

A Deadlock-Free Semantics
for
Shared Memory Concurrency

MOTIVATION

Shared memory concurrency:

- ▶ traditional algorithmic programming style
- ▶ not far from the machine architecture
- ↳ needed at some implementation stage.

Problem: not modular (data-races, deadlocks...)

- ↳ keep the model, restrict the semantics, avoiding deadlocks.

AN EXAMPLE

(1/3)

Bank accounts: a module where an **account** is a reference (pointer) to an integer, with some functions:

► to **deposit** some amount x on account y : $\lambda x \lambda y (y := !y + x)$.

Incorrect:

$$(\{a \mapsto 0\}, (\text{deposit } 100 \ a) \parallel (\text{deposit } 100 \ a)) \xrightarrow{*} (\{a \mapsto 100\}, ())$$

↳ needs an **exclusive access** to y :

$$\text{deposit} = \lambda x \lambda y (\text{lock } y \text{ in } y := !y + x)$$

where $(\text{lock } y \text{ in } e)$ takes the [lock associated with] reference y , and release it upon termination of e .

An EXAMPLE

3
(2/3)

- ▶ to **withdraw** an amount x from account y :

$$\text{withdraw} = \lambda x \lambda y (\text{lock } y \text{ in } (\text{if } !y \geq x \text{ then } (y := !y - x) \\ \text{else error}))$$

- ▶ to **transfer** x from y to z :

$$\text{transfer} = \lambda x \lambda y \lambda z (\text{lock } y \text{ in } (\text{withdraw } xy) ; (\text{deposit } xz))$$

Notice: **reentrant** locks.

A_n EXAMPLE

One may transfer money in any direction:

(transfer 100 a b) || (transfer 10 b a)

where ! $a \geq 100$ and ! $b \geq 10$.

A_n EXAMPLE

(3/3)

One may transfer money in any direction:

$$(\text{transfer } 100 \ a \ b) \parallel (\text{transfer } 10 \ b \ a)$$

where $!a \geq 100$ and $!b \geq 10$.

↳ potential **deadlock**:

$$\xrightarrow{*} \underbrace{(\text{lock } b \text{ in } b := !b + 100)}_{\text{holding } a} \parallel \underbrace{(\text{lock } a \text{ in } a := !a + 10)}_{\text{holding } b}$$

DEADLOCKS: SOLUTIONS

- ▶ deadlock **prevention**: only run code that is guaranteed to be free of deadlocks;
- ▶ deadlock **avoidance**: monitoring the execution so as to avoid dangerous states;
- ▶ deadlock **detection** and **recovery**: supervise execution and rollback (undoing operations) in case of deadlock.

Prevention: by means of static analysis, checking that locks are taken in some **order**.

↳ a **unique** lock for all the bank accounts!

Our SOLUTION

Deadlock avoidance:

- ▶ static analysis by means of a **type and effect** system, anticipating the pointers to lock as the effect,
- ▶ translation of source programs into **annotated** programs

$$(\text{lock } e_0 \text{ in } e_1) \Rightarrow (\text{lock}_{\psi} \bar{e}_0 \text{ in } \bar{e}_1)$$

where ψ is the effect of e_1 = set of pointers to be locked by e_1 ,

- ▶ prudent semantics: to execute $(\text{lock}_{\psi} p \text{ in } e)$ one **does not lock** p if some pointer in ψ is already held by another thread.
- ↳ **type safety**: the annotated programs obtained by translation from typable source programs, executed in the prudent semantics, are **free of deadlocks**.

TECHNICALLY

In analysing $(\text{lock } e_0 \text{ in } e_1)$ one has to have **some information** about the pointer to be locked, i.e. the value of e_0 , to be recorded in the effect, and then used in the types.

- ▶ replace $(\text{ref } e)$ with $(\text{cref } e)$, a **function** f to create a pointer with initial value the value of e ,
- ▶ restricted, by typing, to be used in a particular context, namely

$$(\text{let } x = (f()) \text{ in } e)$$

- ↳ **singleton reference types** θref_x , i.e. locks **univocally** associated with pointers.

SOURCE LANGUAGE

functional + imperative + concurrent:

$v, w \dots ::= x \mid \lambda x e \mid () \mid (\text{cref } v)$	<i>values</i>
$e ::= v \mid (e_1 e_0)$	<i>expressions</i>
$\quad \mid (\text{cref } e) \mid (! e) \mid (e_0 := e_1)$	
$\quad \mid (\text{thread } e) \mid (\text{lock } e_0 \text{ in } e_1)$	

Notation: $(\text{ref } e)$ for $((\text{cref } e)())$.

TARGET LANGUAGE

$p, q \dots$		<i>pointers</i>
$v, w \dots ::=$	$x \mid \lambda x e \mid () \mid p$	<i>values</i>
$e ::=$	$v \mid (e_1 e_0)$	<i>expressions</i>
	$\mid (! e) \mid (e_0 := e_1)$	
	$\mid (\text{thread } e) \mid (\text{lock}_\varphi e_0 \text{ in } e_1)$	
	$\mid (e \setminus p) \mid (\text{new } x \text{ in } e)$	

where φ is an **effect**, that is a finite set of pointer names (either constant or variable)

TYPING and TRANSLATION

(1/3)

Types:

$$\tau, \sigma, \theta \dots ::= \text{unit} \mid \theta \text{ref}_x \mid \theta \text{cref} \mid (\tau \xrightarrow{\varphi} \sigma)$$

In $(\theta \text{ref}_x \xrightarrow{\varphi} \sigma)$ the variable x is **universally quantified**, with scope φ and σ .

