

Polymorphic Behavioural Lock Effects for Deadlock Checking

Ka I Violet Pun, Martin Steffen, Volker Stolz

PMA Group, University of Oslo, Norway

The 23rd Nordic Workshop for Programming Theory - NWPT '11
Västerås, Sweden

26th ~ 28th October, 2011



- Find *potential* deadlocks in programs *statically* by detecting cyclic wait
 - Each of two or more processes, which form a circular chain, wait for a shared resource that is held by the next process in the chain.
 - Shared resources here: *locks*

- Capture *abstract behaviour* as effects with a type and effect system
- Use *program points* π , to characterize locks according to their origin
- *Execute* the abstract behaviour to detect deadlock
- Limit potential infinite state space by:
 - Put an upper bound for reentrant lock counter
 - Transform effects into coarser, tail-recursive effect
 - Don't allow recursive thread/lock creation
- Prove deadlock preservation by defining a *Deadlock and Termination Sensitive Simulation*

$$\begin{aligned} t &::= \text{stop} \mid v \mid \text{let } x:T = e \text{ in } t \\ e &::= t \mid v \, v \mid \text{if } e \text{ then } e \text{ else } e \mid \text{spawn } t \\ &\quad \mid \text{new } L \mid v.\text{lock} \mid v.\text{unlock} \\ v &::= x \mid l \mid \text{fn } x:T.t \mid \text{fun } f:T.x:T.t \end{aligned}$$

Sequential composition $e_1; e_2$ is represented by let-construct

$$\text{let } x:T = e_1 \text{ in } e_2, \quad x \notin \text{fv}(e_2)$$

$$\begin{aligned} t &::= \text{stop} \mid v \mid \text{let } x:T = e \text{ in } t \\ e &::= t \mid v \mid v \mid \text{if } e \text{ then } e \text{ else } e \mid \text{spawn } t \\ &\quad \mid \text{new } L \mid v.\text{lock} \mid v.\text{unlock} \\ v &::= x \mid l \mid \text{fn } x:T.t \mid \text{fun } f:T.x:T.t \end{aligned}$$

Sequential composition $e_1; e_2$ is represented by let-construct

$$\text{let } x:T = e_1 \text{ in } e_2, \quad x \notin \text{fv}(e_2)$$

Dining Philosophers

```
let l1 = new $\pi_1$  L, l2 = new $\pi_2$  L, l3 = new $\pi_3$  L,
    l4 = new $\pi_4$  L, l5 = new $\pi_5$  L in
let grab = fn:L $\times$ L $\rightarrow$ L. (l, r). l.lock; r.lock in
let release = fn:L $\times$ L $\rightarrow$ L. (l, r). l.unlock; r.unlock in
let phil = fun PHIL:L $\times$ L $\rightarrow$ L.(l, r). think; grab(l, r);
    eat; release(l, r); PHIL (l, r) in
    spawn(phil(l1,l2));...;spawn(phil(l5,l1))
```

$$\begin{array}{l} P ::= \emptyset \mid p\langle t \rangle \mid P \parallel P \quad (\text{Processes}) \\ \sigma \vdash P \rightarrow \sigma' \vdash P' \text{ with } \sigma : L \mapsto \{\text{free}, p(n)\} \quad (\text{Configuration}) \end{array}$$

An example run:

$$\emptyset \vdash p_0\langle t \rangle \rightarrow \dots \rightarrow [l_1 \mapsto p_1(1), l_2 \mapsto p_0(1)] \vdash p_1\langle l_2. \text{lock} \rangle \parallel p_0\langle l_1. \text{lock} \rangle$$

Definition (Waiting for a lock)

Given a configuration $\sigma \vdash P$,

$$waits(\sigma \vdash P, p, l)$$

if it is not the case that $\sigma \vdash P \xrightarrow{p\langle l.lock \rangle}$, and furthermore there exists a σ' s.t. $\sigma' \vdash P \xrightarrow{p\langle l.lock \rangle} \sigma'' \vdash P'$.

Definition (Deadlock)

A configuration $\sigma \vdash P$ is *deadlocked* if $\sigma(l_i) = p_i(n_i)$ and furthermore $waits(\sigma \vdash P, p_i, l_{i+k1})$ (where $k \geq 2$ and for all $0 \leq i \leq k-1$).

