

# Distributed race detection

Igor Walukiewicz



## Race prediction

• Operations  $r_t(x), w_t(x)$   
 $acq_t(l), rel_t(l)$

• Lock discipline:  
critical sections of the  
same lock are disjoint

	$t_1$	$t_2$
1	acq( $l$ )	
2	$w(y)$	
3	$w(x)$	
4		$r(x)$
5	rel( $l$ )	
6		acq( $l$ )
7		$w(y)$
8		acq( $l$ )

## Race prediction

• Operations  $r_t(x), w_t(x)$   
 $acq_t(l), rel_t(l)$

• Lock discipline:  
critical sections of the  
same lock are disjoint

• Race  $(w_t(x), r_{t'}(x))$  or  $(w_t(x), w_{t'}(x))$   
the two events are concurrent in  $G$   
(can appear next to each other  
in some sound reordering of  $G$ )

	$t_1$	$t_2$
1	acq( $l$ )	
2	w( $y$ )	
3	w( $x$ )	
4		r( $x$ )
5	rel( $l$ )	
6		acq( $l$ )
7		w( $y$ )
8		acq( $l$ )

## Race prediction

• Operations  $r_t(x), w_t(x)$   
 $acq_t(l), rel_t(l)$

• Lock discipline:  
critical sections of the  
same lock are disjoint

• Race  $(w_t(x), r_{t'}(x))$  or  $(w_t(x), w_{t'}(x))$   
the two events are concurrent in  $\mathcal{G}$   
(can appear next to each other)  
(in some sound reordering of  $\mathcal{G}$ )

Race prediction problem:  
looking at  $\mathcal{G}$  detect if  
there is a race

	$t_1$	$t_2$
1	acq(l)	
2	w(y)	
3	w(x)	
4		r(x)
5	rel(l)	
6		acq(l)
7		w(y)
8		acq(l)

## Race prediction

- Operations  $r_t(x)$ ,  $w_t(x)$   
 $acq_t(\ell)$ ,  $rel_t(\ell)$
- Lock discipline:  
critical sections of the  
same lock are disjoint
- Race  $(w_t(x), r_{t'}(x))$  or  $(w_t(x), w_{t'}(x))$   
the two events are concurrent in  $\sigma$   
(can appear next to each other)  
(in some sound reordering of  $\sigma$ )

Race prediction problem:  
looking at  $\sigma$  detect if  
there is a race

	$t_1$	$t_2$
1	acq( $\ell$ )	
2	$r(x)$	
3		$w(z)$
4	rel( $\ell$ )	
5		$w(x)$

Execution  $\sigma$   
No data race

	$t_1$	$t_2$
1	acq( $\ell$ )	
2		$w(z)$
3	$r(x)$	
4		$w(x)$
5	rel( $\ell$ )	

Execution  $\sigma'$

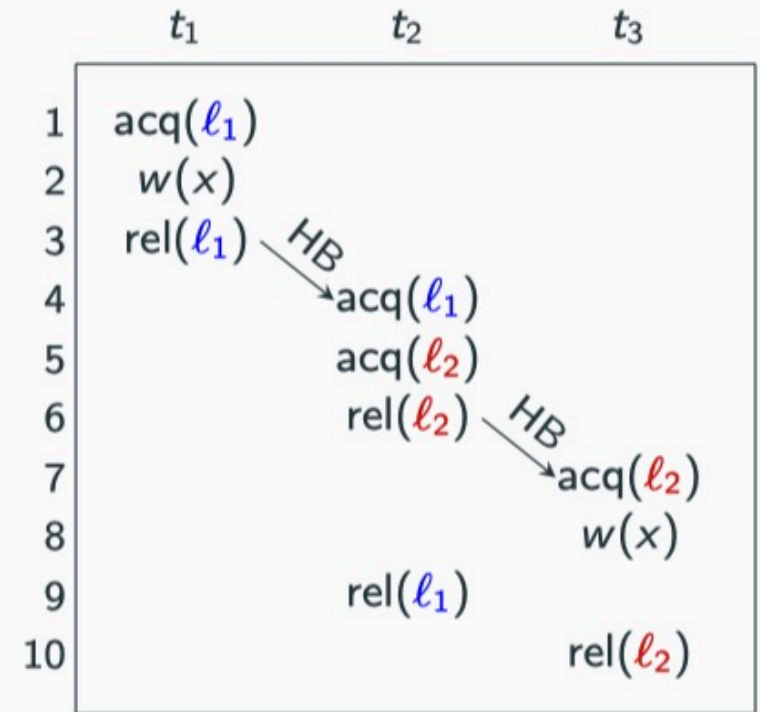
# Happens-before order [Lamport]

What is a sound reordering of  $\tau$ ?

