

Static Analysis of Concurrent Programs

MPRI 2–6: Abstract Interpretation,
application to verification and static analysis

Antoine Miné

year 2013–2014

course 08-A

15 November 2013

Concurrent programming

Idea:

Decompose a program into a **set** of (loosely) interacting processes.

Why concurrent programs?

- **exploit** parallelism in current computers
(multi-processors, multi-cores, hyper-threading)

“Free lunch is over”
change in Moore's law ($\times 2$ transistors every 2 years)
- **exploit** several computers
(distributed computing)
- **ease** of programming
(GUI, network code, reactive programs)

Models of concurrent programs

Many models:

- process calculi: CSP, π -calculus, join calculus
- message passing
- shared memory (threads)
- transactional memory
- combination of several models

Example implementations:

- C, C++, etc. with a thread library (POSIX threads, Win32)
- C, C++, etc. with a message library (MPI, OpenMP)
- Java (native threading API)
- Erlang (based on π -calculus)
- JoCaml (OCaml + join calculus)
- processor-level (interrupts, test-and-set instructions)

Scope

In this talk: **thread model**

- implicit communication through **shared memory**
- explicit communication through **synchronisation** primitives
- **fixed** number of threads (no dynamic creation of threads)
- numeric programs (real-valued variables)

Goal: **static analysis**

- infer numeric program **invariants**
- discover possible **run-time errors** (e.g., division by 0)
- parametrized by a choice of abstract domains

Outline

- State-based analyses
 - sequential programs (reminders)
 - concurrent programs
- Toward thread-modular analyses
 - detour through proof methods (Floyd–Hoare, Owicki–Gries, Jones)
 - rely-guarantee in abstract interpretation form
- Interference-based abstract analyses
 - denotational-style analysis
 - weakly consistent memory models
 - synchronisation

Simple structured numeric language

- finite set of (toplevel) **threads**: $stat_1$ to $stat_n$
- finite set of numeric program variables $X \in \mathbb{V}$
- finite set of statement locations $\ell \in \mathcal{L}$
- finite set of potential error locations $\omega \in \Omega$

Language syntax

$prog ::= \ell stat_1^\ell \parallel \dots \parallel \ell stat_n^\ell$ (parallel composition)

$\ell stat^\ell ::= \ell X \leftarrow expr^\ell$ (assignment)

| $\ell \mathbf{if} \ expr \bowtie 0 \ \mathbf{then} \ \ell stat^\ell$ (conditional)

| $\ell \mathbf{while} \ \ell expr \ \bowtie 0 \ \mathbf{do} \ \ell stat^\ell$ (loop)

| $\ell stat; \ell stat^\ell$ (sequence)

$expr ::= X \mid [c_1, c_2] \mid - expr \mid expr \diamond_\omega expr$

$c_1, c_2 \in \mathbb{R} \cup \{+\infty, -\infty\}$, $\diamond \in \{+, -, \times, /\}$, $\bowtie \in \{=, \leq, \dots\}$

State-based analyses

Sequential program semantics (reminders)

Transition systems

Transition system: $(\Sigma, \tau, \mathcal{I})$

- Σ : a set of program states
- $\tau \subseteq \Sigma \times \Sigma$: transition relation
we note $(\sigma, \sigma') \in \tau$ as $\sigma \rightarrow_{\tau} \sigma'$
- $\mathcal{I} \subseteq \Sigma$: a set of initial states

Traces: sequences of states $\sigma_0, \dots, \sigma_n, \dots$

- Σ^* : finite traces
- Σ^{ω} : infinite countable traces
- $\Sigma^{\infty} \stackrel{\text{def}}{=} \Sigma^* \cup \Sigma^{\omega}$: finite or infinite countable traces
- $u \preceq t$: u is a prefix of t

We view program semantics and properties as sets of traces.

Traces of a transition system

Maximal trace semantics: $\mathcal{M}_\infty \in \mathcal{P}(\Sigma^\infty)$

- set of total executions $\sigma_0, \dots, \sigma_n, \dots$
 - starting in an initial state $\sigma_0 \in \mathcal{I}$ and either
 - **ending** in a blocking state in $\mathcal{B} \stackrel{\text{def}}{=} \{ \sigma \mid \forall \sigma' : \sigma \not\rightarrow_\tau \sigma' \}$
 - or **infinite**

$$\mathcal{M}_\infty \stackrel{\text{def}}{=} \{ \sigma_0, \dots, \sigma_n \mid \sigma_0 \in \mathcal{I} \wedge \sigma_n \in \mathcal{B} \wedge \forall i < n : \sigma_i \rightarrow_\tau \sigma_{i+1} \} \cup \{ \sigma_0, \dots, \sigma_n \dots \mid \sigma_0 \in \mathcal{I} \wedge \forall i : \sigma_i \rightarrow_\tau \sigma_{i+1} \}$$

