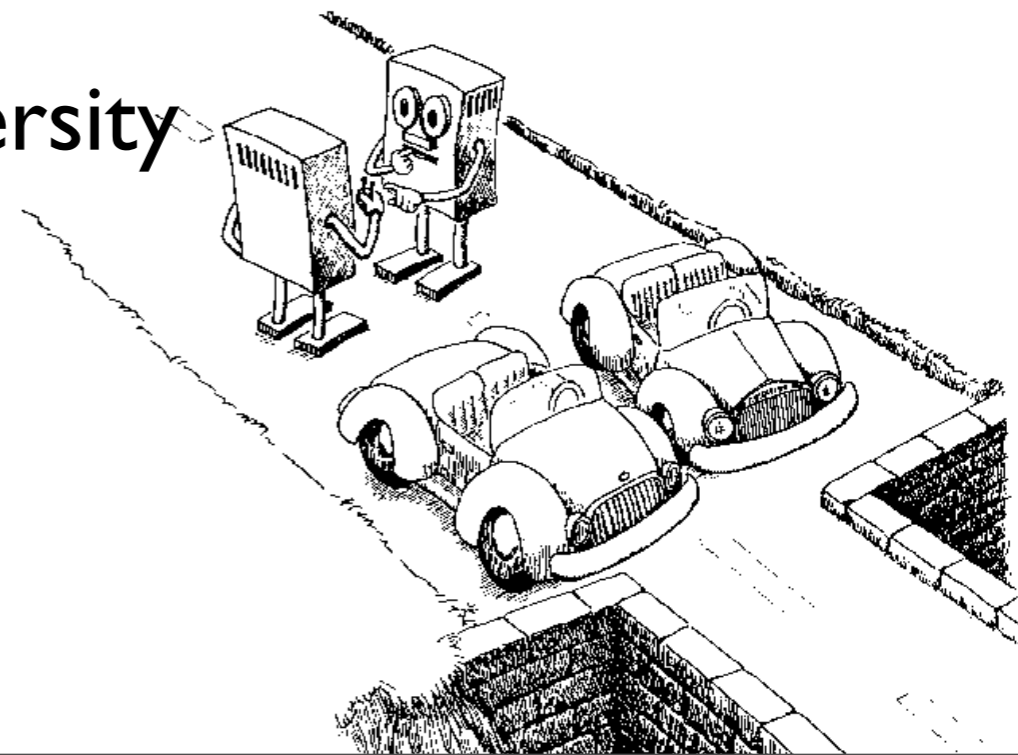


A race-detecting semantics for concurrent programs

Stephen Brookes
Carnegie Mellon University



Concurrency

- processes read and write to shared state
- synchronize via conditional critical regions
- mutually exclusive use of resources

Race condition



- Concurrent write to identifier being read or written by another process

$x := 1 \parallel x := 2$

What's the value of x ?

Depends on granularity...

Solution?



- Ignore races semantically
- Assume known granularity

$$\llbracket x:=1 \parallel x:=2 \rrbracket = \{x:=1 \ x:=2, x:=2 \ x:=1\}$$

High-level

$$x \in \{1, 2\}$$

$$\llbracket x:=1 \parallel x:=2 \rrbracket =$$

$$(x.0:=1 \ x.1:=0) \parallel (x.0:=0 \ x.1:=1)$$

Low-level

$$x \in \{0, 1, 2, 3\}$$

Solution?



- Avoid races syntactically
- Rules for critical variables *Owicki-Gries*
- Semantics just for race-free programs

Solution?



- Avoid races syntactically
- Rules for critical variables *Owicki-Gries*
- Semantics just for race-free programs

$x := 1 \parallel x := 2$ *disqualified*

Solution?



