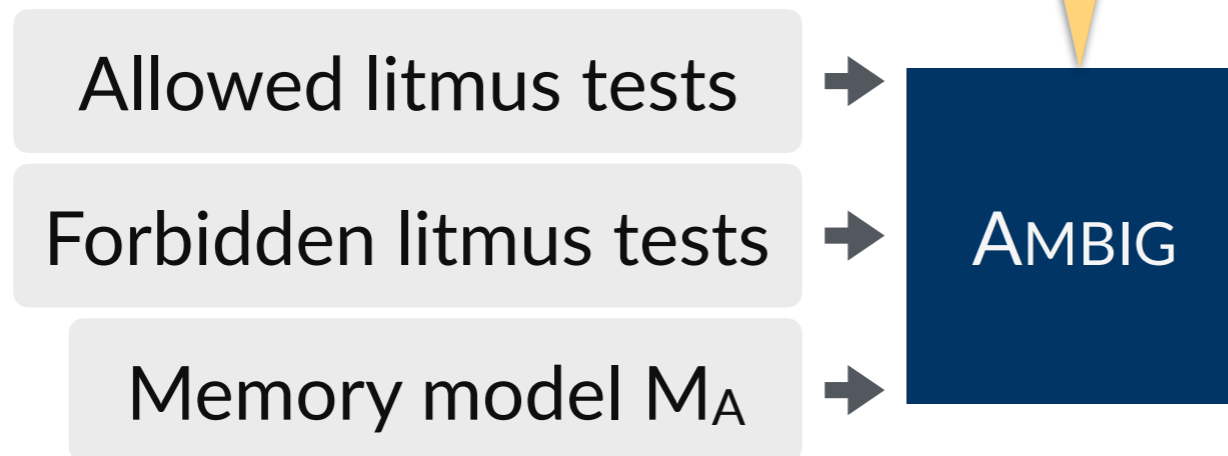
AMBIG



# Ambiguity

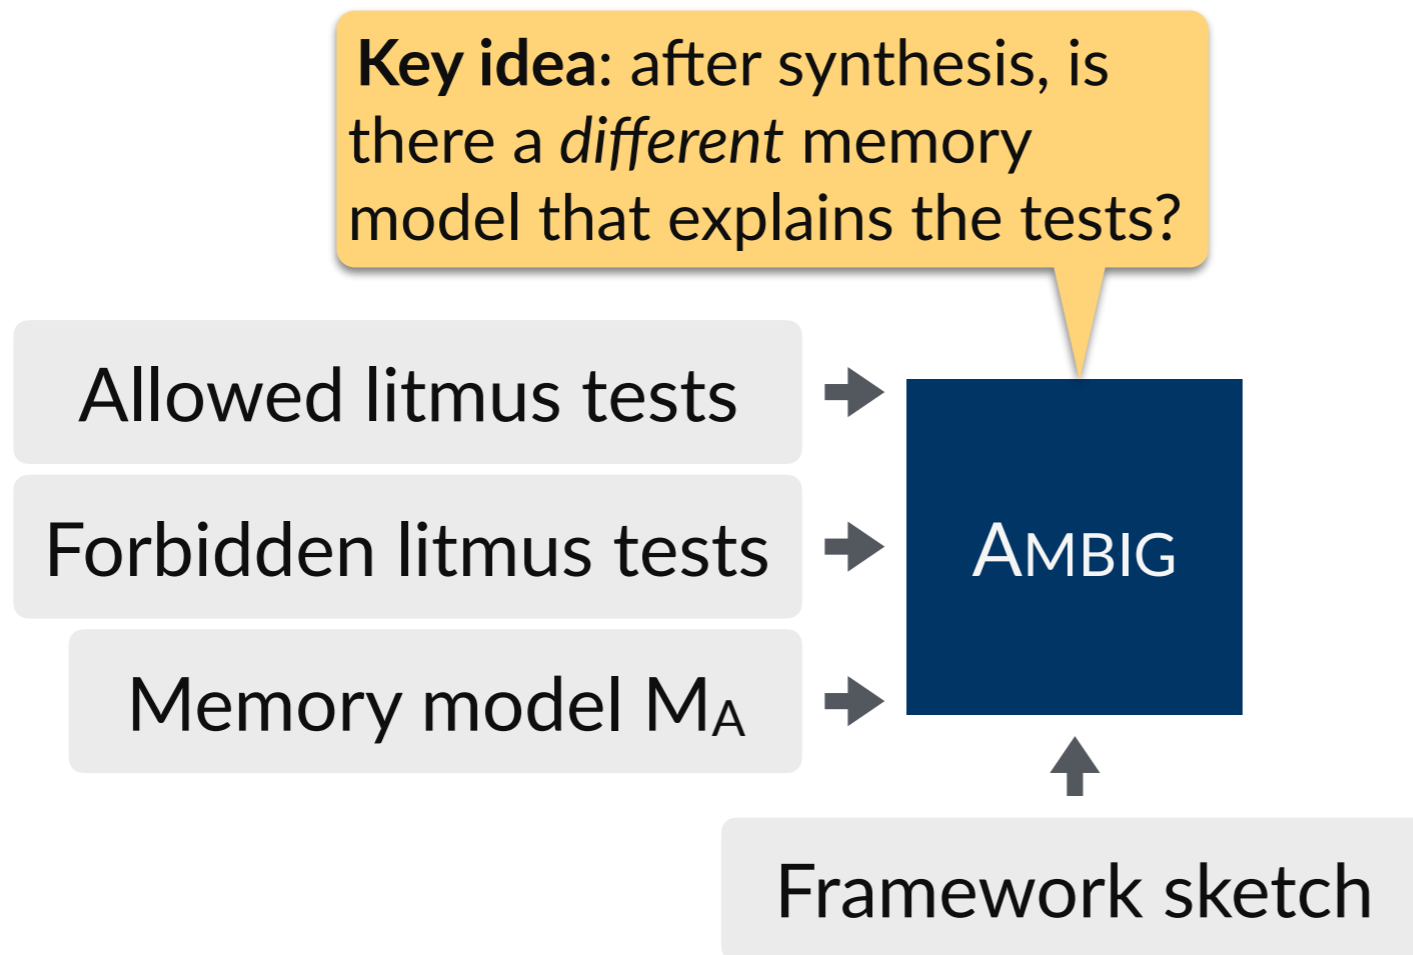
Find a distinguishing litmus test that exposes an ambiguity in a model