- operations on different threads permute unless they concern the same lock

$e_1 \stackrel{hb}{\leq} e_2$  is the smallest partial-order s.t.

- events of the same thread are ordered
- if  $\dots \underset{t}{rel}(l) \dots \underset{t'}{acq}(l) \dots$  then  $\stackrel{hb}{\leq}$



# Happens-before order [Lamport]

What is a sound reordering of  $\sigma$ ?

- operations on different threads permute unless they concern the same lock

$e_1 \leq_{HB}^v e_2$  is the smallest partial-order s.t.

- events of the same thread are ordered
- if  $\dots \underset{t}{rel}(l) \dots \underset{t'}{acq}(l) \dots$  then  $\leq_{HB}^v$

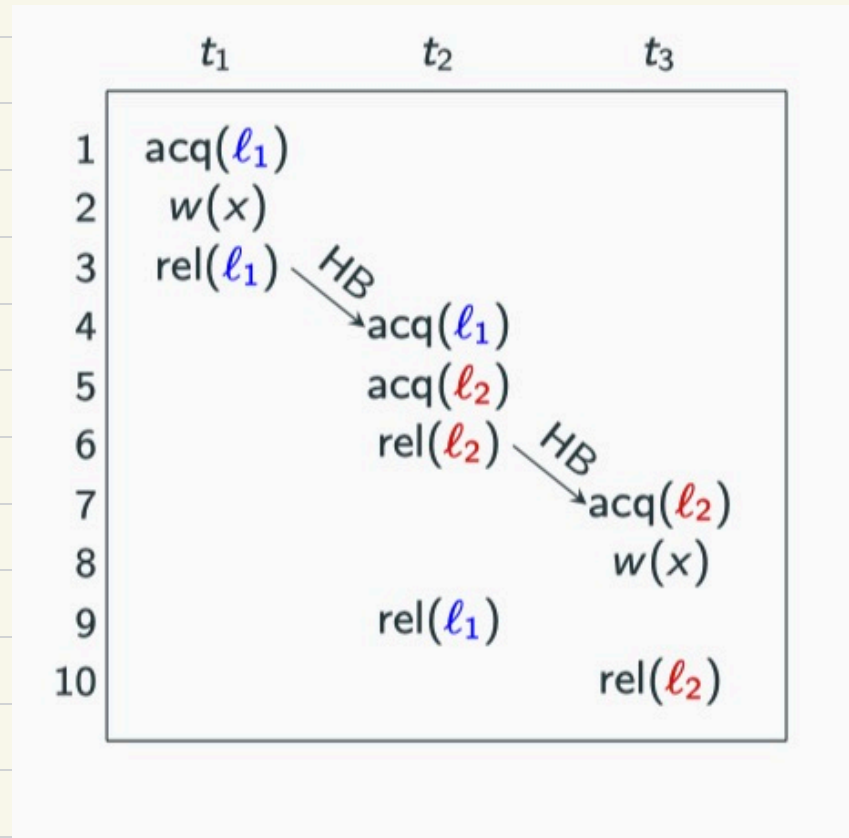
$\sigma \approx_{HB} \sigma'$  if the two have the same set of events and give the same HB order