- Avoid races syntactically
- Rules for critical variables *Owicki-Gries*
- Semantics just for race-free programs

$x := 1 \parallel x := 2$ *disqualified*

with r do $x := 1$ \parallel with r do $x := 2$
 $x \in \{1, 2\}$

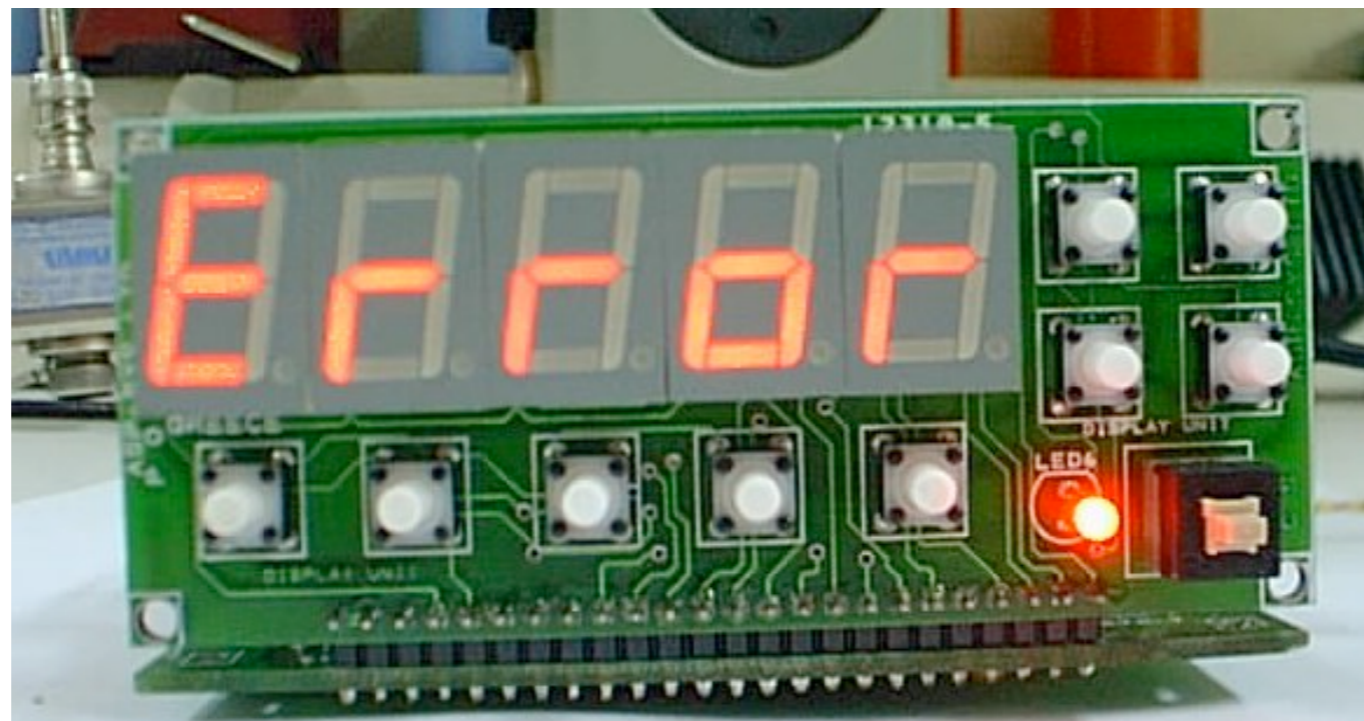
Problems

with the traditional approaches

- Granular semantics
 - combinatorial explosion
 - too specific, not uniform
- Race-avoiding syntax
 - static constraints have limits
 - too draconian

Solution

- Semantics with race-detection
 - potential race treated as catastrophic
cf. Reynolds



Outline

- Trace semantics
 - process denotes set of action traces
- High-level model
 - granularity of integer operations
- Low-level semantics
 - granularity of word operations
- Granularity Theorem
 - high-level consistent with low-level

Actions

High-level model

λ

δ

- idle
- read
- write
- resource
- error

$i=v$

$i:=v$

$try(r), acq(r), rel(r)$

$abort$

$i \in \mathbf{Ide}$ identifiers

$r \in \mathbf{Res}$ resource names

$v \in V$ integers

Traces

- Sequences of actions
- Concatenation

$$\alpha\delta\beta = \alpha\beta$$

$$\alpha \text{ abort } \beta = \alpha \text{ abort}$$

$$\alpha, \beta \in \mathbf{Tr}$$

State

- Global store s
 - maps identifiers to integers
- Resources
 - each process owns a finite set A
 - must be disjoint

changes dynamically...