**Key idea:** after synthesis, is there a *different* memory model that explains the tests?



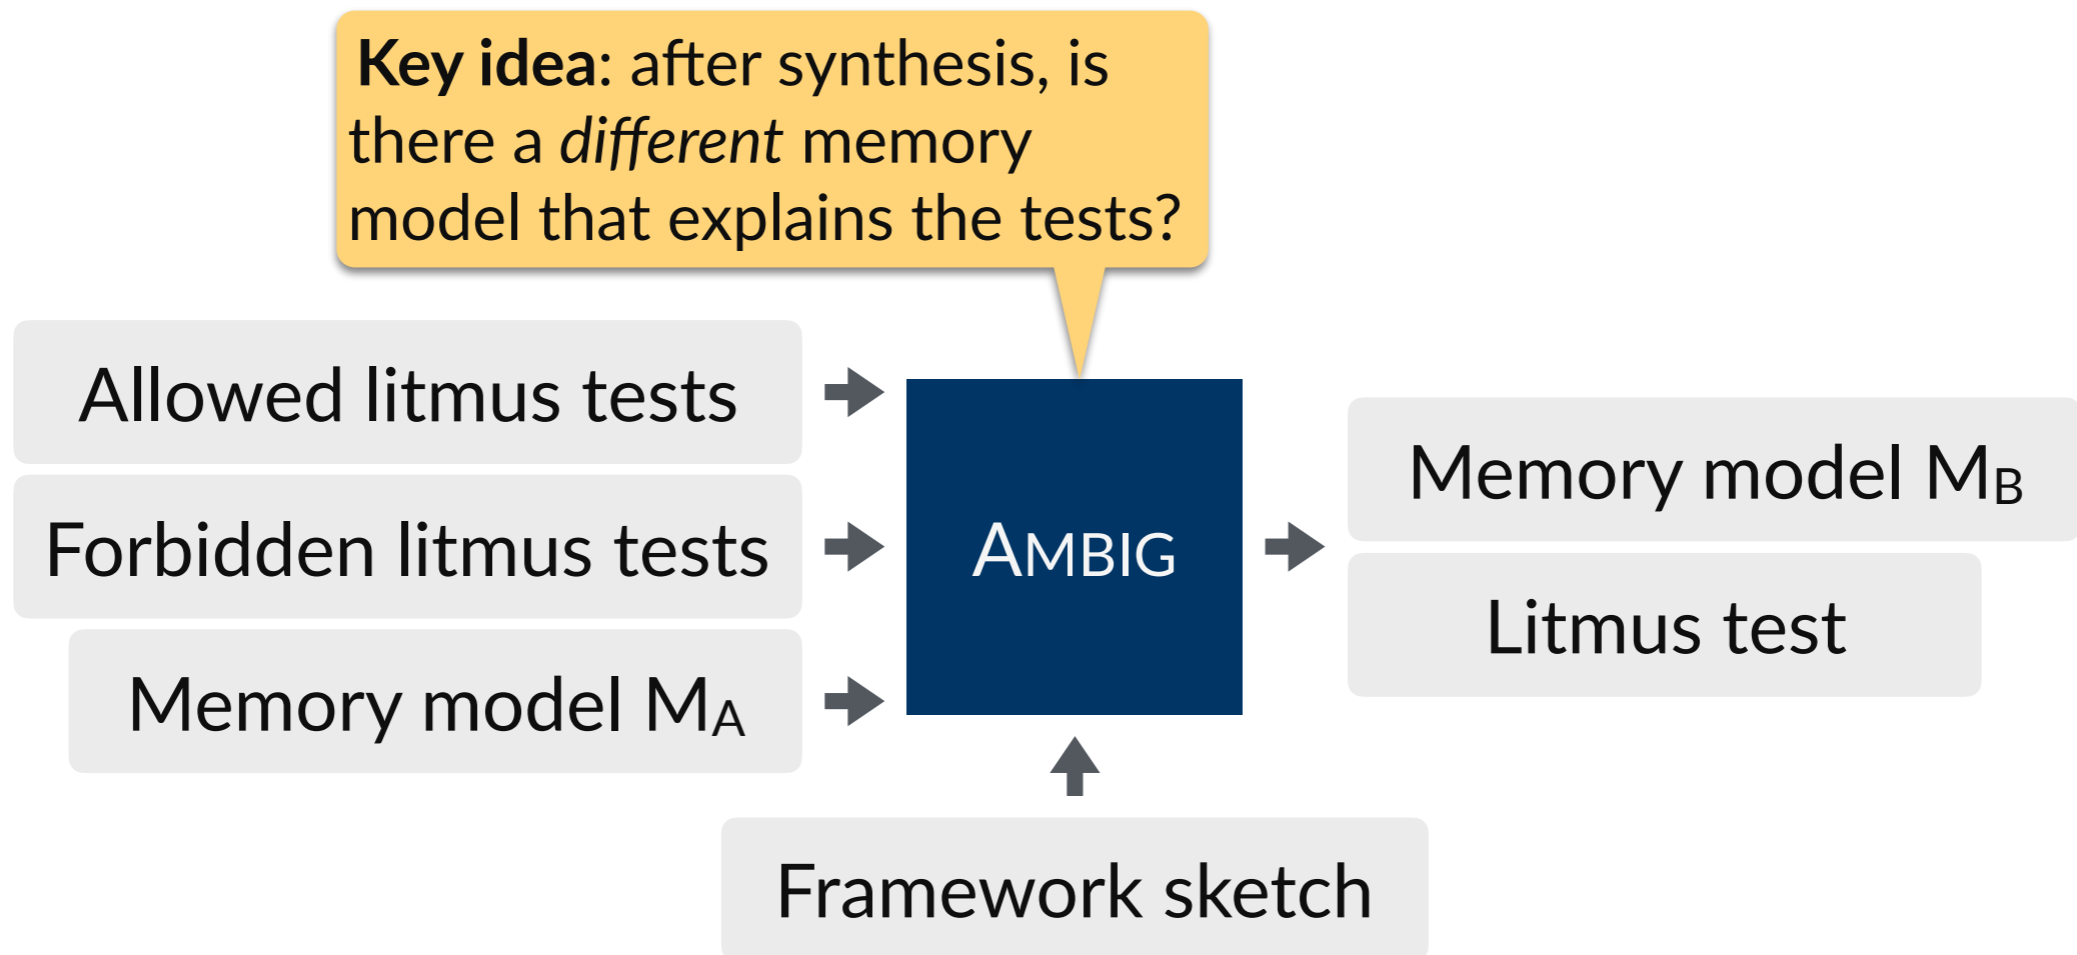
# Ambiguity

Find a distinguishing litmus