Judgements:

$$\Gamma \vdash e : \varphi, \tau \Rightarrow \bar{e}$$

TYPING and TRANSLATION

(2/3)

Main (unusual) rules:

$$\Gamma \vdash_s e_0 : \varphi_0, (\theta \text{ ref}_x \xrightarrow{\varphi_2} \sigma) \Rightarrow \bar{e}_0$$

$$\Gamma \vdash_s e_1 : \varphi_1, \theta \text{ ref}_y \Rightarrow \bar{e}_1$$

$$\Gamma \vdash_s (e_0 e_1) : \varphi_0 \cup \varphi_1 \cup \{x \mapsto y\} \varphi_2, \{x \mapsto y\} \sigma \Rightarrow (\bar{e}_0 \bar{e}_1)$$

$$\Gamma \vdash_s e : \varphi, \theta \Rightarrow \bar{e}$$

$$\Gamma \vdash_s (\text{cref } e) : \varphi, \theta \text{ cref} \Rightarrow (\lambda x \lambda y ((y := x) ; y) \bar{e})$$

i.e.

$$(\text{cref } e) \Rightarrow (\text{let } x = \bar{e} \text{ in } \lambda y ((y := x) ; y))$$

$$\theta \text{ cref} \Rightarrow (\theta \text{ ref}_y \xrightarrow{\emptyset} \theta \text{ ref}_y)$$

TYPING and TRANSLATION

(3/3)

$$\Gamma \vdash_s e_0 : \varphi_0, \theta \text{ cref} \quad \Rightarrow \quad \bar{e}_0$$

$$\Gamma, x : \theta \text{ ref}_x \vdash_s e_1 : \varphi_1, \tau \quad \Rightarrow \quad \bar{e}_1$$

$$\Gamma \vdash_s (\lambda x e_1(e_0())) : \varphi_0 \cup (\varphi_1 - \{x\}), \tau \quad \Rightarrow \quad (\text{new } y \text{ in } (\lambda x \bar{e}_1(\bar{e}_0 y)))$$

where y is fresh and $x \notin \Gamma, \varphi_0, \tau$. Notice: the only way to create a reference is $(\text{let } x = (e_0()) \text{ in } e_1)$ where e_0 has type $\theta \text{ cref}$.

$$\Gamma \vdash_s e_0 : \varphi_0, \theta \text{ ref}_x \quad \Rightarrow \quad \bar{e}_0$$

$$\Gamma \vdash_s e_1 : \varphi_1, \tau \quad \Rightarrow \quad \bar{e}_1$$

$$\Gamma \vdash_s (\text{lock } e_0 \text{ in } e_1) : \{x\} \cup \varphi_0 \cup \varphi_1, \tau \quad \Rightarrow \quad (\text{lock}_{\varphi_1} \bar{e}_0 \text{ in } \bar{e}_1)$$

EXAMPLE

$$\Gamma \vdash \text{deposit} : \emptyset, \text{int} \xrightarrow{\emptyset} (\text{int ref}_y \xrightarrow{\{y\}} \text{unit})$$

$$\Gamma \vdash \text{transfer} : \emptyset, \text{int} \xrightarrow{\emptyset} (\text{int ref}_y \xrightarrow{\emptyset} (\text{int ref}_z \xrightarrow{\{y,z\}} \text{unit}))$$

(polymorphic types) with translations

$$\lambda x \lambda y (\text{lock}_{\emptyset} y \text{ in } y := !y + x)$$

$$\lambda x \lambda y \lambda z (\text{lock}_{\{y,z\}} y \text{ in } (\text{withdraw } xy) ; (\text{deposit } xz))$$

and one can type

let create_account = $\lambda x(\text{cref } x)$ in

let $a = (\text{create_account } 100)()$ in

let $b = (\text{create_account } 10)()$ in ...

where a and b have **distinct** types, int ref_a and int ref_b .

PRUDENT SEMANTICS

Main (unusual) rules:

$$\begin{aligned}
 (S, L, \mathbf{E}[(\text{lock}_\psi p \text{ in } e)] \parallel T) &\rightarrow (S, L, \mathbf{E}[e] \parallel T) && p \in [\mathbf{E}] \\
 (S, L, \mathbf{E}[(\text{lock}_\psi p \text{ in } e)] \parallel T) &\rightarrow (S, L', \mathbf{E}[(e \setminus p)] \parallel T) && p \notin [\mathbf{E}] \\
 &&& \& (\spadesuit)
 \end{aligned}$$

$$(S, L, \mathbf{E}[(v \setminus p)] \parallel T) \rightarrow (S, L - \{p\}, \mathbf{E}[v] \parallel T)$$

where $p \in [\mathbf{E}]$ means that p is currently locked by the thread, and

$$(\spadesuit) \quad L \cap (\{p\} \cup (\psi - [\mathbf{E}])) = \emptyset, \quad L' = L \cup \{p\}$$

vs standard condition: $L \cap \{p\} = \emptyset$.

RESULTS

- ▶ **Type Safety:** if $\Gamma \vdash e : \varphi, \tau \Rightarrow \bar{e}$ then evaluating \bar{e} in the prudent semantics is free of deadlocks.
- ▶ **modularity:** composing systems of (typable) threads is safe – no deadlock.
- ▶ **fine grained** locking policy: each pointer has its own lock (the programmer does not have to think about locks at run time – only pointers).
- ▶ **simple** “pessimistic” semantics: only local (i.e. per thread) conditions, no global analysis of the current state, no rollback.