Effects

- State *enables* certain actions
- Action has an effect

$$(s, A) \xRightarrow{\lambda} (s', A')$$

$$(s, A) \xRightarrow{\lambda} \text{abort}$$

Effects

$$(s, A) \xRightarrow{\delta} (s, A)$$

$$(s, A) \xRightarrow{i=v} (s, A)$$

if $(i, v) \in s$

$$(s, A) \xRightarrow{i:=v} ([s/i:v], A)$$

if $i \in \text{dom}(s)$

Effects

$(s, A) \xRightarrow{i=v} \text{abort}$ if $i \notin \text{dom}(s)$

$(s, A) \xRightarrow{i:=v} \text{abort}$ if $i \notin \text{dom}(s)$

$(s, A) \xRightarrow{\text{abort}} \text{abort}$

$\text{abort} \xRightarrow{\lambda} \text{abort}$

Effects

$$(s, A) \xRightarrow{\text{try}(r)} (s, A)$$

$$(s, A) \xRightarrow{\text{acq}(r)} (s, A \cup \{r\}) \quad \text{if } r \notin A$$

$$(s, A) \xRightarrow{\text{rel}(r)} (s, A - \{r\}) \quad \text{if } r \in A$$

Semantics of expressions

$$\llbracket e \rrbracket \subseteq \mathbf{Tr} \times V$$

$$\llbracket i \rrbracket = \{(i=v, v) \mid v \in V\}$$

$$\llbracket e+e' \rrbracket = \{(\rho\rho', v+v') \mid \\ (\rho, v) \in \llbracket e \rrbracket \ \& \ (\rho', v') \in \llbracket e' \rrbracket\}$$

Semantics

of commands

$$\llbracket c \rrbracket \subseteq \mathbf{Tr}$$

$$\llbracket i := e \rrbracket = \{ \rho \ i := v \mid (\rho, v) \in \llbracket e \rrbracket \}$$

$$\llbracket c ; c' \rrbracket = \llbracket c \rrbracket \llbracket c' \rrbracket$$

$$\llbracket c \parallel c' \rrbracket = \llbracket c \rrbracket \overset{\text{mutex}}{\parallel} \llbracket c' \rrbracket$$

*mutex fairmerge
with race-detection*

Mutual exclusion

- At most one process holds each

$(acq(r) x:=1 rel(r)) \parallel (acq(r) x:=2 rel(r))$

only includes

$acq(r) x:=1 rel(r) acq(r) x:=2 rel(r)$

and

$acq(r) x:=2 rel(r) acq(r) x:=1 rel(r)$

Mutex fairmerge

Traditional Definition

$$\begin{aligned} & (\lambda_1 \alpha)_{A_1} \parallel_{A_2} (\lambda_2 \beta) \\ = & \{ \lambda_1 \gamma \mid (A_1, A_2) \xrightarrow{\lambda_1} (A'_1, A_2) \ \& \ \gamma \in \alpha_{A'_1} \parallel_{A_2} (\lambda_2 \beta) \} \\ & \cup \{ \lambda_2 \gamma \mid (A_2, A_1) \xrightarrow{\lambda_2} (A'_2, A_1) \ \& \ \gamma \in (\lambda_1 \alpha)_{A_1} \parallel_{A'_2} \beta \} \end{aligned}$$