test that exposes an ambiguity in a model



# Ambiguity

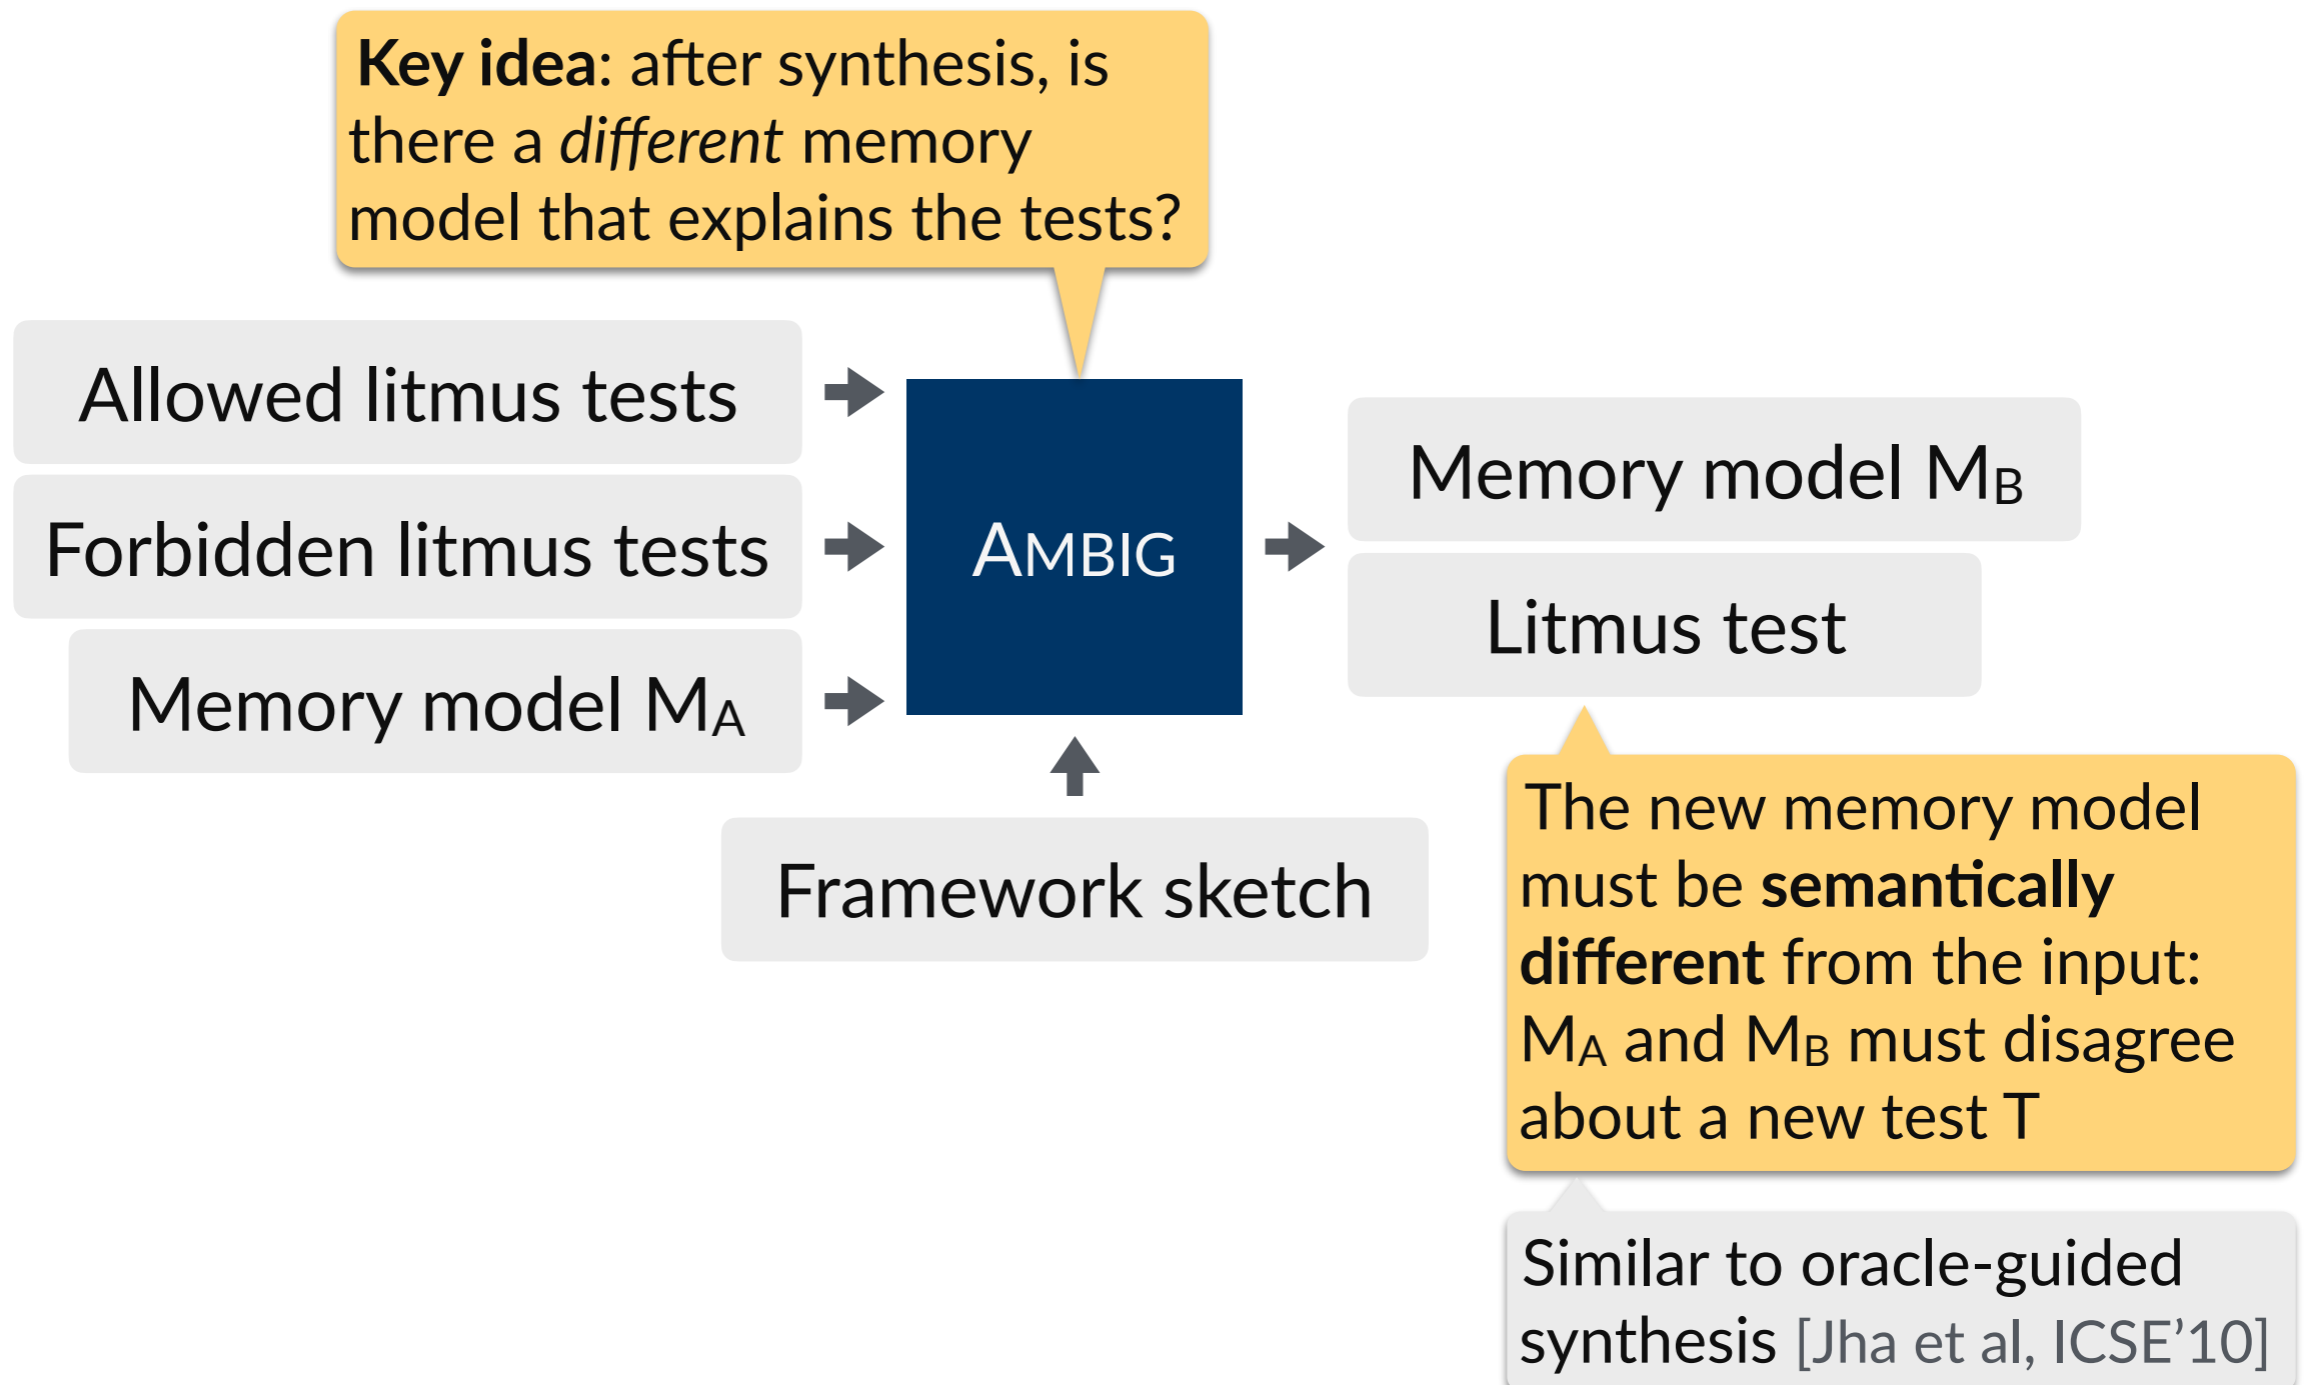
Find a distinguishing litmus test that exposes an ambiguity in a model





