

Weak-memory
local reasoning

Ian Wehrman

September 26, 2012

Overview

Single-threaded program behavior w.r.t. an idealized computer model is complex.

Multi-threaded program behavior w.r.t. a **realistic** computer model is *really* complex.

Memory models

Specify interaction between programs and memory.

Description:

Notion of state (an abstract representation of memory);

Explanation of how values are read from/written in a given state.

Memory models

Different programs require different MMs:

Sequential imperative programs w/statically allocated memory:

$\text{State} \triangleq \text{Stack}$ where $\text{Stack} \triangleq \text{Variable} \rightarrow \text{Value}$

Sequential (or well-locked concurrent) imperative programs with dynamically allocated memory:

$\text{State} \triangleq \text{Stack} \times \text{Heap}$ where $\text{Heap} \triangleq \text{Address} \rightarrow \text{Value}$

Racy concurrent imperative programs:

$\text{State} \triangleq \dots$ depends on the architecture.

Racy programs

Not all racy programs are broken:

e.g., lock-free concurrent data structures.

x86 MM

A weak, x86-like memory model:

State \triangleq Stack \times Heap \times WriteBufferArray \times Lock

WriteBufferArray \triangleq Processor \rightarrow WriteBuffer

WriteBuffer \triangleq Queue[Write]

Write \triangleq Address \times Value

Lock \triangleq Processor + \perp

x86 MM

On processor i :

`store` enqueues a new write on i^{th} buffer;

`load` returns value of most recent write in i^{th} buffer;
if none, then value in heap;

`fence` flushes all writes on i^{th} buffer to the heap;

acquire (`lock`) or release (`unlock`) the global lock.

all but `store` block while $j \neq i$ holds lock.

Writes may commit nondeterministically!

Hoare logic

Program specifications: $\vdash \{ P \} c \{ Q \}$

command c is a sequential static-memory command;

precondition P describes an initial set of states;

postcondition Q describes a final set of states.

Meaning:

if c executes from a P -state, it terminates in a Q -state or diverges.

Separation logic

Extension of Hoare Logic: $\vdash \{ P \} c \{ Q \}$

enables sound reasoning about dynamic-memory commands;

additional assertions used to describe heap values;

all proved programs are memory-safe.

Concurrent separation logic

An extension of separation logic: $J \vdash \{ P \} c \{ Q \}$

c is a *concurrent* dynamic-memory command;

P and Q describe thread-private states;

invariant J describes environment-shared states;

all proved programs are well-locked and race-free.

(Required by simple memory model!)

Project

Goal: A program logic with an x86-like model.

Why?

Existing logics insufficient or unsound for racy programs.

Eventually wish to prove racy programs correct.

Explore concurrent reasoning in weak vs. strong MMs.

Project

Result: an x86-like variant of CSL.

Components:

- 1) a programming language;
- 2) an assertion logic;
- 3) a specification logic.

Project components

(1/3) Programming language:

C-like w/assignment, load, store, fence & locking primitives.

x86-like semantics.



Project components

(2/3) Assertion logic:

Like the assertion language of SL/CSL, but more expressive.



Describe heaps **and** write buffers **and** the global lock.

x86-like semantics.



Ideally also a proof theory, but that's future work.

Project components

(3/3) Specification logic:

CSL-like specifications.

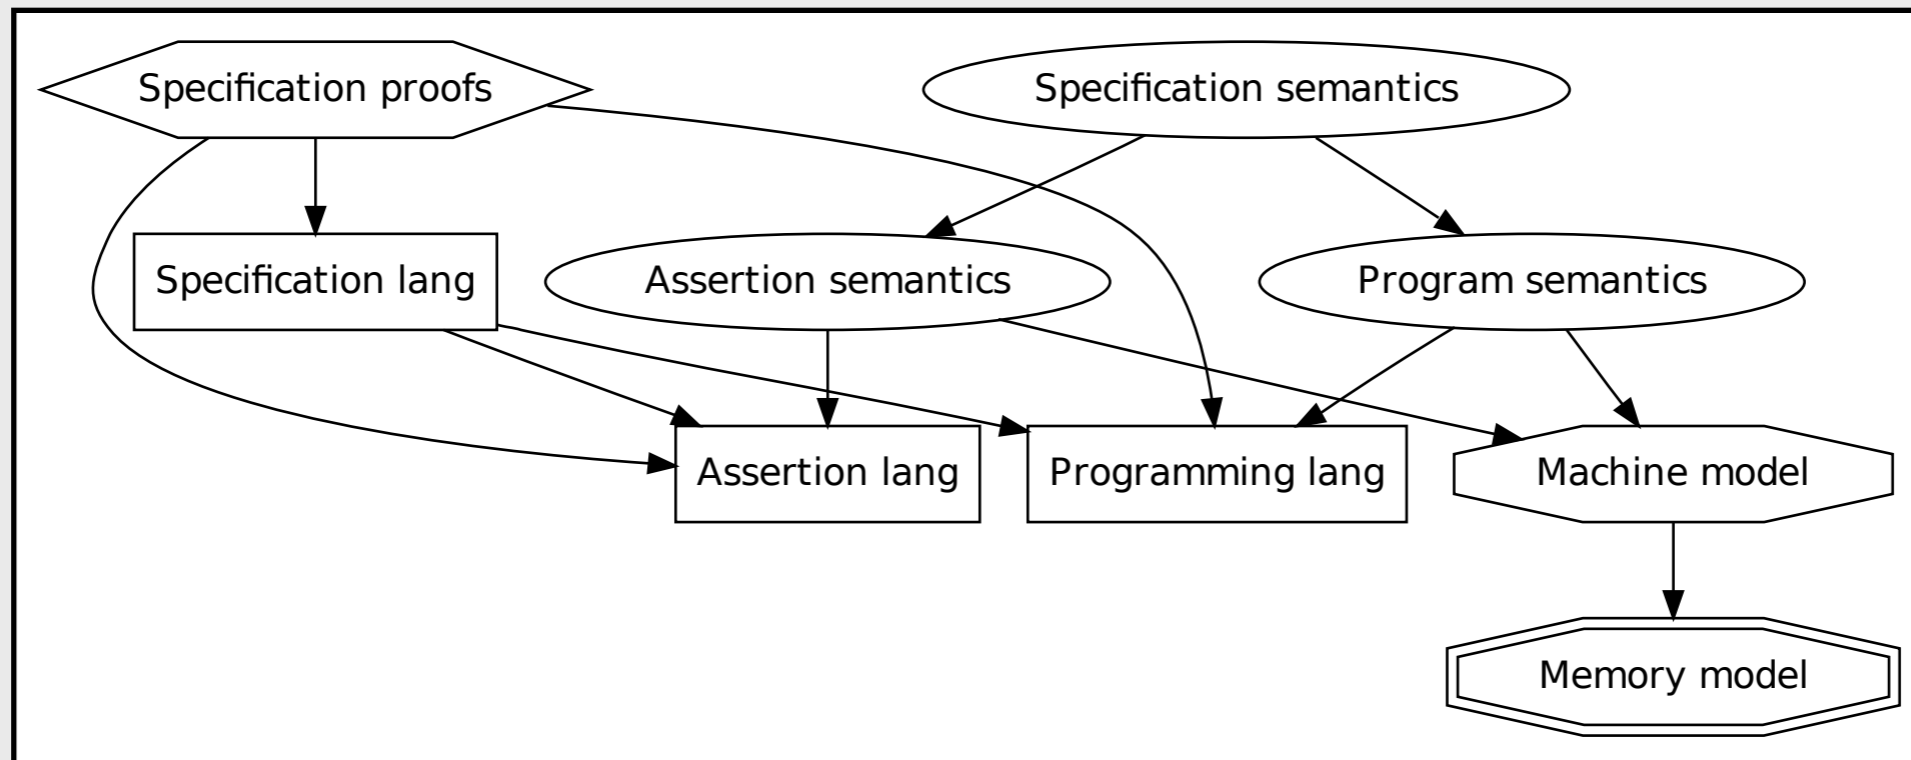
CSL-like proof theory, but with x86-specific adjustments.



x86-like semantics.



Project components



Agenda

Assertion language and models:

Language extends FOL and SL/CSL language;

New formulas for new state elements.

Design constraints from specification logic:

Expressive enough to describe x86 commands;

Constrained enough for sound, local reasoning.

Local reasoning

The **big idea** in SL/CSL:

Restrict reasoning to a small, relevant part of system state;

Then generalize to a complete description of system state.