*each process constrained by the other
to maintain disjoint sets of resources*

Race-detecting mutex fairmerge

Definition

$$(\lambda_1 \alpha)_{A_1} \parallel_{A_2} (\lambda_2 \beta) = \{abort\}$$

if λ_1 and λ_2 interfere

$$= \{ \lambda_1 \gamma \mid (A_1, A_2) \xrightarrow{\lambda_1} (A'_1, A_2) \ \& \ \gamma \in \alpha_{A'_1} \parallel_{A_2} (\lambda_2 \beta) \}$$
$$\cup \{ \lambda_2 \gamma \mid (A_2, A_1) \xrightarrow{\lambda_2} (A'_2, A_1) \ \& \ \gamma \in (\lambda_1 \alpha)_{A_1} \parallel_{A'_2} \beta \}$$

otherwise

Interference

Definition

$$\lambda_1 \bowtie \lambda_2$$

if

$$free(\lambda_1) \cap writes(\lambda_2) \neq \{\}$$

or

$$free(\lambda_2) \cap writes(\lambda_1) \neq \{\}$$

*concurrent write to identifier
being used by other process*

Semantics

$$\llbracket c \parallel c' \rrbracket = \llbracket c \rrbracket \parallel_{\emptyset} \llbracket c' \rrbracket$$

*mutex fairmerge
with
race detection*

Example

$\llbracket x := x+1 \parallel x := x+1 \rrbracket$

$= \{x=v \text{ abort} \mid v \in V\}$

Semantics

[[with r when b do c]] =

wait enter* \cup *wait* ^{ω}

wait = *acq*(r) **[[b]]**_{false} *rel*(r) \cup {*try*(r)}

enter = *acq*(r) **[[b]]**_{true} **[[c]]** *rel*(r)

critical region protected by resource

Semantics

$$\llbracket \text{resource } r \text{ in } c \rrbracket = \{ \alpha \backslash r \mid \alpha \in \llbracket c \rrbracket_r \}$$

$\llbracket c \rrbracket_r$ *traces sequential for r*

$\alpha \backslash r$ *hide actions on r*

statically scoped local resource

Examples

[[with r do $x := x+1$]] =

$try(r)^\infty \{acq(r) x=v x:=v+1 rel(r) \mid v \in V \}$

[[resource r in (...) || (...)]] =

[[$x := x+1$; $x := x+1$]]

Respect for resources

Lemma

If $\alpha \in \llbracket c \rrbracket$ and

$$(s, \emptyset) \xRightarrow{\alpha} (s', A')$$

then $A' = \emptyset$

Race-free programs

Definition

c is race-free from s

iff

$\neg \exists \alpha \in \llbracket c \rrbracket. s \xRightarrow{\alpha} \text{abort}$

Low-level model

- Word size M
- Integer represented as *list of words*
- Word-level actions

$i.len=n$

$i.len:=n$

$i.0=w, \dots, i.n=w$

$i.0:=w, \dots, i.n:=w$

$$0 \leq w < 2^M$$

Low-level states

- Global store S
 - maps identifiers to *lists of words*
- Each process has set of resources A
 - pairwise disjoint

Low-level semantics of expressions

$\llbracket i \rrbracket =$

$$\{(i.len=n \ i.0=w_0 \ \dots \ i.n=w_n, [w_0, \dots, w_n]) \mid \\ n \geq 0 \ \& \ w_0, \dots, w_n \in W\}$$

$\llbracket e+e' \rrbracket = \{(\rho\rho', L \oplus L') \mid$

$$(\rho, L) \in \llbracket e \rrbracket \ \& \ (\rho', L') \in \llbracket e' \rrbracket\}$$

word-level arithmetic

Low-level semantics of commands

$$\llbracket i := e \rrbracket = \left\{ \rho \mid i.\text{len} := n \quad i.0 := w_0 \dots i.n := w_n \right. \\ \left. \mid (\rho, [w_0, \dots, w_n]) \in \llbracket e \rrbracket \right\}$$