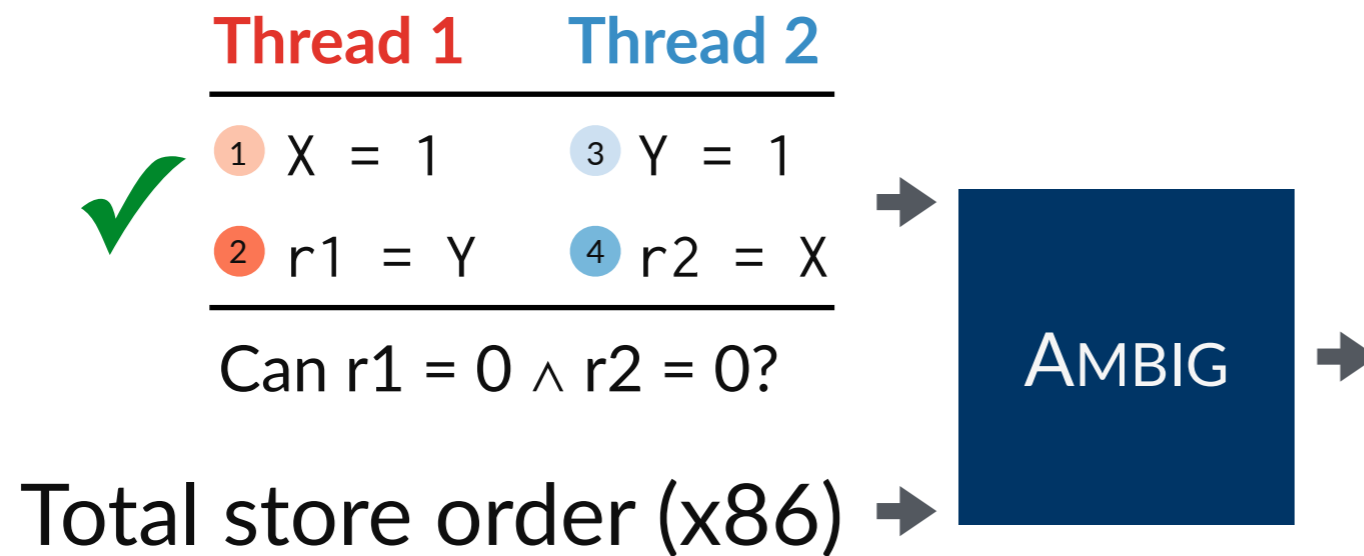
# Ambiguity

Find a distinguishing litmus test that exposes an ambiguity in a model



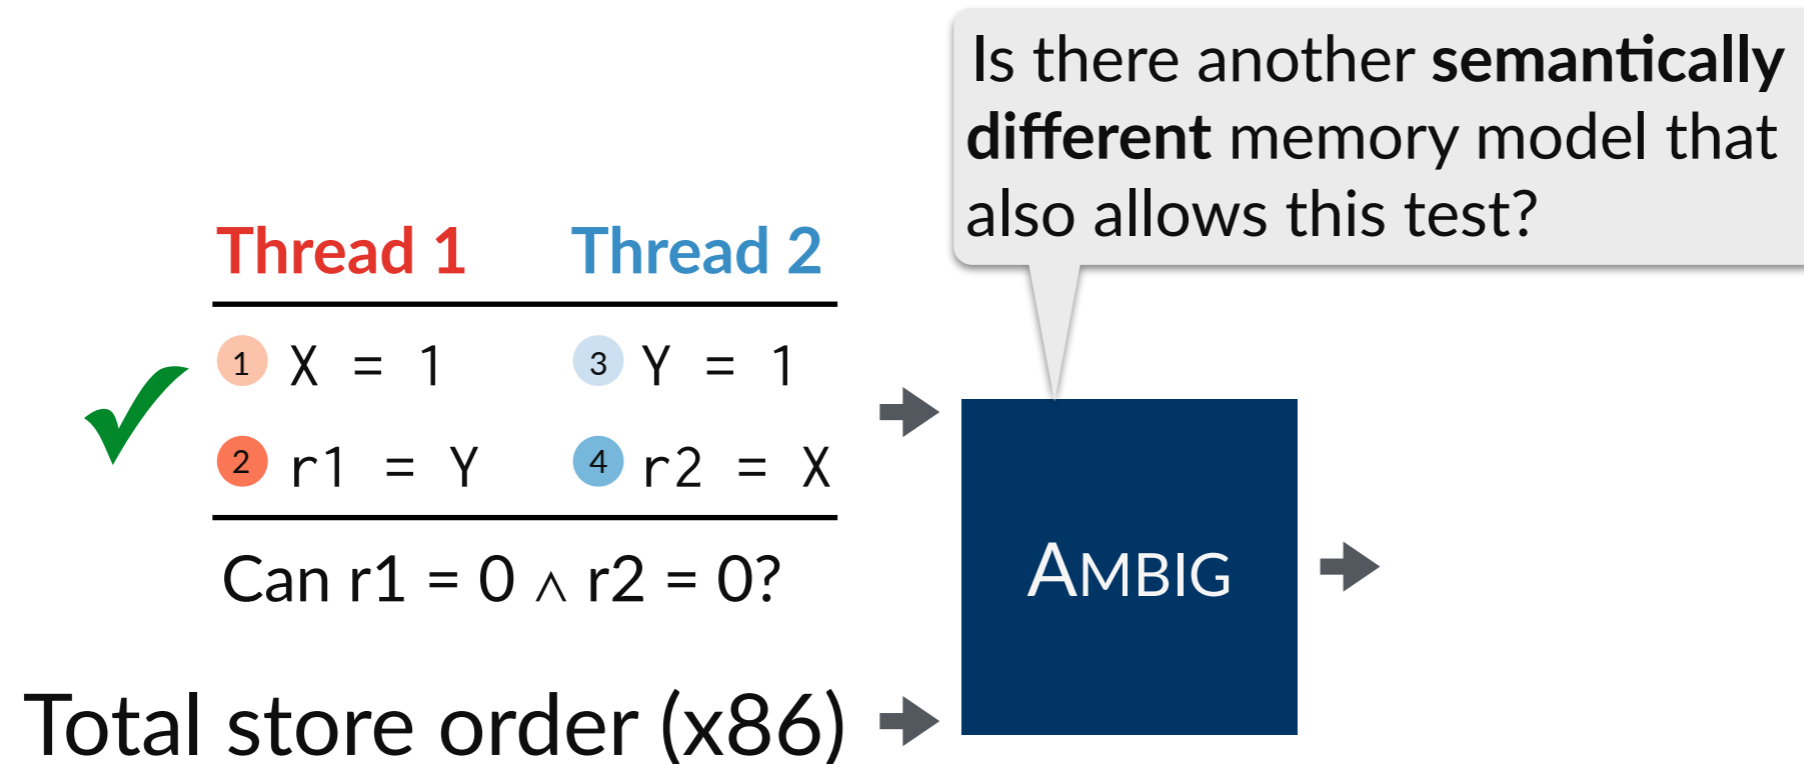
# Ambiguity

Find a distinguishing litmus test that exposes an ambiguity in a model



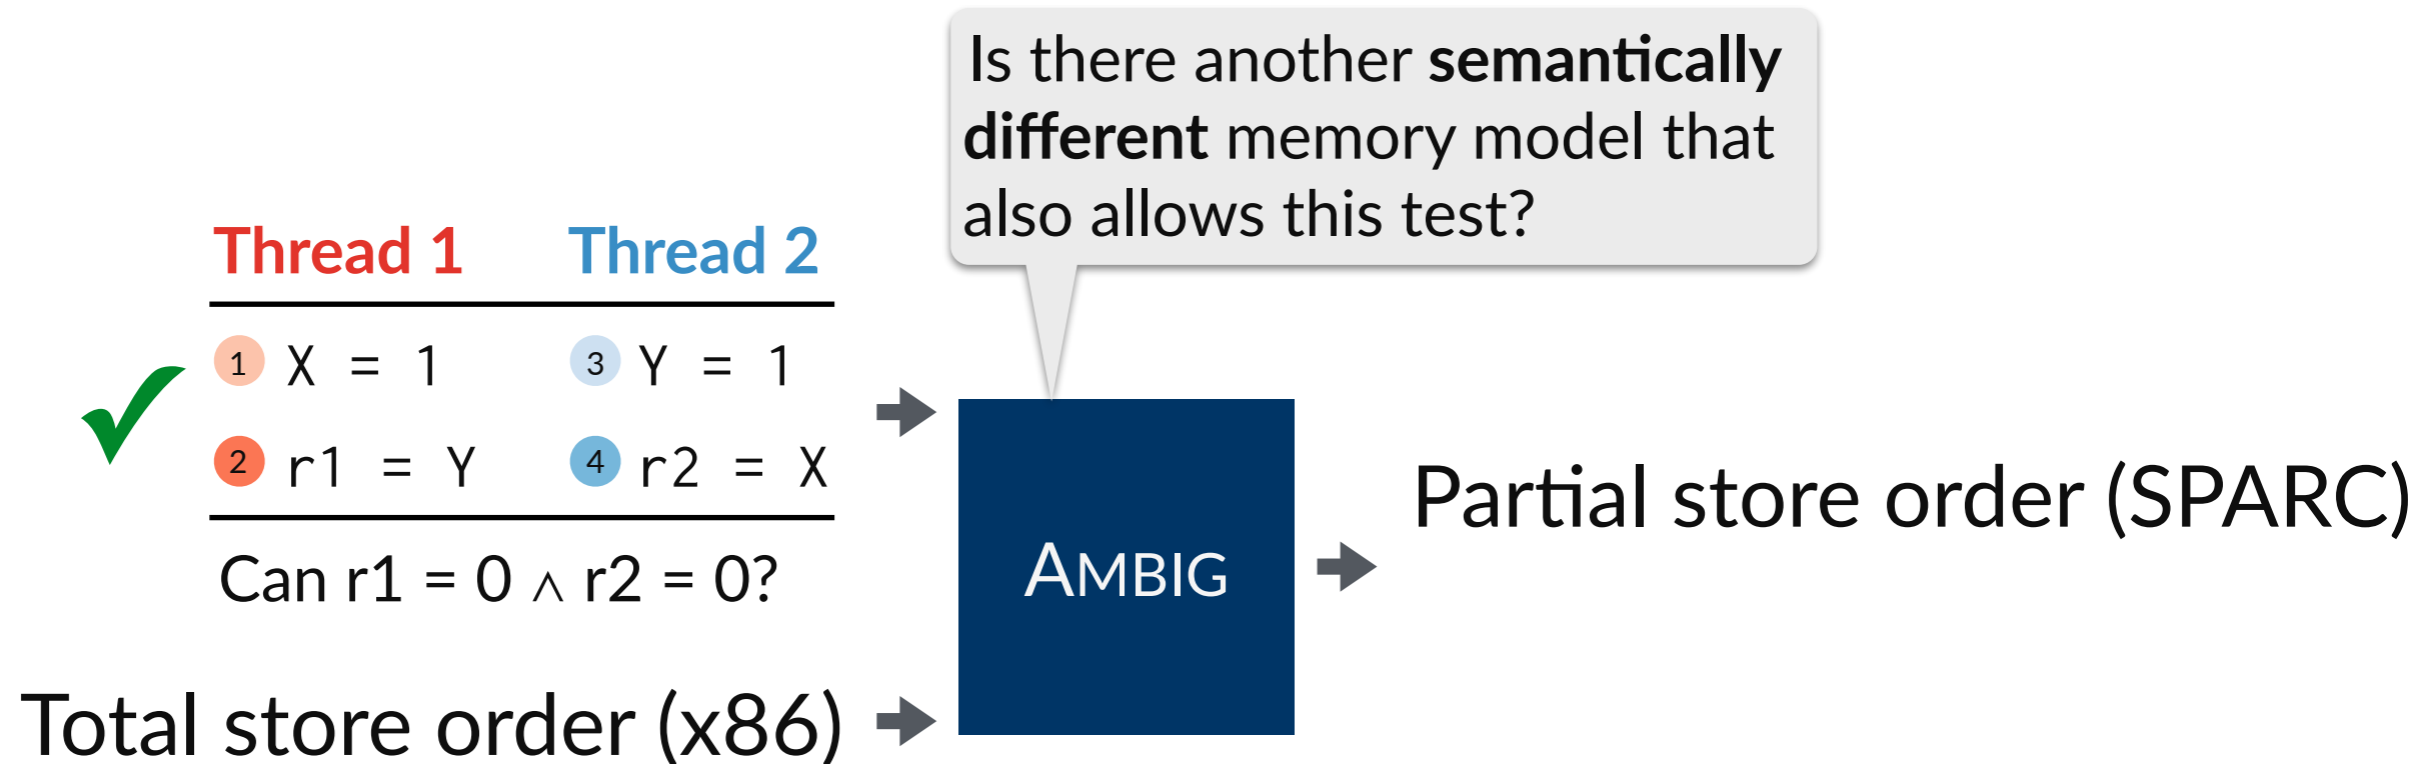
# Ambiguity

Find a distinguishing litmus test that exposes an ambiguity in a model



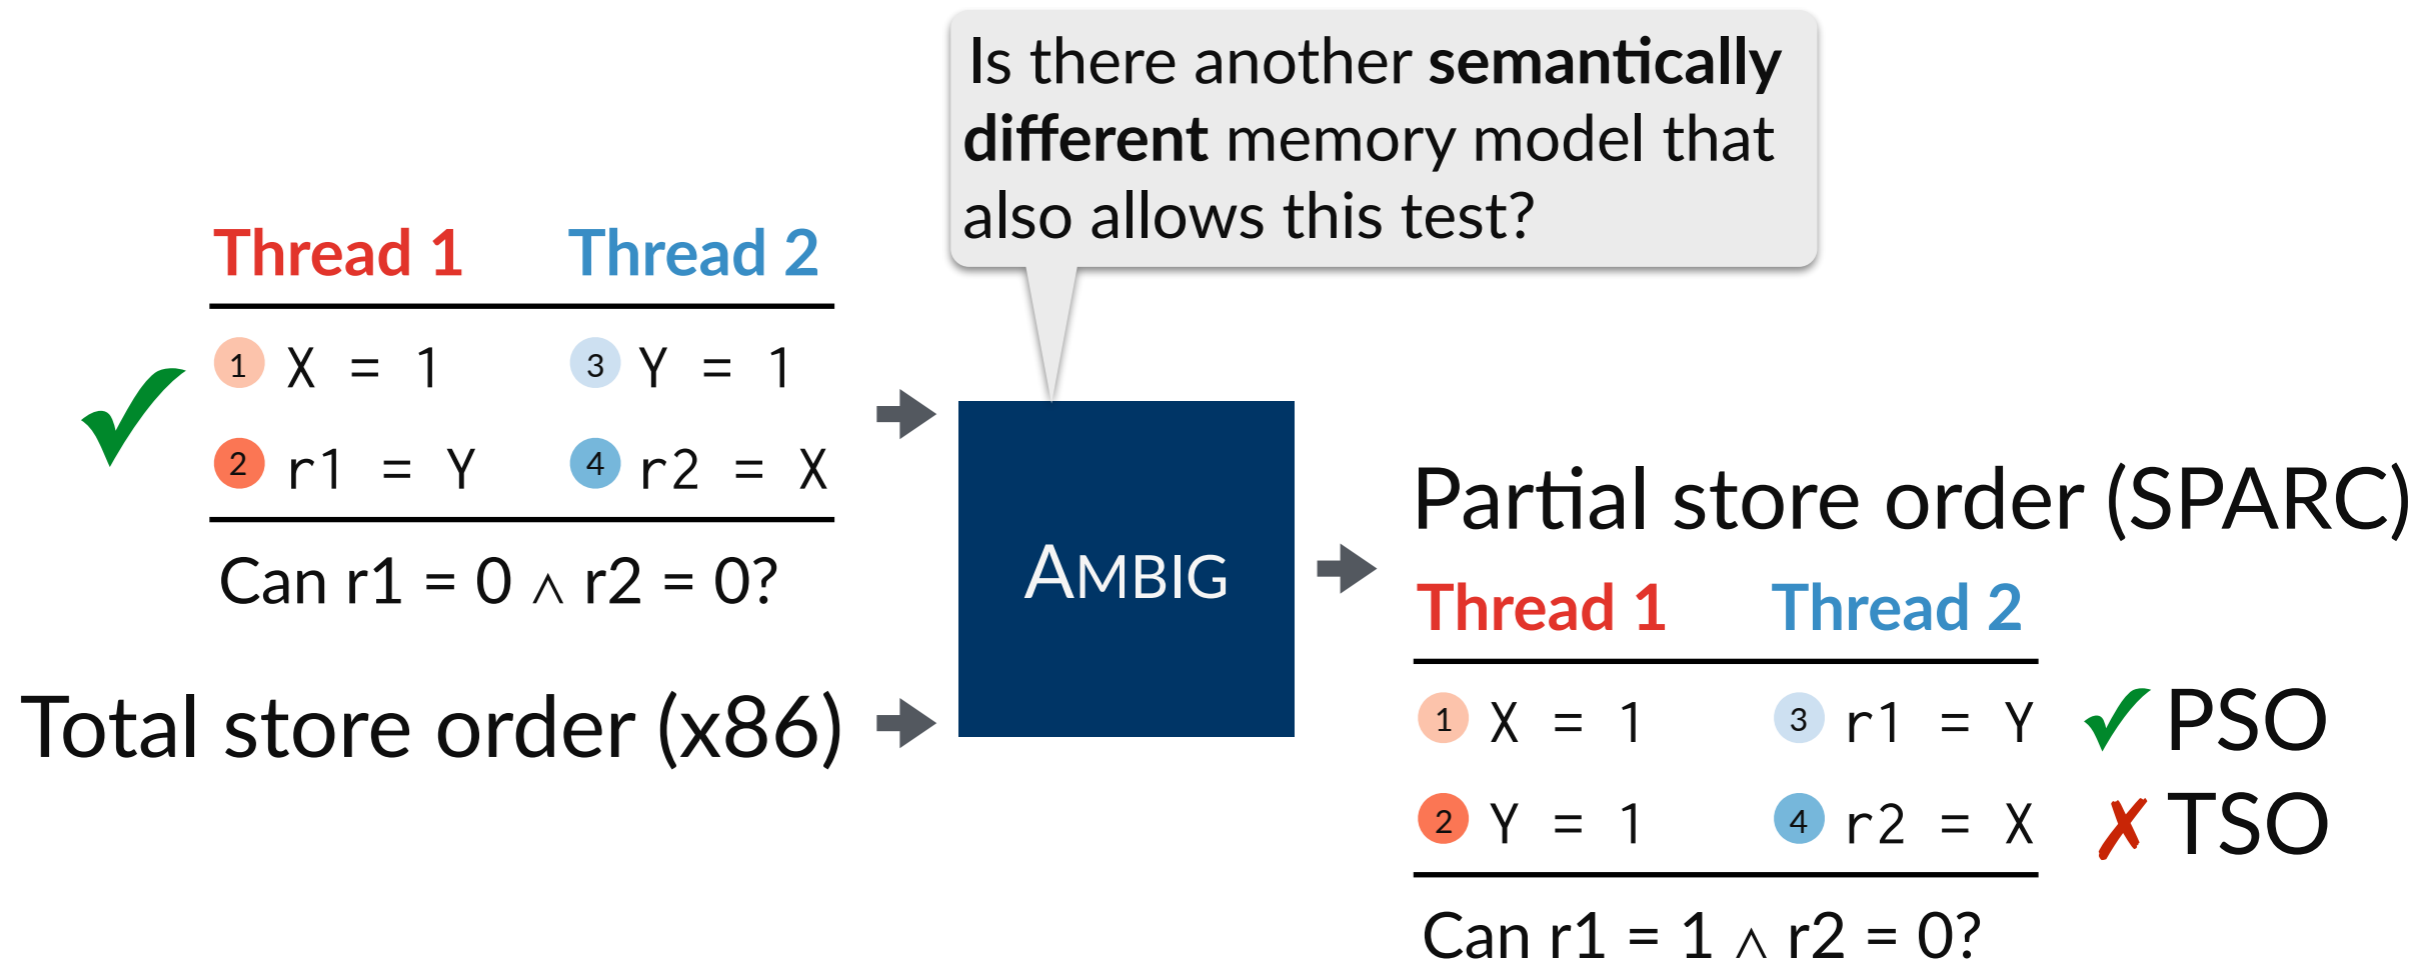
# Ambiguity

Find a distinguishing litmus test that exposes an ambiguity in a model




# Ambiguity

Find a distinguishing litmus test that exposes an ambiguity in a model



# The Synthesis-Ambiguity Cycle



**Litmus tests**

# The Synthesis-Ambiguity Cycle



Documentation