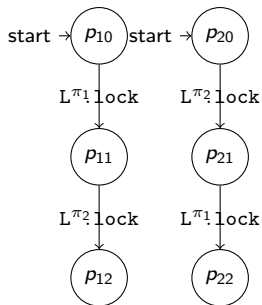


Figure: Deadlock




Figure: Wait-for graph

Type and Effect System

The judgment of our type and effect system is given by:

$$\Gamma \vdash e : T :: \varphi$$

Types and effects are described by:

$$U ::= \text{Bool} \mid \text{Int} \mid \textcolor{blue}{L}^r \mid \text{Thread}$$

basic types

$$T ::= U \mid \vec{U} \xrightarrow{\varphi} U \mid \forall_{\varrho}. T$$

types

$$r ::= \pi \mid \varrho$$

location annotations

Type and Effect System

The judgment of our type and effect system is given by:

$$\Gamma \vdash e : T :: \varphi$$

Types and effects are described by:

$$U ::= \text{Bool} \mid \text{Int} \mid \textcolor{blue}{L}^r \mid \text{Thread}$$

basic types

$$T ::= U \mid \vec{U} \xrightarrow{\varphi} U \mid \forall \varrho. T$$

types

$$r ::= \pi \mid \varrho$$

location annotations

$$\Phi ::= \mathbf{0} \mid p\langle\varphi\rangle \mid \Phi \parallel \Phi$$

effects (global)

$$a ::= \text{spawn } \varphi \mid \nu L^r \mid L^r.\text{lock} \mid L^r.\text{unlock}$$

labels/basic effects

$$\alpha ::= a \mid \tau$$

transition labels

Type and Effect System

The judgment of our type and effect system is given by:

$$\Gamma \vdash e : T :: \varphi$$

Types and effects are described by:

$$U ::= \text{Bool} \mid \text{Int} \mid \textcolor{blue}{L}^r \mid \text{Thread}$$

basic types

$$T ::= U \mid \vec{U} \xrightarrow{\varphi} U \mid \forall \varrho. T$$

types

$$r ::= \pi \mid \varrho$$

location annotations

$$\Phi ::= \mathbf{0} \mid p\langle\varphi\rangle \mid \Phi \parallel \Phi$$

effects (global)

$$\varphi ::= \epsilon \mid X \mid \varphi; \varphi \mid \varphi + \varphi \mid \text{rec } X. \varphi \mid \alpha$$

effects (local)

$$a ::= \text{spawn } \varphi \mid \nu L^r \mid L^r.\text{lock} \mid L^r.\text{unlock}$$

labels/basic effects

$$\alpha ::= a \mid \tau$$

transition labels

Deadlock Checking

To detect a deadlock in a program, we execute the abstract behaviour of the program. In our example:

```
let l1 = new $\pi_1$  L, l2 = new $\pi_2$  L, l3 = new $\pi_3$  L,  
    l4 = new $\pi_4$  L, l5 = new $\pi_5$  L in  
let grab = fn:L×L→L. (l, r). l.lock; r.lock in  
let release = fn:L×L→L. (l, r). l.unlock; r.unlock in  
let phil = fun PHIL:L×L→L.(l, r). think; grab(l, r);  
    eat; release(l, r); PHIL (l, r) in  
    spawn(phil(l1,l2));...;spawn(phil(l5,l1))
```

We have the effect:

$$\nu L^{\pi_1}; \dots; \nu L^{\pi_5}; \text{spawn } (\varphi_p(\pi_1, \pi_2)); \dots; \text{spawn } (\varphi_p(\pi_5, \pi_1))$$
$$\varphi_p(\varrho_1, \varrho_2) = \text{rec } X. \quad \text{think; } L^{\varrho_1}.\text{lock; } L^{\varrho_2}.\text{lock; eat; } \\ L^{\varrho_1}.\text{unlock; } L^{\varrho_2}.\text{unlock; } X$$

Deadlock Checking

$$\begin{array}{l}
 [] \vdash p \langle \nu L^{\pi_1}; \dots; \nu L^{\pi_5}; \text{spawn}(\varphi_p(\pi_1, \pi_2)); \dots; \text{spawn}(\varphi_p(\pi_5, \pi_1)) \rangle \\
 [\pi_1 \mapsto \text{free}] \dots [\pi_5 \mapsto \text{free}] \vdash p \langle \text{spawn}(\varphi_p(\pi_1, \pi_2)); \dots; \text{spawn}(\varphi_p(\pi_5, \pi_1)) \rangle \\
 \vdots \\
 [\pi_1 \mapsto \text{free}] \dots [\pi_5 \mapsto \text{free}] \vdash p_1 \langle L^{\pi_1}. \text{lock}; L^{\pi_2}. \text{lock}; L^{\pi_1}. \text{unlock}; L^{\pi_2}. \text{unlock}; \\
 \quad \text{rec } X. L^{\pi_1}. \text{lock}; \dots \rangle \parallel \dots \parallel \\
 \quad p_5 \langle L^{\pi_5}. \text{lock}; L^{\pi_1}. \text{lock}; L^{\pi_5}. \text{unlock}; L^{\pi_1}. \text{unlock}; \\
 \quad \quad \text{rec } X. L^{\pi_5}. \text{lock}; \dots \rangle \\
 \vdots \\
 [\pi_1 \mapsto p_1(1)][\pi_2 \mapsto p_2(2)] \dots [\pi_5 \mapsto p_5(1)] \vdash p_1 \langle L^{\pi_2}. \text{lock}; L^{\pi_1}. \text{unlock}; L^{\pi_2}. \text{unlock}; \\
 \quad \text{rec } X. L^{\pi_1}. \text{lock}; \dots \rangle \parallel \dots \parallel \\
 \quad p_5 \langle L^{\pi_1}. \text{lock}; L^{\pi_5}. \text{unlock}; L^{\pi_1}. \text{unlock}; \\
 \quad \quad \text{rec } X. L^{\pi_5}. \text{lock}; \dots \rangle
 \end{array}
 \begin{array}{l}
 \xrightarrow{p \langle \nu L^{\pi_1}; \dots; \nu L^{\pi_5} \rangle} \\
 \xrightarrow{p \langle \text{spawn}(\varphi_p(\pi_1, \pi_2)); \dots \rangle} \\
 \\
 \\
 \xrightarrow{p_1 \langle L^{\pi_1} \text{ lock} \rangle} \\
 \\
 \xrightarrow{p_1 \langle L^{\pi_2} \text{ lock} \rangle}
 \end{array}$$