Race detection: race prediction for  $\approx_{HB}$

By looking at  $\sigma$  detect if there is

$\sigma' \approx_{HB} \sigma$  with a race situation

$w_t(x) \ r_{t'}(x)$





# Vector clocks [Lamport]


Clocks measuring HB-time

$$C_t: J \rightarrow \mathbb{N}$$

$C_t(t)$  - program counter of  $t$

$C_t(t')$  - the last program counter of  $t'$  seen by  $t$

$$\Gamma_t(x), \omega_t(x)$$



# Vector clocks [Lamport]


Clocks measuring HB-time

$$C_t: J \rightarrow \mathbb{N}$$

$C_t(t)$  - program counter of  $t$


$C_t(t')$  - the last program counter of  $t'$  seen by  $t$

$$\Gamma_t(x), \omega_t(x)$$



$$C_t(t) := C_t(t) + 1$$

$$acq_t(e)$$



$$acq_t(e)$$

$$C_t(t) := C_e(t) + 1$$

$$\forall t' \neq t \quad C_t(t') = \max(C_t(t'), C_e(t'))$$

$$\forall t' \quad C_e(t') = \text{---} \text{---} \text{---}$$

# Vector clocks [Lamport]


Clocks measuring HB-time

$$C_t: J \rightarrow \mathbb{N}$$

$C_t(t)$  - program counter of  $t$


$C_t(t')$  - the last program counter of  $t'$  seen by  $t$

$\Gamma_t(x), \omega_t(x)$



$$C_t(t) := C_x(t) + 1$$

$acq_t(e)$




$$C_t(t) := C_e(t) + 1$$

$acq_t(e)$

$$\forall t' \neq t \quad C_t(t') = \max(C_t(t'), C_e(t'))$$

$$\forall t' \quad C_e(t') = \text{---} \text{---} \text{---}$$

$rel_t(e)$



$$C_t(t) := C_e(t) + 1$$

$$\forall t': C_e(t') = C_e(t')$$

## Race detection

$$C_t: T \rightarrow \mathbb{N}$$

$C_t(t)$  - program counter of  $t$

$C_t(t')$  - the last program counter of  $t'$  seen by  $t$

$$W^x \in T \times \mathbb{N}$$

thread that made the last write to  $x$ , and its PC

$$R^x: T \rightarrow \mathbb{N}$$

for every thread last read from  $x$

## Race check

$r_t(x)$  if  $W^x = (t', i)$  and  $C_t(t') < i$  then race else  $R^x(t) := C_t(t)$

$w_t(x)$

if  $\xrightarrow{\hspace{10em}} \parallel \xrightarrow{\hspace{10em}}$

or  $\exists t'. C_t(t') < R^x(t')$  then race else  $W^x = (t, C(t))$

## Race detection

Lamport. 1978. Time, Clocks, and the Ordering of Events in a Distributed System. Commun. ACM 21

Itzkovitz, Schuster, Zeev-Ben-Mordehai. 1999. Toward Integration of Data Race Detection in DSM Systems. J. Parallel Distrib. Comput.

Flanagan, Freund. 2009. FastTrack: Efficient and Precise Dynamic Race Detection. PLDI '09

Thokair, Zhang, Mathur, Viswanathan, Dynamic Race Detection with  $O(1)$  Samples, POPL'23