Random/systematic  
generation




Architects

3 5

1 2 4

## Litmus tests


# The Synthesis-Ambiguity Cycle




**Litmus tests**



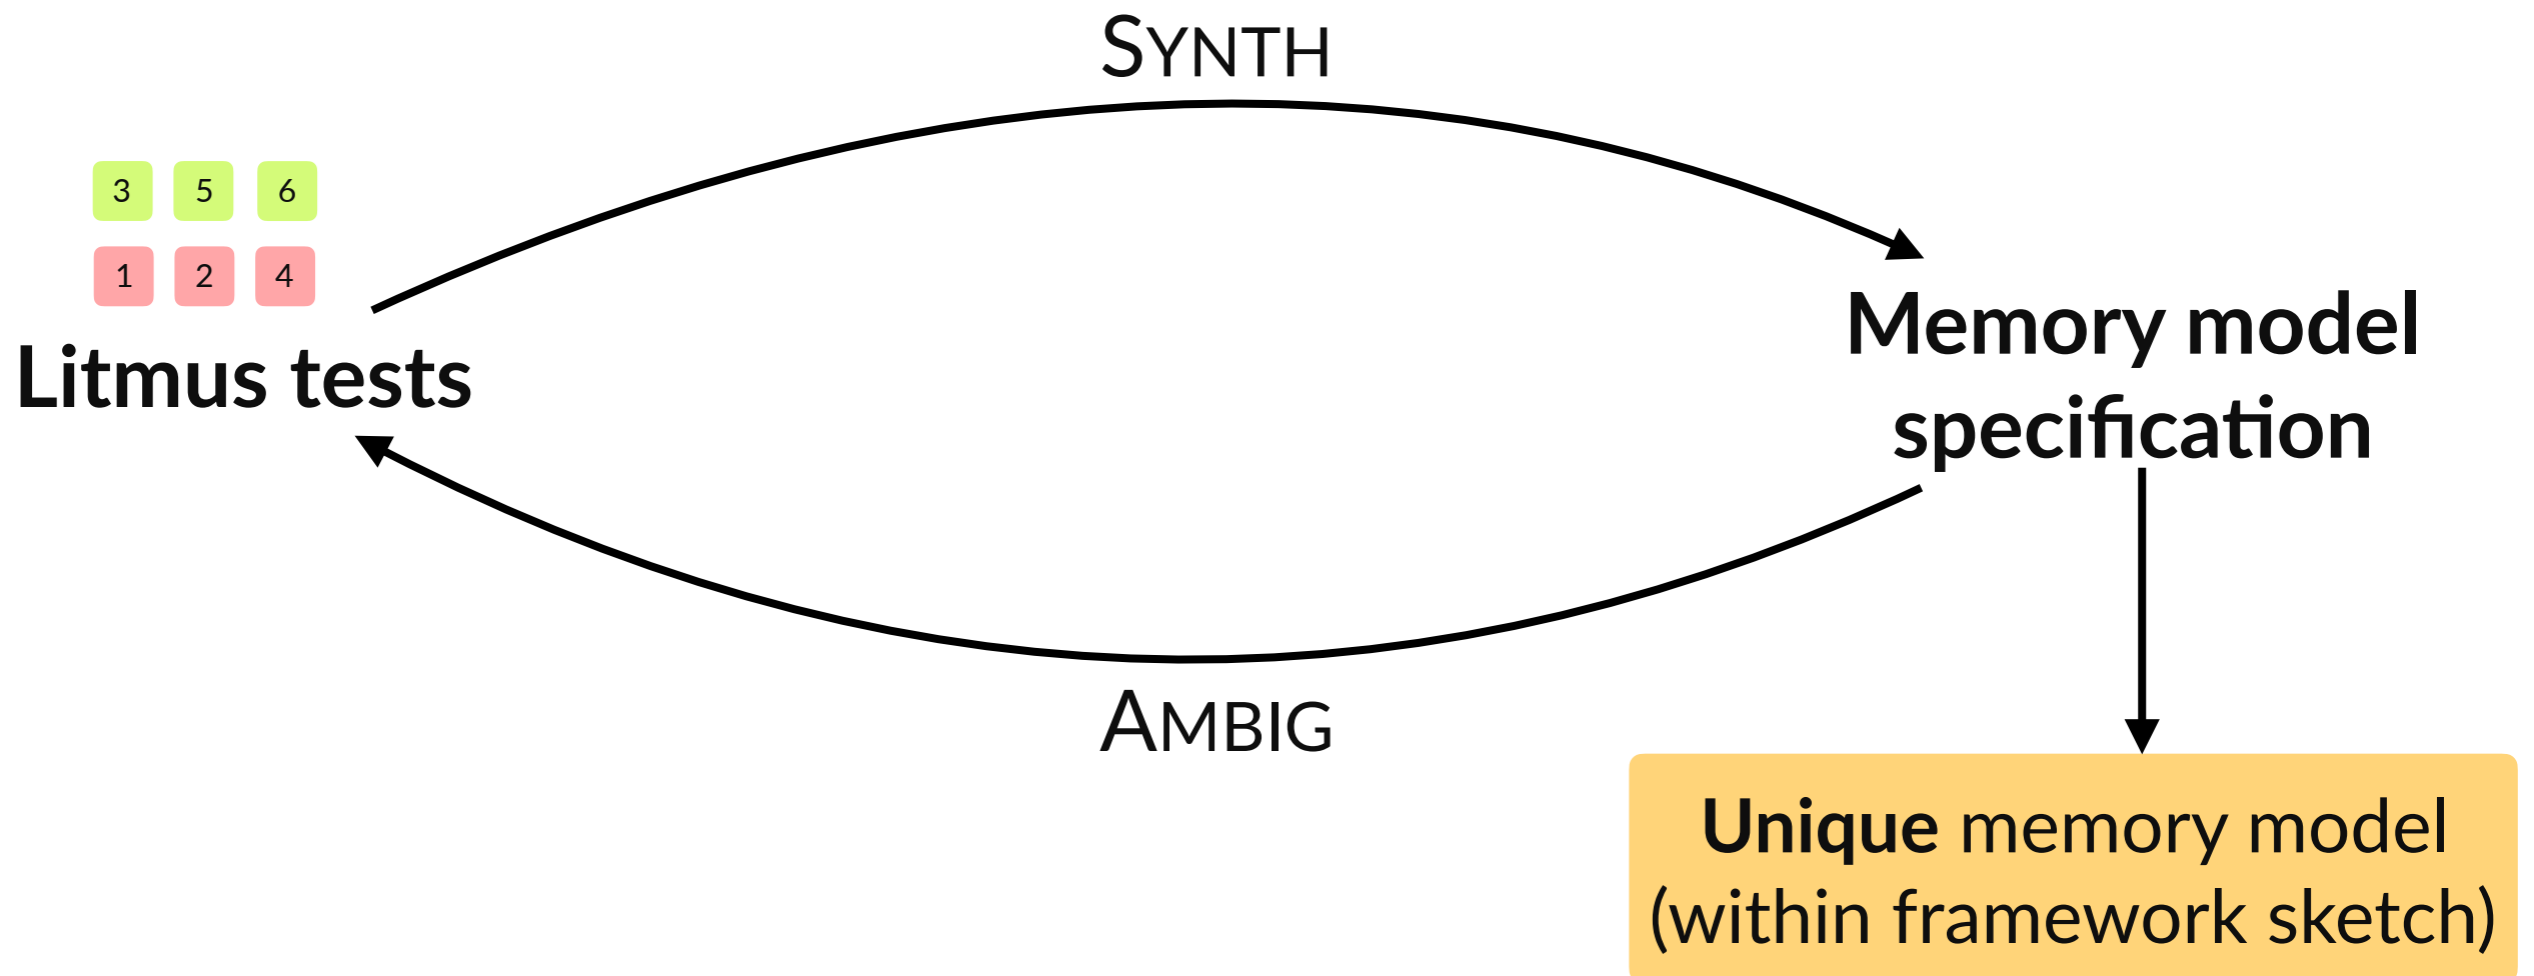
# The Synthesis-Ambiguity Cycle



# The Synthesis-Ambiguity Cycle



# The Synthesis-Ambiguity Cycle



# Results

# Synthesizing existing memory models

PowerPC

x86

# Synthesizing existing memory models

**PowerPC**

768 tests  
[Alglave et al, CAV'10]

**x86**

10 tests



# Synthesizing existing memory models

## Synthesis

**PowerPC**

768 tests  
[Alglave et al, CAV'10]

✓ 12 seconds  
Search space:  $2^{1406}$

**x86**

10 tests



✓ 2 seconds  