Embodied by the **frame rule**:

$$\frac{J \vdash \{P\} c \{Q\}}{J \vdash \{R * P\} c \{R * Q\}}$$

Separation

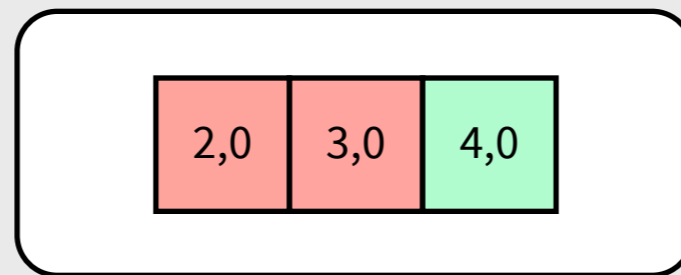
In SL/CSL:

$P * Q$ is the **separating conjunction** of assertions P and Q .

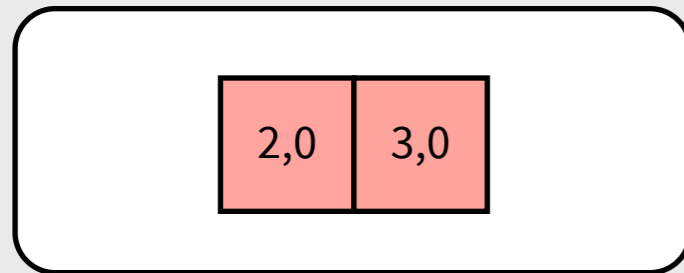