Deadlock and termination sensitive simulation $\lesssim^D / \lesssim^{DT}$

$$\begin{array}{c} \sigma_1 \vdash \Phi_1 \text{ --- } R \text{ --- } \sigma_2 \vdash \Phi_2 \\ \downarrow p\langle \tau \rangle \qquad \qquad \qquad \vdots = /p\langle \tau \rangle \\ \sigma_1 \vdash \Phi'_1 \text{ } R \text{ } \sigma_2 \vdash \Phi'_2 \end{array}$$

(a)

$$\begin{array}{c} \sigma_1 \vdash \Phi_1 \text{ --- } R \text{ --- } \sigma_2 \vdash \Phi_2 \\ \downarrow p\langle a \rangle \qquad \qquad \qquad \vdots p\langle a \rangle \\ \sigma'_1 \vdash \Phi'_1 \text{ } R \text{ } \sigma'_2 \vdash \Phi'_2 \end{array}$$

(b)

$$\begin{array}{c} \sigma_1 \vdash \Phi_1 \text{ --- } R \text{ --- } \sigma_2 \vdash \Phi_2 \\ \downarrow \neg p\langle L^\pi \text{lock} \rangle \qquad \qquad \qquad \vdots \neg p\langle L^\pi \text{lock} \rangle \\ \sigma'_1 \vdash \Phi'_1 \qquad \qquad \qquad \sigma'_2 \vdash \Phi'_2 \end{array}$$

(c)

$$\begin{array}{c} \sigma_1 \vdash \Phi_1 \text{ --- } R \text{ --- } \sigma_2 \vdash \Phi_2 \\ \downarrow p\langle \surd \rangle \qquad \qquad \qquad \vdots p\langle \surd \rangle \\ \sigma_1 \vdash \Phi'_1 \text{ } R \text{ } \sigma_2 \vdash \Phi'_2 \end{array}$$

(d)

Two sources of infinity

- Unboundedness of *reentrant* lock counters
- Unboundedness of the “control stack” of *non-tail recursive* behaviour descriptions

Problem in state space:

Unbounded lock counters counting uuuuuuupppppp
(with recursion)...

Solution:


Fix upper bound; unlocking from upper bound becomes non-deterministic.

Lemma

Given a configuration $\sigma \vdash \Phi$, and let further denote $\sigma_1 \vdash^{n_1} \Phi$ and $\sigma_2 \vdash^{n_2} \Phi$ the corresponding configurations under the lock-counter abstraction. If $n_1 \geq n_2$, then $\sigma_1 \vdash^{n_1} \Phi \lesssim^D \sigma_2 \vdash^{n_2} \Phi$.

Lemma (Ω is maximal wrt. \lesssim^{DT})

Assume φ over a set of locations r , then $\sigma \vdash p\langle\varphi\rangle \lesssim^{DT} \sigma \vdash p\langle\Omega\rangle$.



Theorem (Finite abstractions)

The lock counter abstraction and behavior abstraction (when abstracting all locks and recursions) results in a finite state space.

Theorem (Soundness of the abstraction)

Given $\Gamma \vdash P : ok :: \Phi$ and two heaps $\sigma_1 \equiv \sigma_2$. Further, $\sigma'_2 \vdash \Phi'$ is obtained by lock-counter resp. behavior abstraction of $\sigma_2 \vdash \Phi$. Then if $\sigma'_2 \vdash \Phi'$ is deadlock free then so is $\sigma_1 \vdash P$.

- Conclusion:
 - We have proven that our type systems is correct in the aspect of capturing behavior of a program
 - Abstract behavior correctly over-approximates the concrete one
 - Deadlocks in a program are correctly detected in the abstract run. . .
 - Inference algorithm is partially formalized with `Ott` and `Coq`
- Future Work:
 - Applying to communication analysis of asynchronous systems
 - Relaxing the condition (e.g. lock creation in loop)
 - Abstracting processes
 - Implement our algorithm with model checker for real language
 - CEGAR - Counter-Example Guided Abstraction Refinement