Search space:  $2^{624}$

# Synthesizing existing memory models

## Synthesis

PowerPC

768 tests  
[Alglave et al, CAV'10]

✓ 12 seconds

Search space:  $2^{1406}$

Not equivalent to  
published model!

x86

10 tests



✓ 2 seconds

Search space:  $2^{624}$



# Synthesizing existing memory models

## Synthesis

PowerPC

768 tests  
[Alglave et al, CAV'10]

✓ 12 seconds

Search space:  $2^{1406}$

Not equivalent to  
published model!

x86

10 tests




✓ 2 seconds

Search space:  $2^{624}$

Not equivalent to  
TSO!

# Synthesizing existing memory models

		<u>Synthesis</u>	<u>Ambiguity</u>
<b>PowerPC</b>	768 tests [Alglave et al, CAV'10]	✓ 12 seconds Search space: $2^{1406}$ <b>Not equivalent to published model!</b>	9 new tests sync, lwsync, etc.
<b>x86</b>	10 tests 	✓ 2 seconds Search space: $2^{624}$ <b>Not equivalent to TSO!</b>	4 new tests mfence, xchg

# Other results

Implemented another **framework sketch** [Mador-Haim et al, DAC'11]

Found typo in paper; couldn't fix by hand, but synthesized repair

# Other results

Implemented another **framework sketch** [Mador-Haim et al, DAC'11]

Found typo in paper; couldn't fix by hand, but synthesized repair

Order of magnitude faster than the Alloy **general-purpose relational solver** for verification and equivalence

Ocelot offers finer-grained control over relational constraints

# Other results


Implemented another **framework sketch** [Mador-Haim et al, DAC'11]

Found typo in paper; couldn't fix by hand, but synthesized repair

Order of magnitude faster than the Alloy **general-purpose relational solver** for verification and equivalence

Ocelot offers finer-grained control over relational constraints

Comparable performance to existing **custom memory model tool** for verification (Herd [Alglave et al, CAV'10])



## Framework sketches

define a class of memory models

## MemSynth engine

verification, equivalence, synthesis, ambiguity

## Results

synthesize real-world memory model specs

[memsynth.uwplse.org](http://memsynth.uwplse.org)