Describes heaps that can be partitioned into sub-heaps:

a sub-heap described by P and a sub-heap described by Q .

Separation



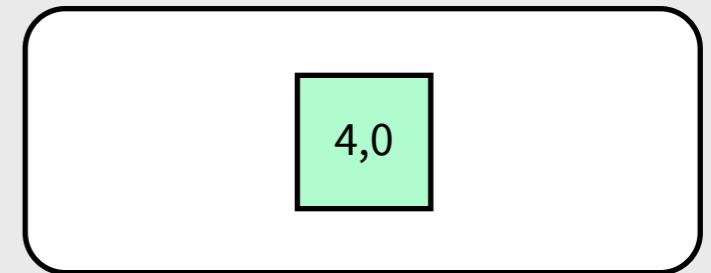
h



h_0

$*$

h_1



Spatial separation

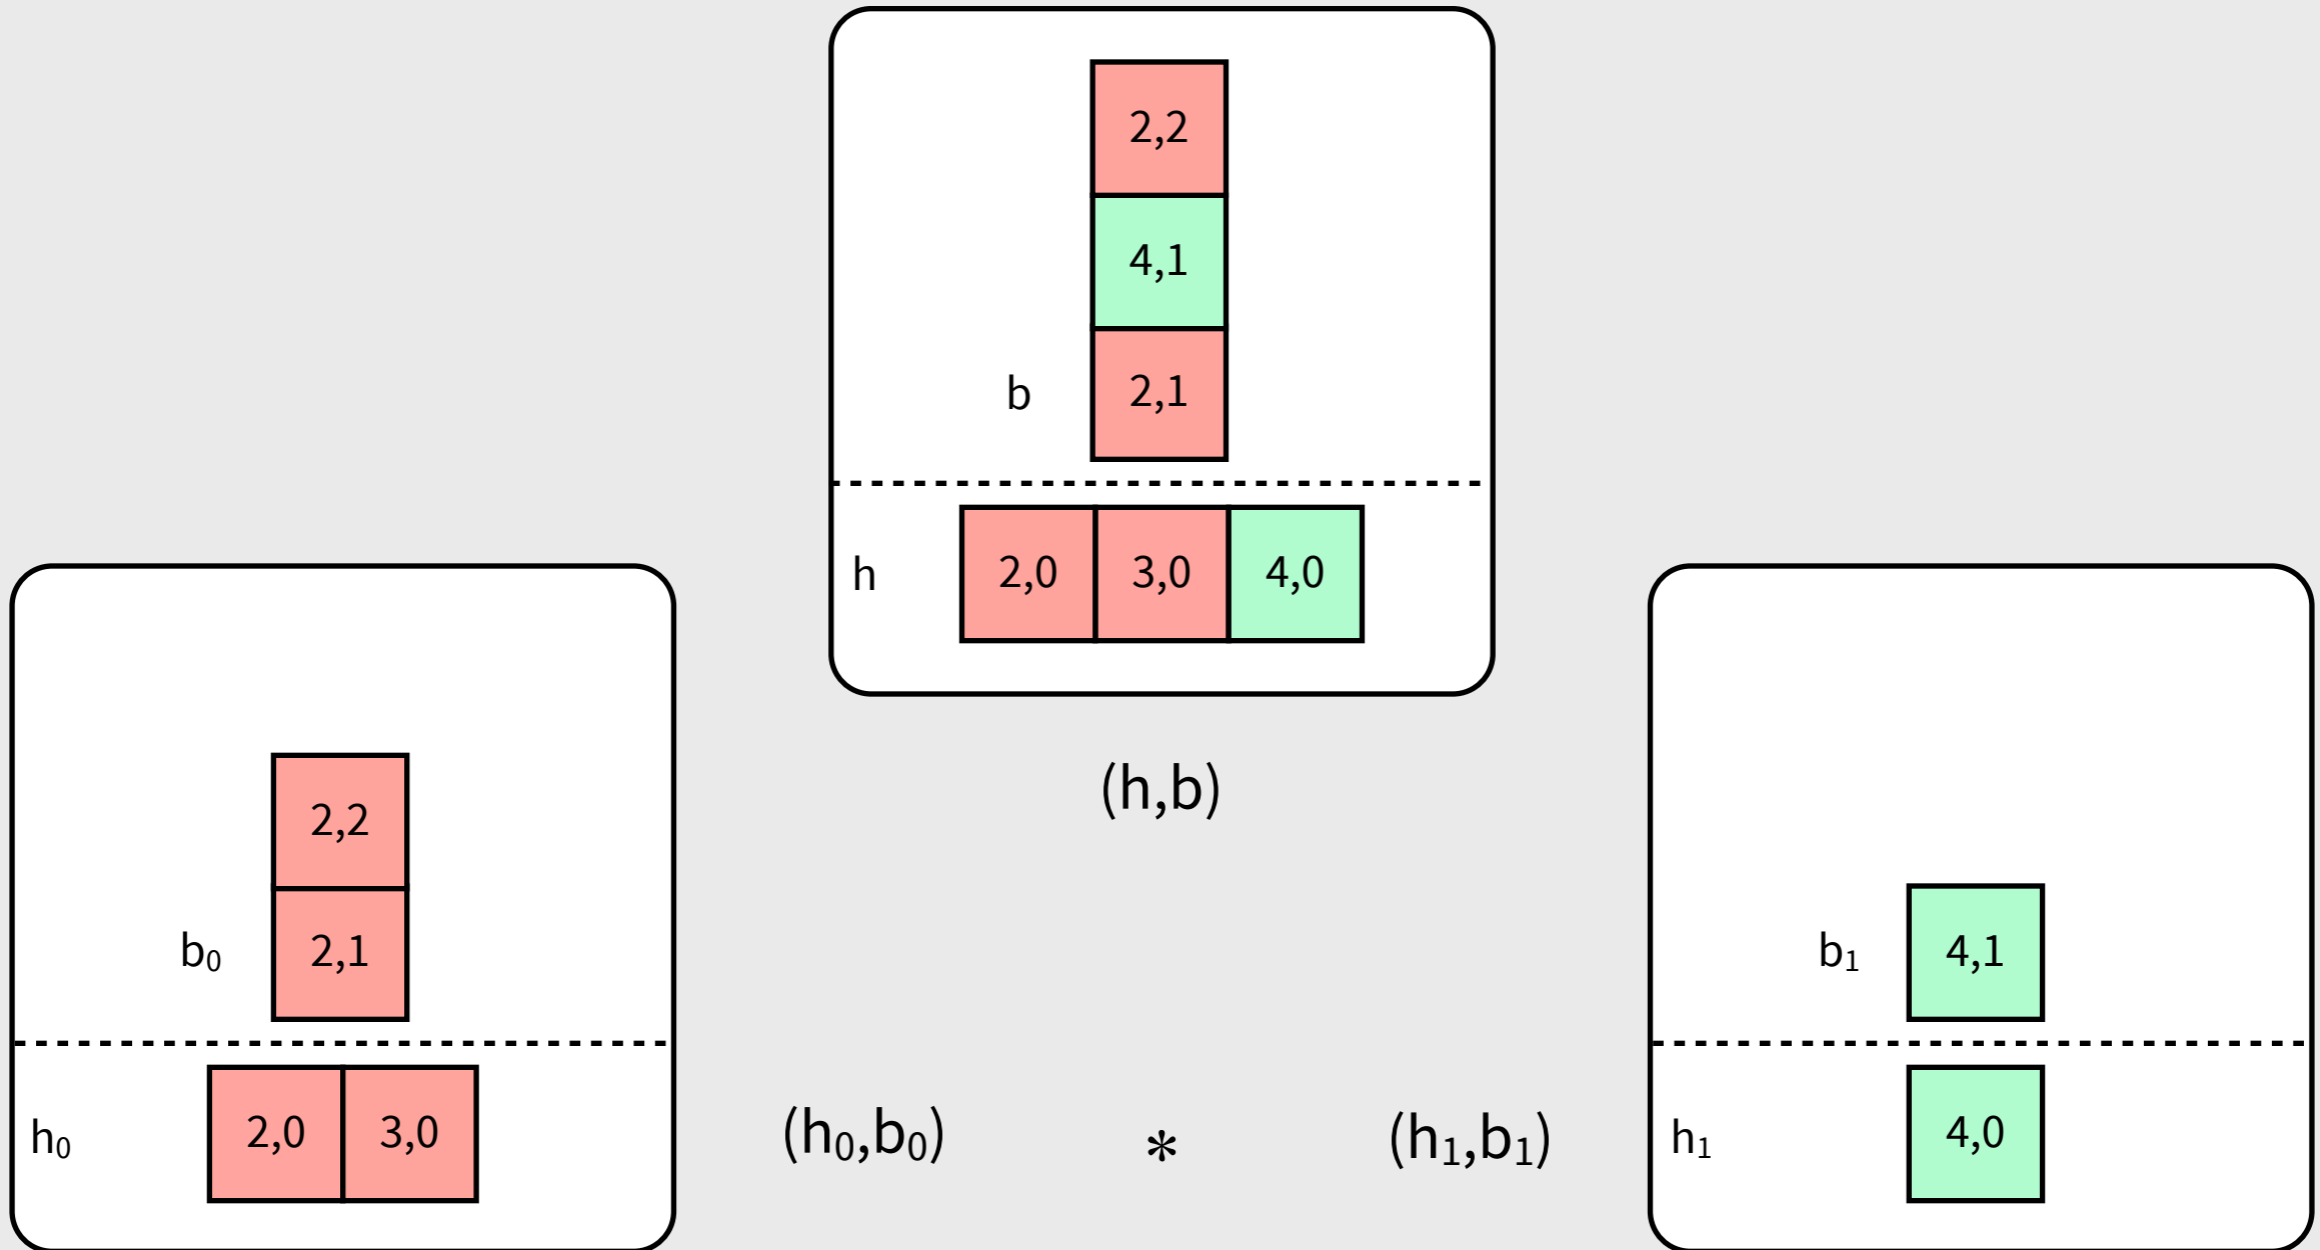
In x86-CSL:

$P * Q$ is called the **spatial separating conjunction** of P and Q .

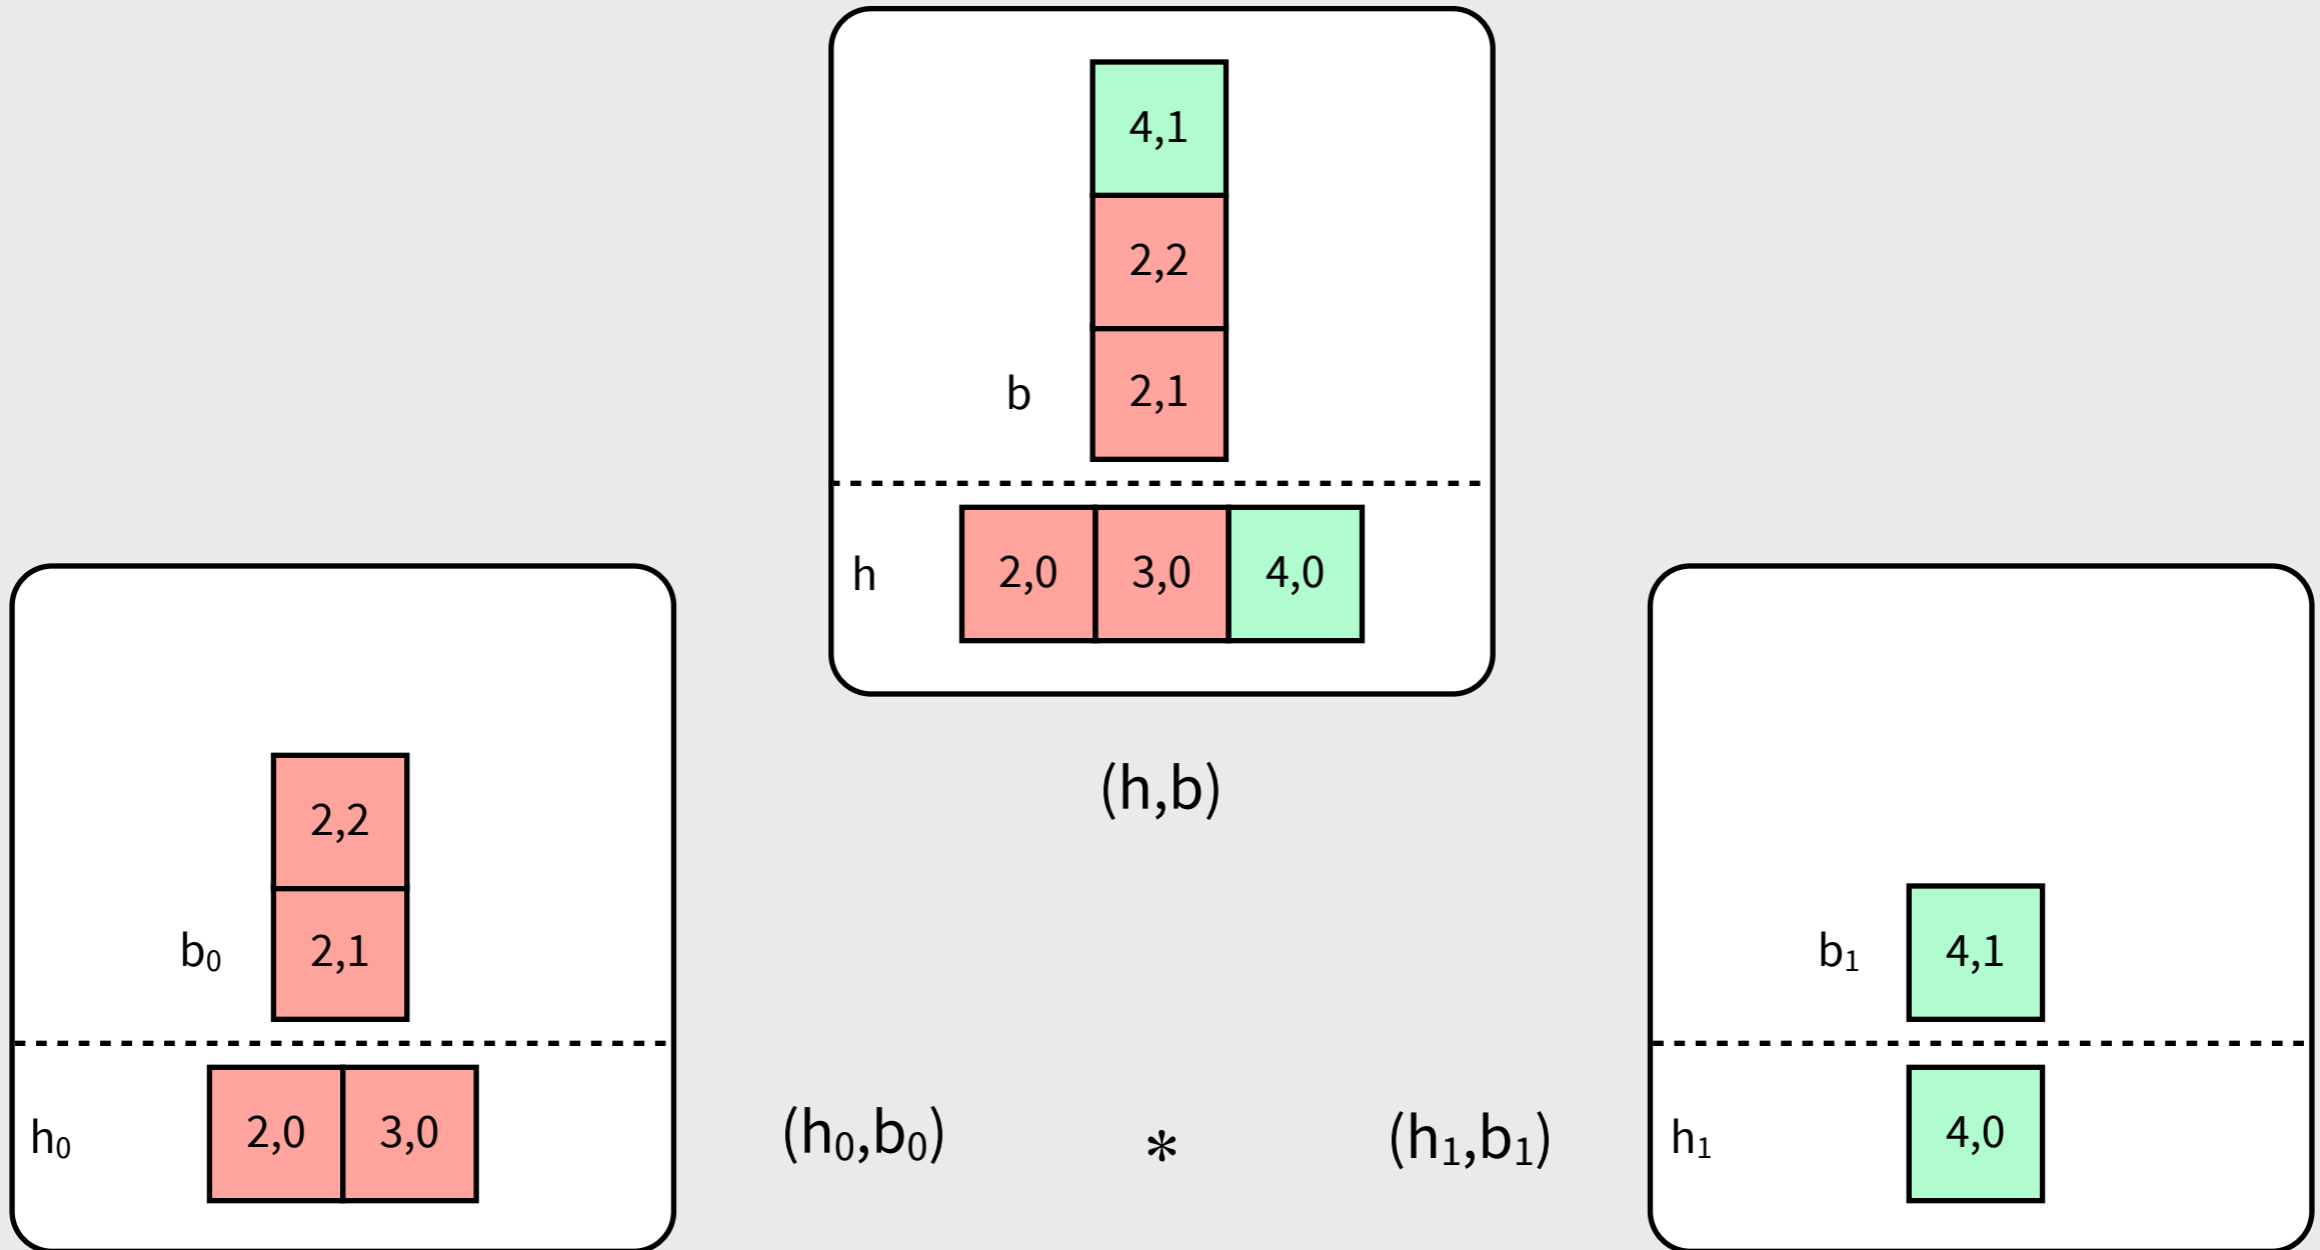
Describes x86 states that are separable by address:

a sub-state described by P and a sub-state described by Q .

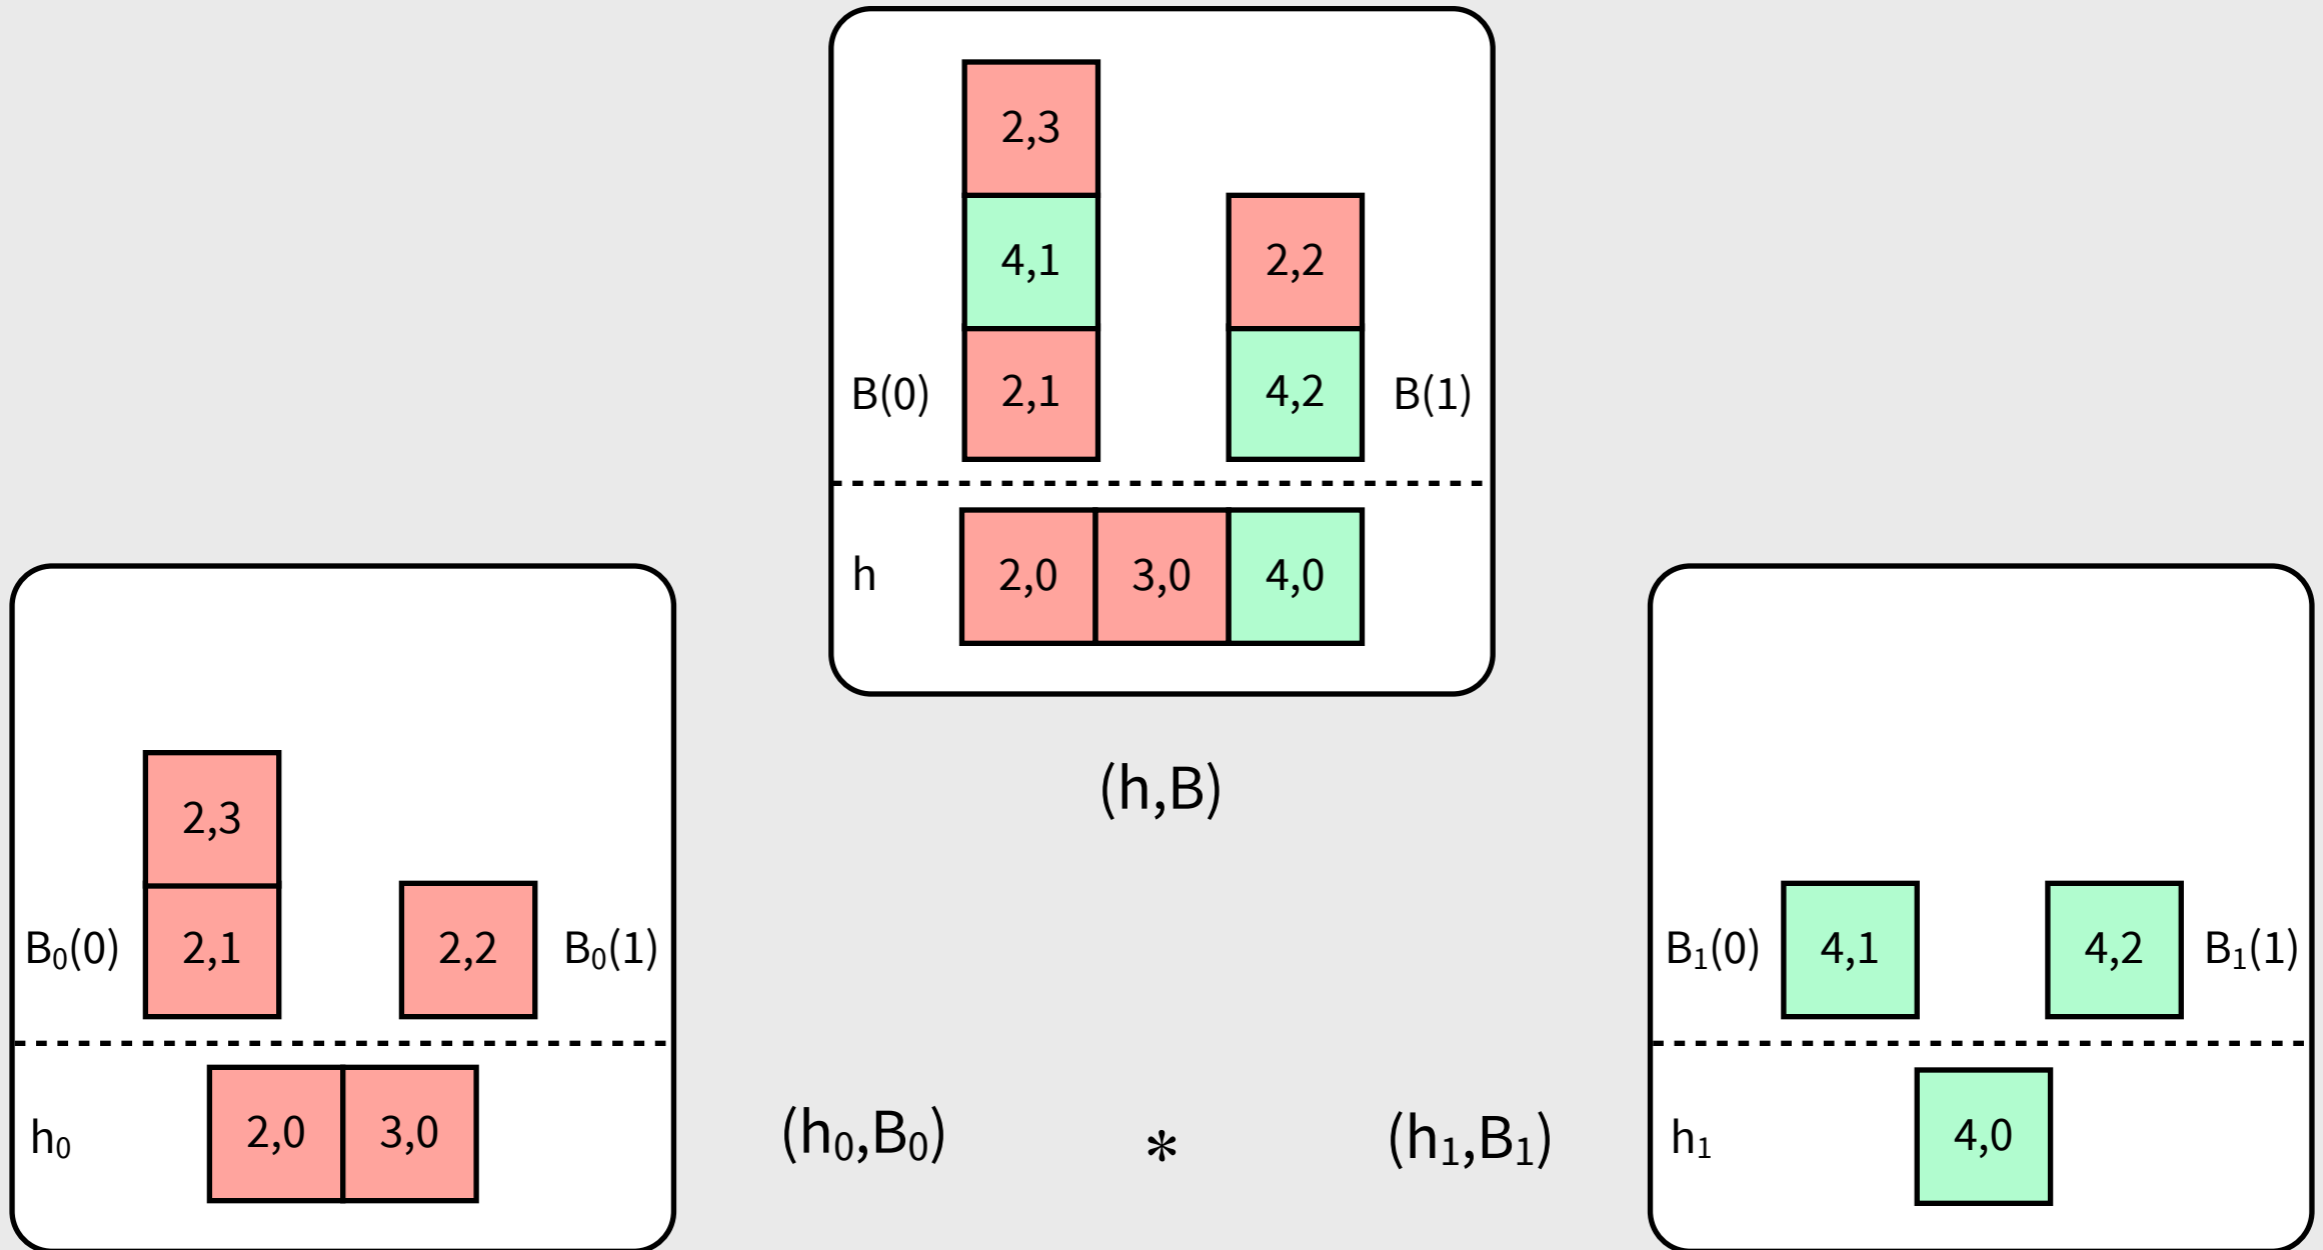
Spatial separation



Spatial separation



Spatial separation

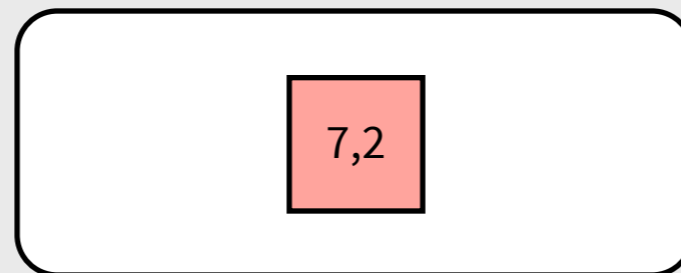


Heap values

In SL/CSL:

The **points-to assertion** describes a **single heap value**.

$7 \mapsto 2$

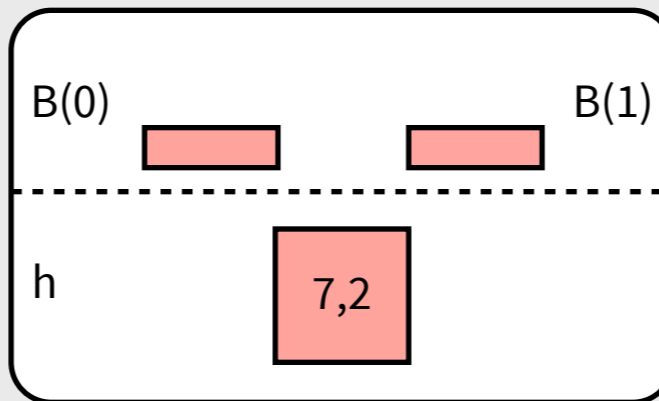


Heap values

In x86-CSL:

The points-to assertion describes a heap value and **empty buffers**.