Interference

at low level

$$\textit{writes}(i.j:=v) = \{i.j\}$$

$$\textit{writes}(i.len:=n) = \{\}$$

$$\textit{writes}(i.j=v) = \{\}$$

$$\textit{reads}(i.j:=v) = \{\}$$

$$\textit{reads}(i.len:=n) = \{\}$$

$$\textit{reads}(i.j=v) = \{i.j\}$$

Representation

Definitions

- Word lists represent integers

$$[w_0, w_1, \dots, w_n]_M = w_0 + 2^M w_1 + \dots + 2^{nM} w_n$$

- Low-level states represent high-level states

$$s_{low} \approx s_{high}$$

iff

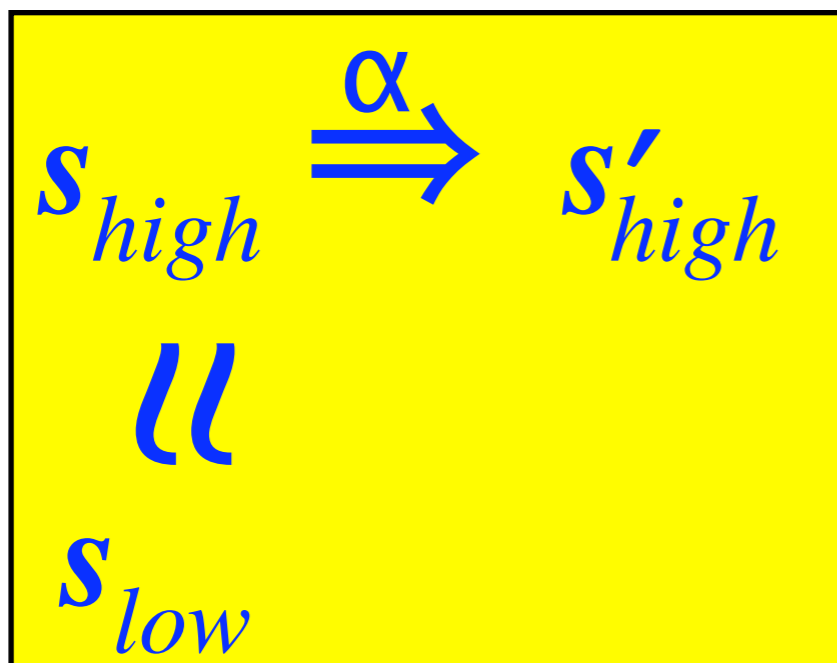
$$[s_{low}^{(i)}]_M = s_{high}^{(i)}$$

for all identifiers i

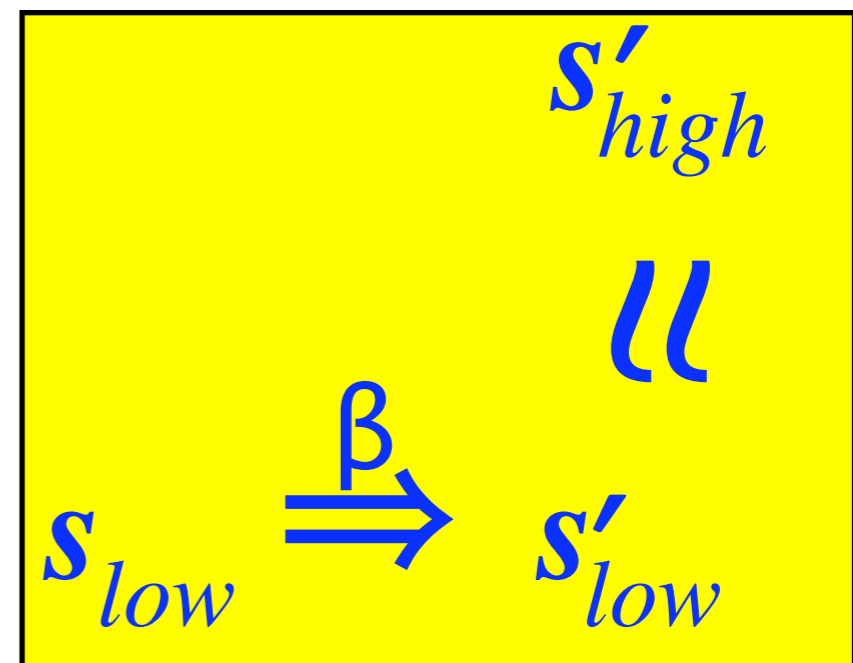
Granularity Theorem (I)