# Race detection (distributed)

$$\mathcal{C}_t: \mathcal{T} \rightarrow \mathbb{N}$$

$\mathcal{C}_t(t)$  - program counter of  $t$

$\mathcal{C}_t(t')$  - the last program counter of  $t'$  seen by  $t$


$$W_t^x \in \mathcal{T} \times \mathbb{N}$$

thread that made the last write to  $x$ , and its PC

$$R_t^x: \mathcal{T} \rightarrow \mathbb{N}$$

for every thread last read from  $x$

## Race check



$$\exists x \ W_t(x) = (t', i) \quad \mathcal{C}_e(t') < i$$

$$W_e(x) = (t'', j) \quad \mathcal{C}_t(t'') < j$$

# Race detection (distributed)

$$\mathcal{C}_t: \mathcal{T} \rightarrow \mathbb{N}$$

$\mathcal{C}_t(t)$  - program counter of  $t$

$\mathcal{C}_t(t')$  - the last program counter of  $t'$  seen by  $t$


$$W_t^x \in \mathcal{T} \times \mathbb{N}$$

thread that made the last write to  $x$ , and its PC

$$R_t^x: \mathcal{T} \rightarrow \mathbb{N}$$


for every thread last read from  $x$

## Race check



$$\exists x \ W_e(x) = (t', i) \quad \mathcal{C}_e(t') < i$$

$$W_e(x) = (t'', j) \quad \mathcal{C}_e(t'') < j$$



$$\exists x \ W_e(x) = (t', i) \quad \mathcal{C}_e(t') < i$$

$$\exists t'' \ R(x, t'') > \mathcal{C}_e(t')$$

## Race detection (distributed)

$$\tau_t: \mathcal{T} \rightarrow \mathbb{N}$$

$\tau_t(t)$  - program counter of  $t$


$\tau_t(t')$  - the last program counter of  $t'$  seen by  $t$

$$W_t^x \in \mathcal{T} \times \mathbb{N}$$

thread that made the last write to  $x$ , and its PC

$R_t^x: \mathcal{T} \rightarrow \mathbb{N} \cup \{\perp\}$  for every thread last read from  $x$

Update is difficult



for  $x \in \text{Var}$

$$W_e^x := (t', i)$$

if  $\tau_e(t') < i$  then  $W_e^x := (t', i)$

for  $t' \in \mathcal{T}$


if  $R_t^x(t') > \tau_e(t')$  then  $R_e^x(t') = R_t^x(t')$



## Distributed race detection, take 2

$$W_t : J \rightarrow (W \times \text{Vars})^*$$

$$R_t : J \rightarrow (W \times \text{Vars} \times \text{Vals})^*$$



Race if there is  $x \in \text{Vars}$   $t_1, t_2 \in J$  such that

- $(i, x) \in W_t(t_1)$
- $e_e(t_1) < i$
- $(j, x, d) \in R_e(t_2)$
- $e_t(t_2) < j$


# Distributed race detection, take 2

$$W_t : J \rightarrow (W \times \text{Vars})^*$$



$W_t(t')$  sorted by  $W$

$$R_t : J \rightarrow (W \times \text{Vars} \times \text{Vals})^*$$




Race if there is  $x \in \text{Vars}$   $t_1, t_2 \in J$  such that

- $(i, x) \in W_t(t_1)$
- $e_e(t_1) < i$
- $(j, x, \alpha) \in R_e(t_2)$
- $e_t(t_2) < j$

$t_1$  moved in  $t \downarrow - e \downarrow$  if  $e_e(t_1) > \tau_e(t_1) \mapsto t_1 \in MI_t$   
 $x$  written in  $t \downarrow - e \downarrow$  if  $(i, x) \in W_t(t_1)$   $i > \tau_e(t_1) \mapsto WX_t$

for  $t_2 \in MI_e$


for  $(i, x, \alpha) \in R_e(t_2)$  with  $i > \tau_t(t_2)$   
 if  $x \in WX_t$  then race



# Distributed race detection, take 2

$$W_t : J \rightarrow (W \times \text{Vars})^*$$

$$R_t : J \rightarrow (W \times \text{Vars} \times \text{Vals})^*$$



Race if there is  $x \in \text{Vars}$   $t_1, t_2 \in J$  such that


- $(i, x) \in W_t(t_1)$
- $e_e(t_1) < i$
- $(j, x, \alpha) \in R_e(t_2)$
- $e_t(t_2) < j$

$t_1$  moved in  $t \downarrow - e \downarrow$  if  $e_e(t_1) > \tau_e(t_1) \mapsto t_1 \in MI_t$   
 $x$  written in  $t \downarrow - e \downarrow$  if  $(i, x) \in W_t(t_1)$   $i > \tau_e(t_1) \mapsto Wx_t$

for  $t_2 \in MI_e$

for  $(i, x, \alpha) \in R_e(t_2)$  with  $i > \tau_x(t_2)$


if  $x \in Wx_t$  then race



The work is proportional to the size of symmetric difference.

# TODO


- Faster distributed monitoring
- Distributed monitoring for other properties
  - serializability



- any  $\Sigma_2$  property

# TODO

- Faster distributed monitoring
- Distributed monitoring for other properties
  - serializability



• any  $\Sigma_2$  property

- Race prediction reads from equivalence