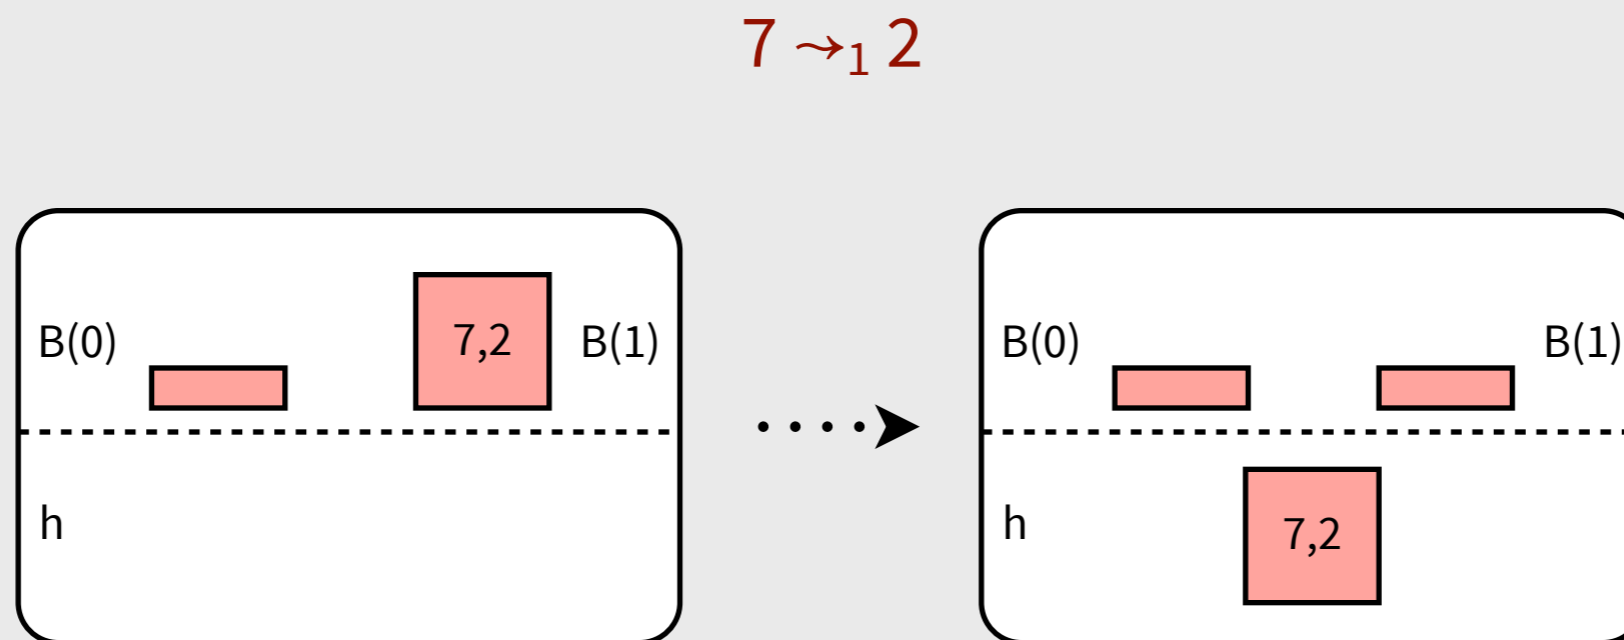
7 ↦ 2



Buffered writes

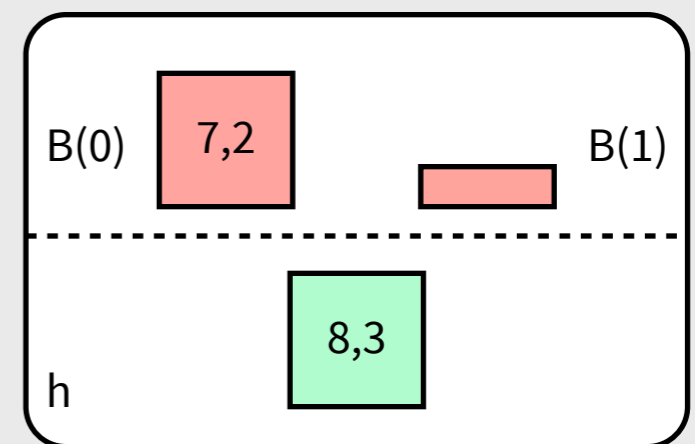
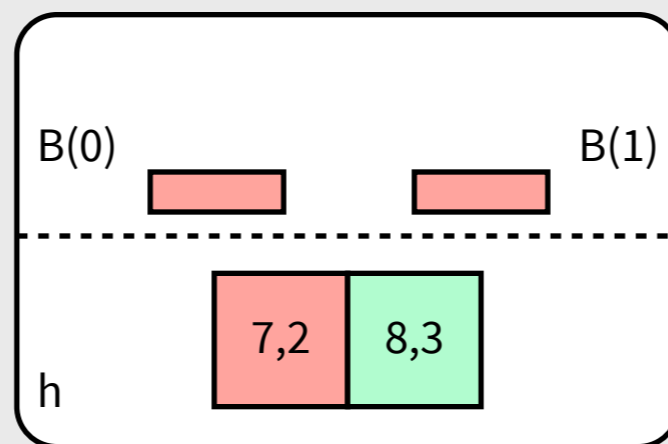
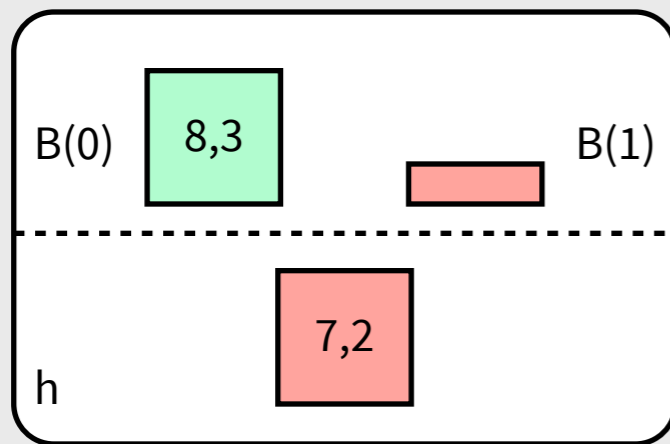
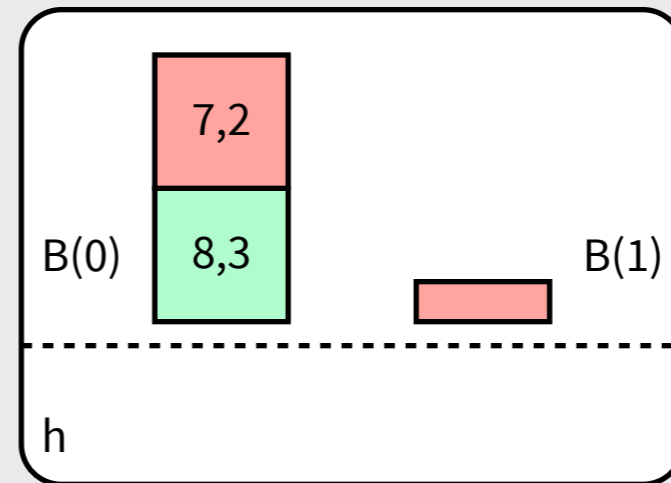
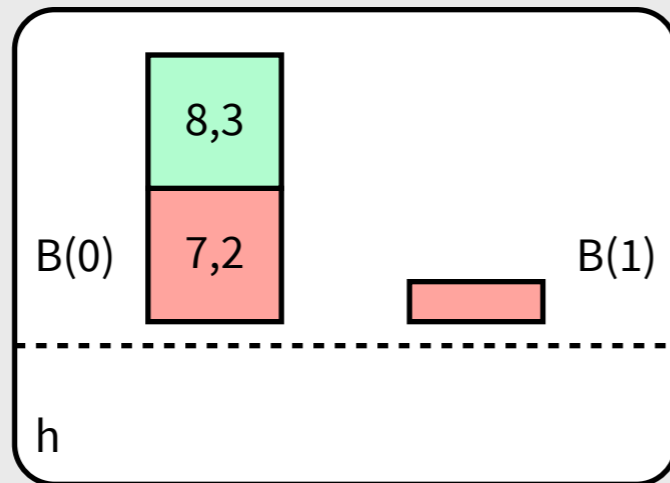
In x86-CSL:

The **leads-to assertion** describes a **single buffered write**.



Buffered writes

7 \rightarrow_0 2 * 8 \rightarrow_0 3



Buffered writes

7 \rightarrow_0 2 * 7 \rightarrow_0 3

Inconsistent!

Spatial separating conjunction can't:

describe writes to the same location;

describe writes in any particular order.

Temporal separation

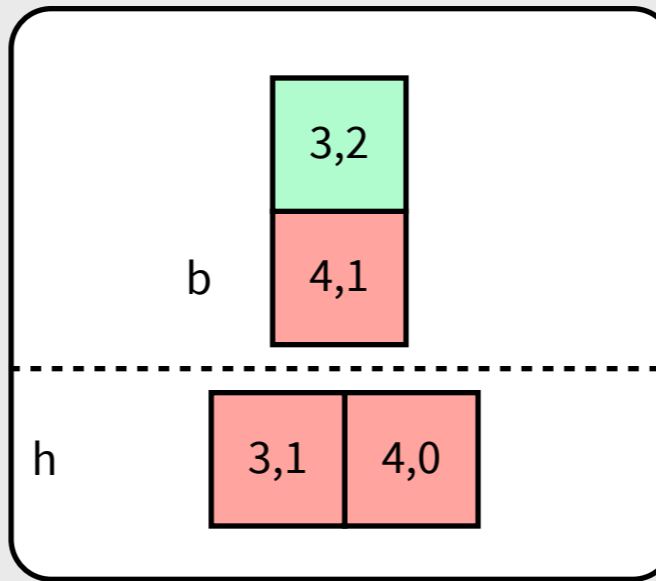
$P \triangleleft Q$: **temporal separating conjunction** of P and Q .

Describes *ordered* sequences of writes to *non-disjoint* addresses.

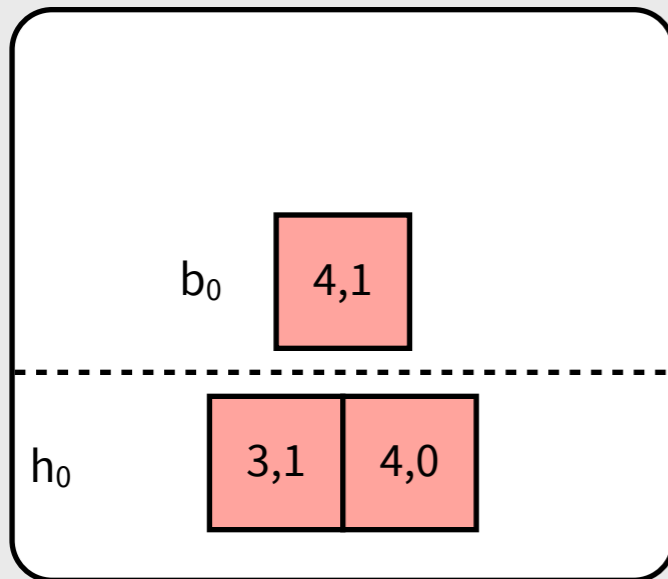
Separates x86 states according to time:

writes described by P must *occur before* writes described by Q .