- Every high-level error-free computation simulates a low-level computation

For all $\alpha \in \llbracket c \rrbracket_{high}$, s_{high} , s'_{high} , s_{low}



IMPLIES

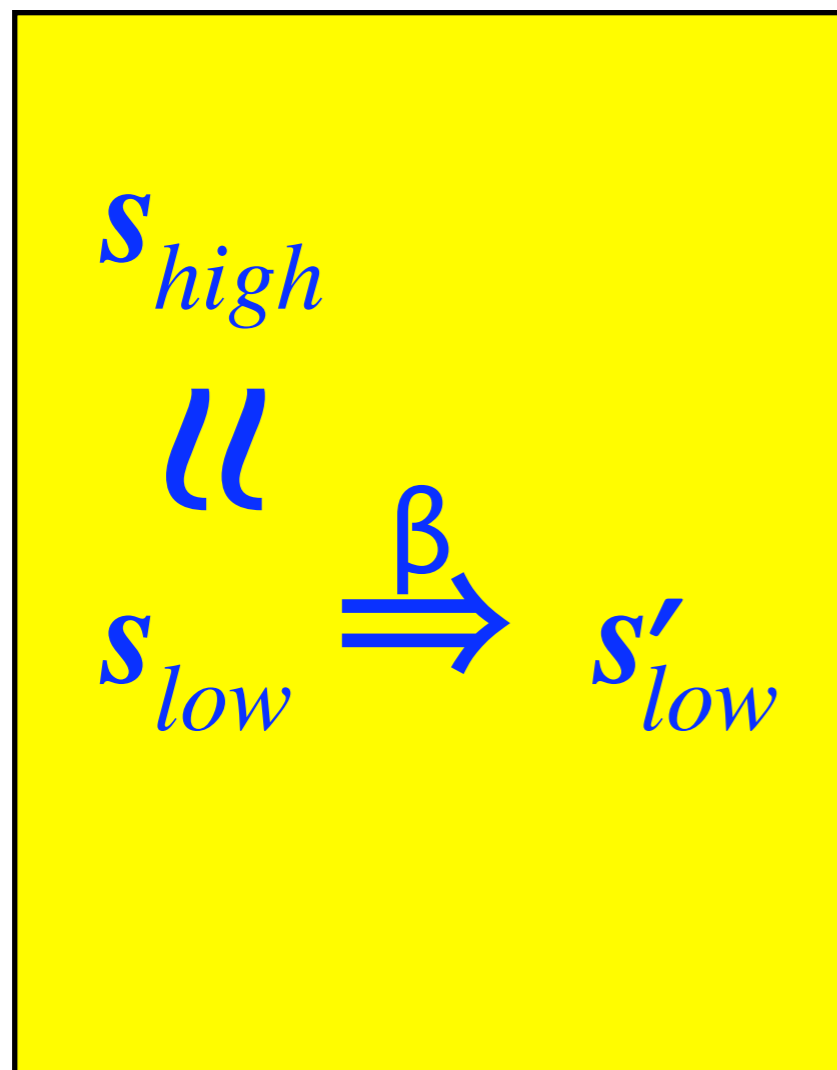


for some $\beta \in \llbracket c \rrbracket_{low}$, s'_{low}

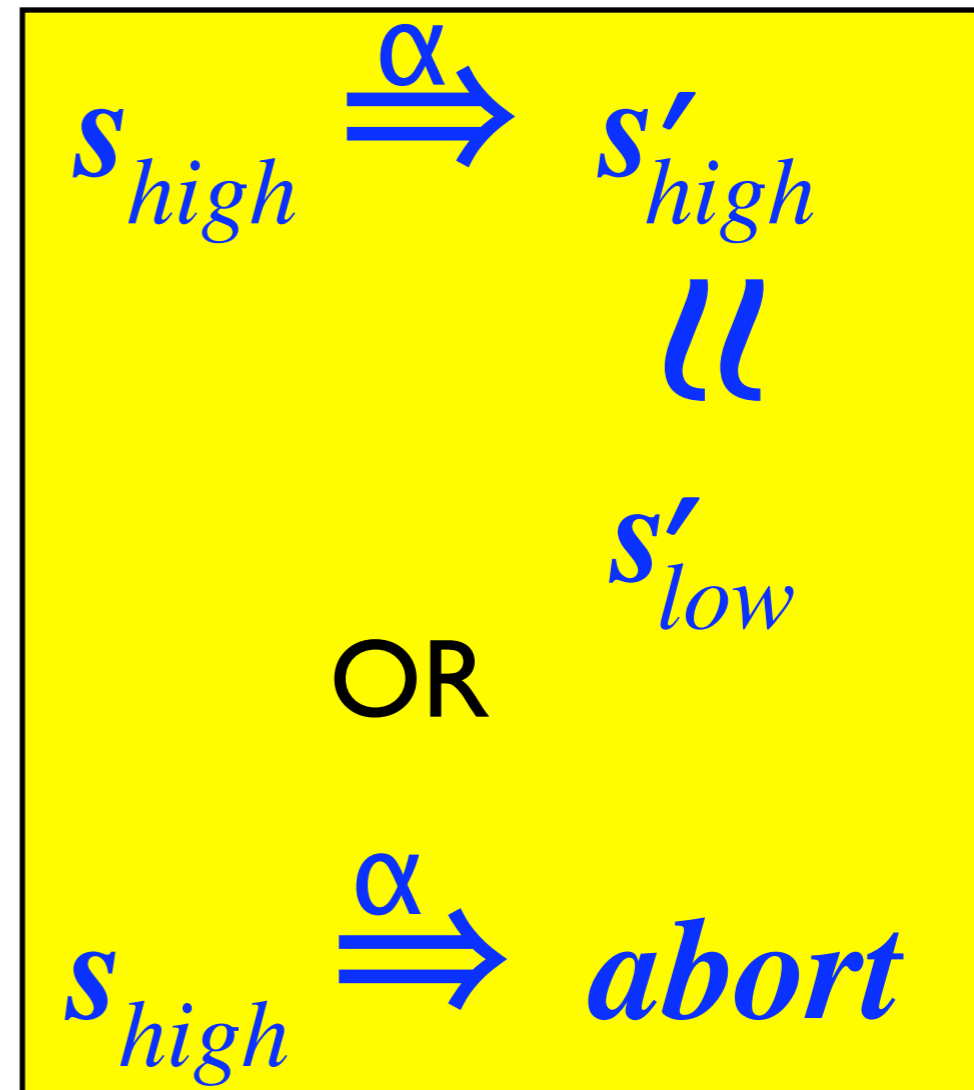
Granularity Theorem⁽²⁾