Temporal separation



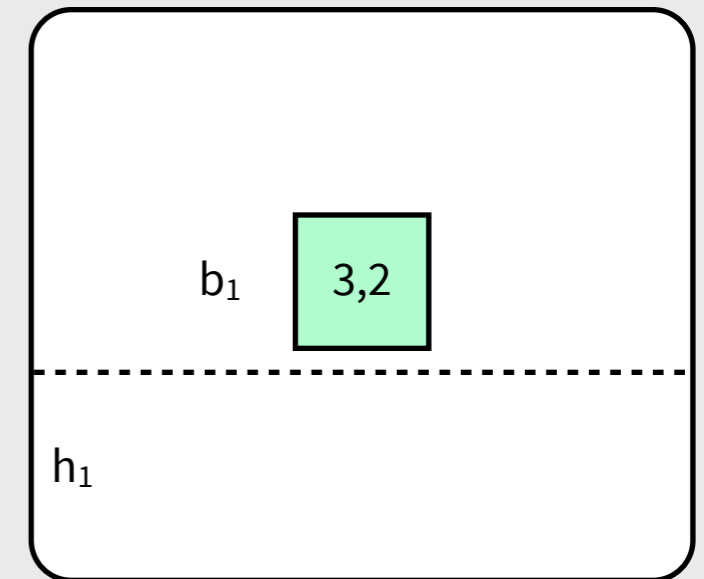
(h, b)



(h_0, b_0)

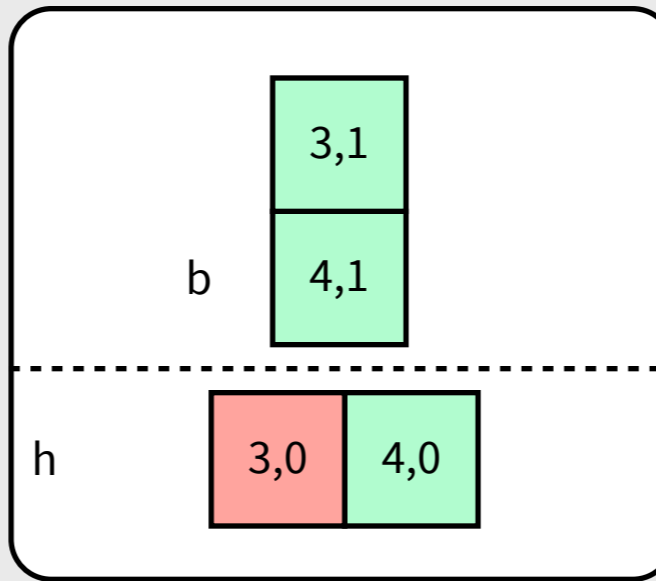


(h_1, b_1)

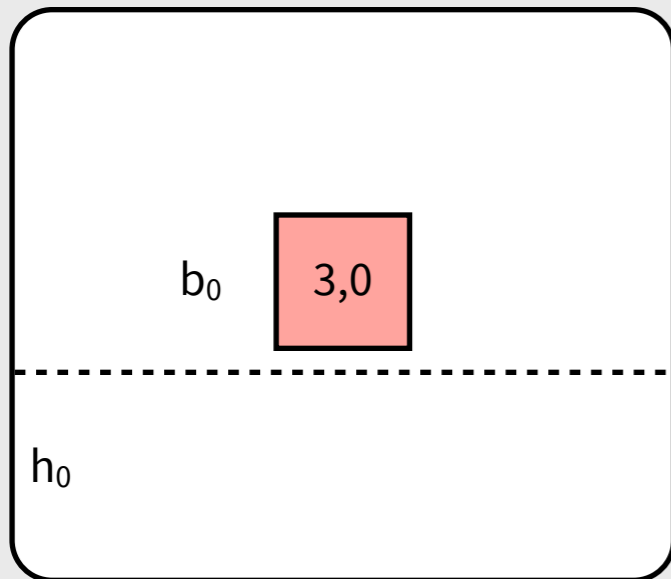


h_1

Temporal separation



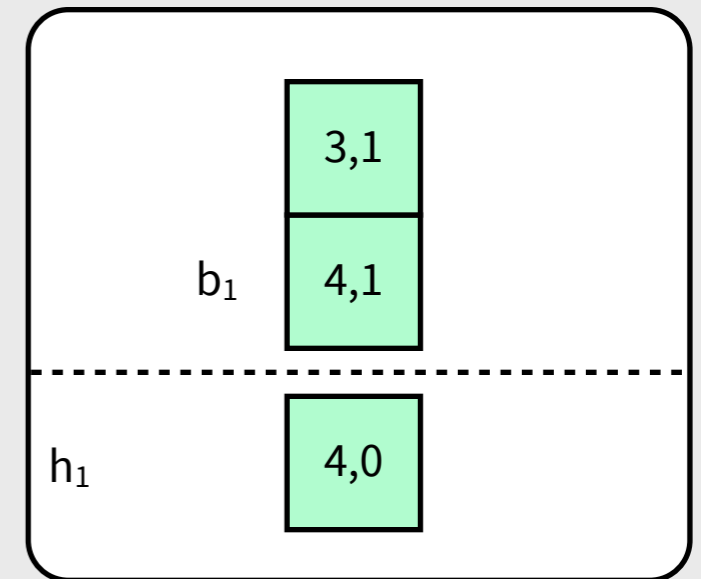
(h, b)



(h_0, b_0)

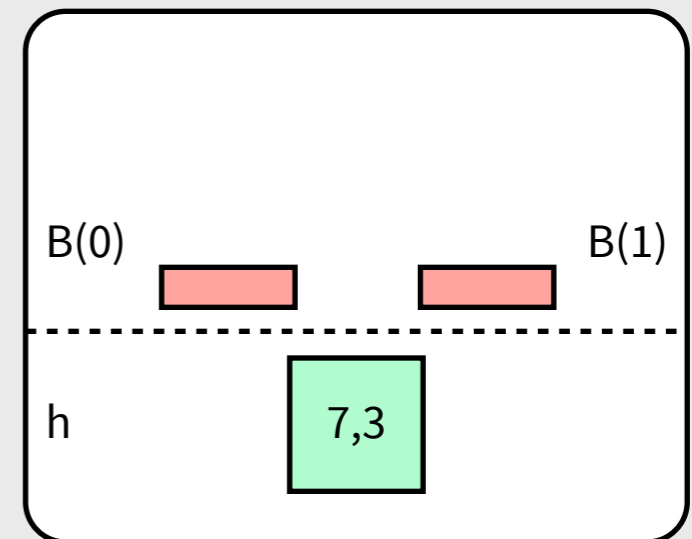
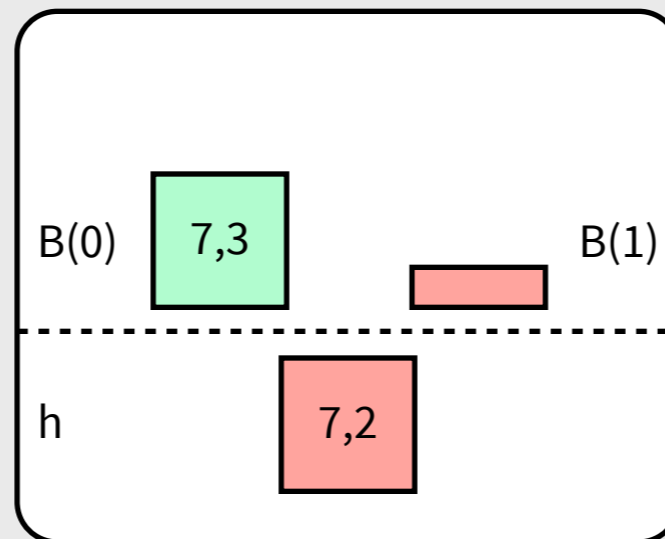
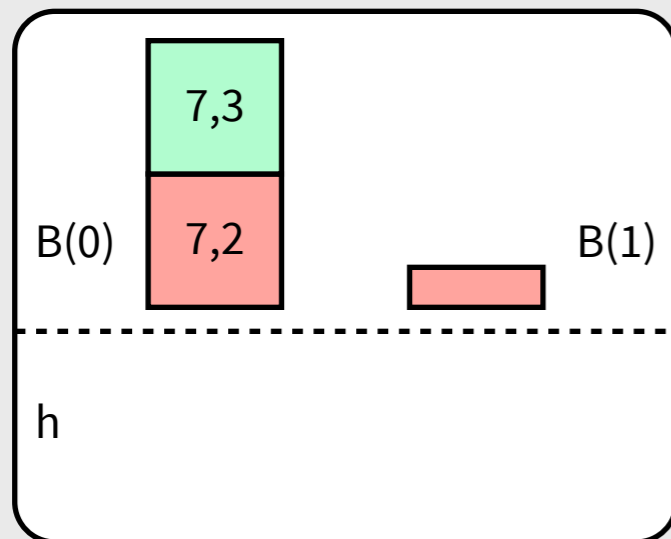


(h_1, b_1)



Write sequences

$7 \rightarrow_0 2 \triangleleft 7 \rightarrow_0 3$



Temporal locality

Commands are local in space **and time**.

Consider a load $x := [7]$:

Assigns to x value of the *most recent* write to address 7.

Earlier writes are irrelevant.

Temporal frame rule:

$$\frac{J \vdash \{ P \} c \{ Q \}}{J \vdash \{ R \triangleleft P \} c \{ R \triangleleft Q \}}$$

Strong temporal separation

$P \blacktriangleleft Q$: strong temporal separating conjunction.

Separates in both time and space;

$$P \blacktriangleleft Q \triangleq (P * Q) \wedge (P \triangleleft Q)$$

Strong temporal frame rule:

$$\frac{J \vdash \{P\} c \{Q\}}{J \vdash \{R \blacktriangleleft P\} c \{R \blacktriangleleft Q\}}$$

Load and store

Load axiom:

Good: $J \vdash \{ e \rightsquigarrow_i f \} x := [e]_i \{ e \rightsquigarrow_i f \wedge x = e \}$

Better: $J \vdash \{ e \rightsquigarrow_i f \blacktriangleleft P \} x := [e]_i \{ (e \rightsquigarrow_i f \blacktriangleleft P) \wedge x = e \}$

Store axiom:

Good: $J \vdash \{ e \rightsquigarrow_i e' \} [e] := f_i \{ e \rightsquigarrow_i e' \blacktriangleleft e \rightsquigarrow_i f \}$

Better: $J \vdash \{ e \rightsquigarrow_i e' \blacktriangleleft P \} [e] := f_i \{ e \rightsquigarrow_i e' \blacktriangleleft P \blacktriangleleft e \rightsquigarrow_i f \}$

Conclusion

Contributions:

A programming language with an x86-like model.

An assertion logic with an x86-like model.

A CSL-style logic for local reasoning about x86-like programs.

(Examples indicate reasoning might not be significantly harder than in CSL.)

Lots of work left!

Some important meta-theory remains (e.g., soundness).

Proof theory of specifications *must* be strengthened.

Thank you

Advisors:

Warren Hunt and J Moore.

Committee, etc.:

Josh Berdine, Allen Emerson, Don Fussell, Tony Hoare and Mohamed Gouda.

Everyone else!

Barrier assertions

emp : empty state

$$P * \mathbf{emp} \equiv P \triangleleft \mathbf{emp} \equiv P \equiv \mathbf{emp} \triangleleft P \equiv \mathbf{emp} * P$$

bar_i : result of flushing i^{th} write buffer

$P \triangleleft \mathbf{bar}_i$: like P but with empty i^{th} buffer

Expresses fence axiom:

$$J \vdash \{ \mathbf{emp} \} \text{ fence}_i \{ \mathbf{bar}_i \}$$

$$J \vdash \{ P \} \text{ fence}_i \{ P \triangleleft \mathbf{bar}_i \}$$

Lock assertions

lock_i describes states in which processor **i** holds lock.

$i \neq j \wedge (\mathbf{lock}_i * \mathbf{lock}_j)$: inconsistent because lock is exclusive.

$i \neq j \wedge (\mathbf{lock}_i * e \rightsquigarrow_j f)$: buffered write *only* because **j** is blocked by **i**.

Lock axioms

Good:

$$J \vdash \{ \mathbf{emp} \} \text{lock}_i \{ \mathbf{lock}_i \}$$
$$J \vdash \{ \mathbf{lock}_i \} \text{unlock}_i \{ \mathbf{emp} \}$$

Better:

$$J \vdash \{ \mathbf{emp} \} \text{lock}_i \{ \mathbf{lock}_i * \mathbf{bar}_i \}$$
$$J \vdash \{ \mathbf{lock}_i \} \text{unlock}_i \{ \mathbf{bar}_i \}$$

Concurrency

Accessing shared state:

$$\frac{\text{emp} \vdash \{ J * P * \text{lock}_i \} \text{c} \{ J * Q * \text{lock}_i \}}{J \vdash \{ P * \text{lock}_i \} \text{c} \{ Q * \text{lock}_i \}}$$

Concurrent composition:

$$\frac{J \vdash \{ P \} \text{c} \{ Q \} \text{ and } J \vdash \{ P' \} \text{c}' \{ Q' \}}{J \vdash \{ P * P' \} \text{c} \parallel \text{c}' \{ Q * Q' \}}$$

Sharing private state:

$$\frac{J \vdash \{ P \} \text{c} \{ Q \}}{\text{emp} \vdash \{ J * P \} \text{c} \{ J * Q \}}$$

Closure

Assertions denote sets that are closed under flushing:

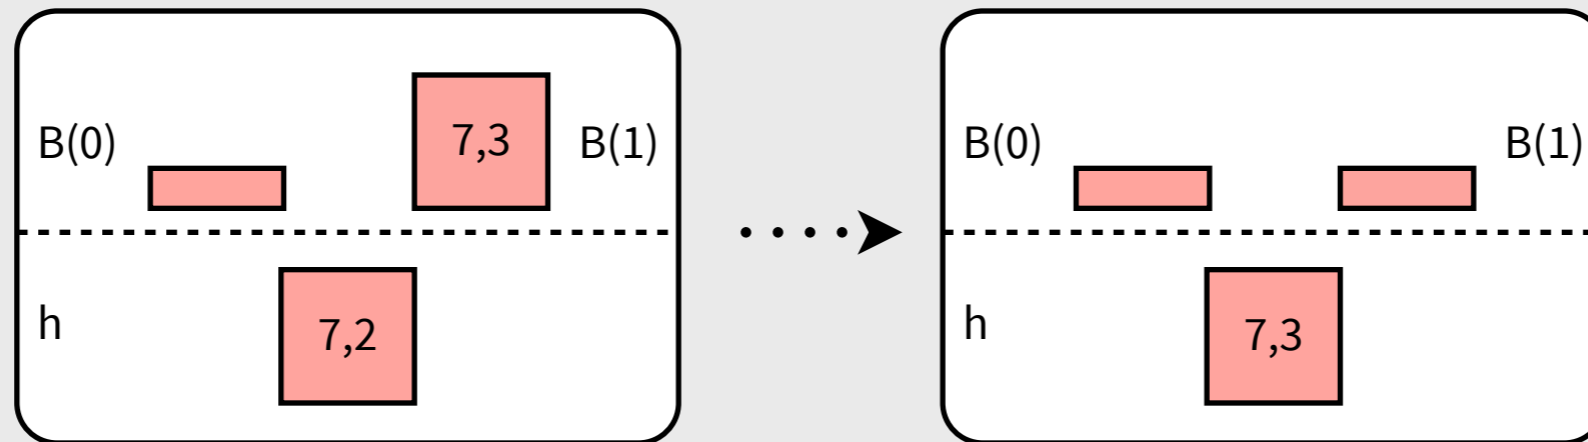
if $\sigma \models P$ and σ can flush writes to yield σ' then also $\sigma' \models P$.

Nondeterministic flushing is hidden by the logic;
no explicit reasoning about flushing.

Important for soundness:

$$J \vdash \{ P \} \text{skip}_i \{ P \}$$

Races and disjunction



$x := [7]_0$ is a racy load; a true post-condition is: $x = 2 \vee x = 3$

Can we use the disjunction rule to reason about racy loads?

$$\frac{J \vdash \{P\} c \{Q\} \text{ and } J \vdash \{P'\} c \{Q\}}{J \vdash \{P \vee P'\} c \{Q\}}$$

No: the former state, alone, is not closed under flushing!