- Every low-level computation is (weakly) simulated by a high-level computation

For all $\beta \in \llbracket c \rrbracket_{low}$, s_{high} , s'_{low} , s_{low}



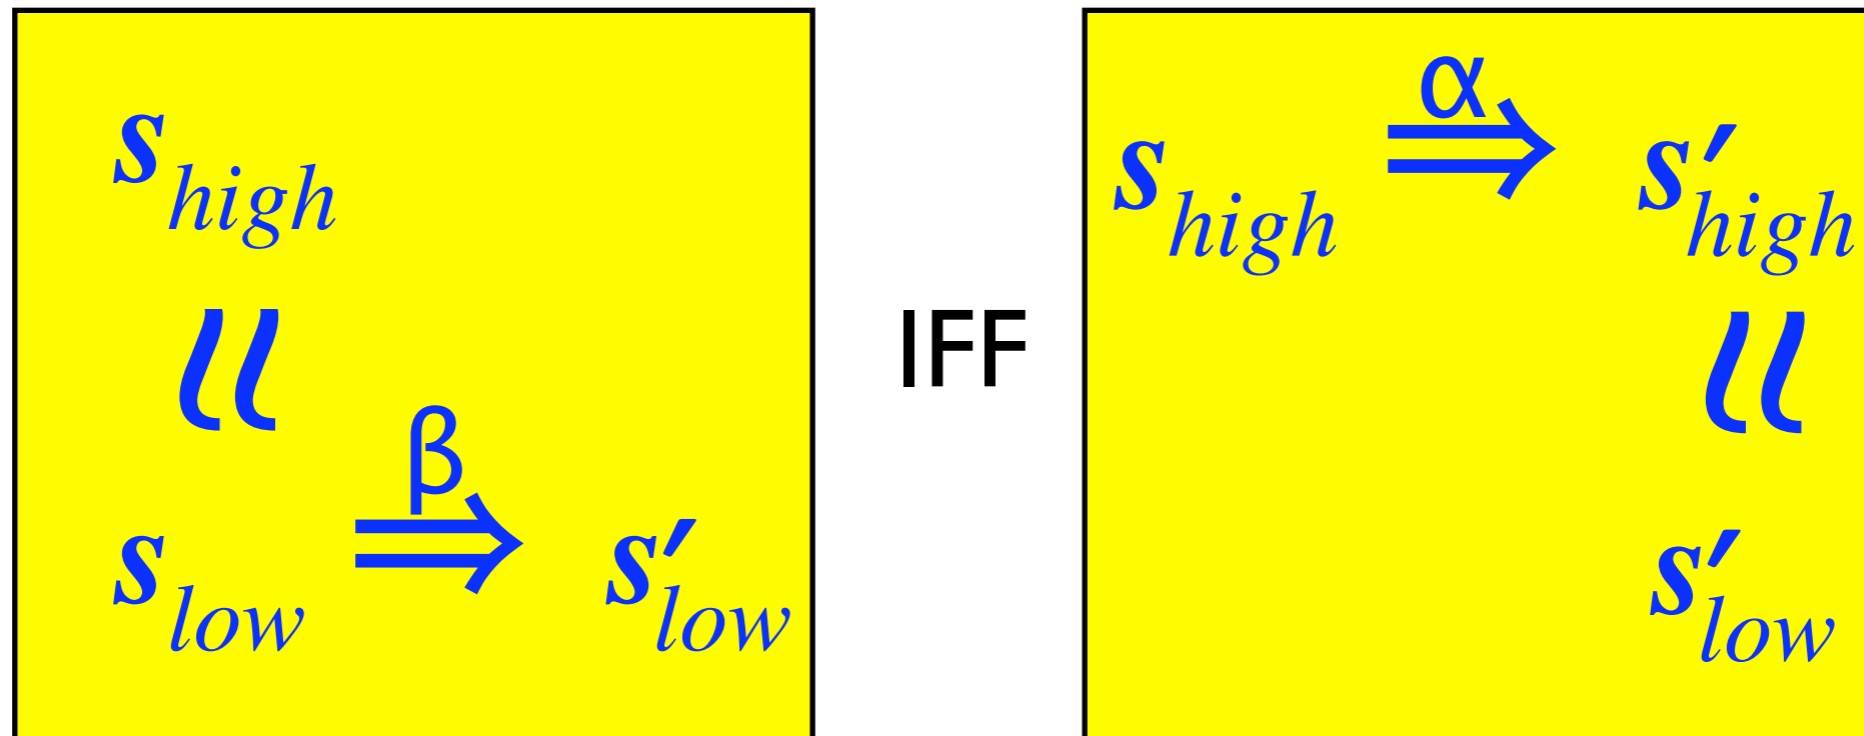
IMPLIES



for some $\alpha \in \llbracket c \rrbracket_{high}$, s'_{high}

Race-free case

- Simulation both ways if c is race-free



More concrete

- Low-level state = store + heap
- Effect of *$i.len := n$* includes *allocation* and/or *deallocation* of heap cells
- Results still go through

Further work

- Soundness of Owicki-Gries logic
- Ideas extend to include pointers
- concurrent separation logic

O'Hearn, Brookes

to appear at

CONCUR '04