- able to express many properties of programs, e.g.:
 - **safety:** $\mathcal{M}_\infty \subseteq S^\infty$ (executions stay in S)
 - **ordering:** $\mathcal{M}_\infty \subseteq S_1^\infty \cdot S_2^\infty$ (S_2 can only occur after S_1)
 - **termination:** $\mathcal{M}_\infty \subseteq \Sigma^*$ (executions are finite)
 - **inevitability:** $\mathcal{M}_\infty \subseteq \Sigma^* \cdot S \cdot \Sigma^\infty$ (executions pass through S)

Traces of a transition system

Finite prefix trace semantics: $\mathcal{T}_p \subseteq \mathcal{P}(\Sigma^*)$

- set of **finite prefixes** of executions:

$$\mathcal{T}_p \stackrel{\text{def}}{=} \{ \sigma_0, \dots, \sigma_n \mid \sigma_0 \in \mathcal{I}, \forall i < n: \sigma_i \rightarrow_{\tau} \sigma_{i+1} \}$$

- \mathcal{T}_p is an abstraction of the maximal trace semantics

$$\mathcal{T}_p = \alpha_{* \preceq}(\mathcal{M}_{\infty}) \text{ where } \alpha_{* \preceq}(X) \stackrel{\text{def}}{=} \{ t \in \Sigma^* \mid \exists u \in X: t \preceq u \}$$

- \mathcal{T}_p can prove **safety** properties: $\mathcal{T}_p \subseteq S^*$ (executions stay in S)

$$\mathcal{T}_p \text{ can prove } \text{ordering} \text{ properties: } \mathcal{T}_p \subseteq S_1^* \cdot S_2^*$$

(if S_1 and S_2 occur, S_2 occurs after S_1)

- \mathcal{T}_p **cannot** prove **termination** nor **inevitability** properties

- **fixpoint characterisation:** $\mathcal{T}_p = \text{lfp } F_p$ where

$$F_p(X) = \mathcal{I} \cup \{ \sigma_0, \dots, \sigma_{n+1} \mid \sigma_0, \dots, \sigma_n \in X \wedge \sigma_n \rightarrow_{\tau} \sigma_{n+1} \}$$

State abstraction

Reachable state semantics: $\mathcal{R} \subseteq \mathcal{P}(\Sigma)$

- set of states **reachable** in any execution:

$$\mathcal{R} \stackrel{\text{def}}{=} \{ \sigma \mid \exists \sigma_0, \dots, \sigma_n : \sigma_0 \in \mathcal{I}, \forall i < n : \sigma_i \rightarrow_{\tau} \sigma_{i+1} \wedge \sigma = \sigma_n \}$$

- \mathcal{R} is an abstraction of the finite trace semantics: $\mathcal{R} = \alpha_p(\mathcal{T}_p)$

$$\text{where } \alpha_p(X) \stackrel{\text{def}}{=} \{ \sigma \mid \exists \sigma_0, \dots, \sigma_n \in X : \sigma = \sigma_n \}$$

- \mathcal{R} **can** prove **safety** properties: $\mathcal{R} \subseteq S$ (executions stay in S)
 \mathcal{R} **cannot** prove **ordering**, **termination**, **inevitability** properties

- **fixpoint characterisation:** $\mathcal{R} = \text{lfp } F_{\mathcal{R}}$ where

$$F_{\mathcal{R}}(X) = \mathcal{I} \cup \{ \sigma \mid \exists \sigma' \in X : \sigma' \rightarrow_{\tau} \sigma \}$$

States of a sequential program

Simple sequential numeric program: $prog = li\ stat^{lx}$.

Program states: $\Sigma \stackrel{\text{def}}{=} (\mathcal{L} \times \mathcal{E}) \cup \Omega$

- a **control** state in \mathcal{L}
- a **memory** state: an environment in $\mathcal{E} \stackrel{\text{def}}{=} \mathbb{V} \rightarrow \mathbb{R}$
- an **error** state in Ω

Initial states:

start at the first control point li , and with variables set to 0:

$$\mathcal{I} \stackrel{\text{def}}{=} \{ (li, \lambda V.0) \}$$

Note that $\mathcal{P}(\Sigma) \simeq (\mathcal{L} \rightarrow \mathcal{P}(\mathcal{E})) \times \mathcal{P}(\Omega)$:

- a state property in $\mathcal{P}(\mathcal{E})$ at each program point in \mathcal{L}
- and a set of errors in $\mathcal{P}(\Omega)$

Expression semantics with errors

Expression semantics: $E[\![expr]\!] : \mathcal{E} \rightarrow (\mathcal{P}(\mathbb{R}) \times \mathcal{P}(\Omega))$

$$E[\![X]\!] \rho \stackrel{\text{def}}{=} \langle \{ \rho(X) \}, \emptyset \rangle$$

$$E[\![[c_1, c_2]]\!] \rho \stackrel{\text{def}}{=} \langle \{ x \in \mathbb{R} \mid c_1 \leq x \leq c_2 \}, \emptyset \rangle$$

$$E[\![-e_1]\!] \rho \stackrel{\text{def}}{=} \text{let } \langle V_1, O_1 \rangle = E[\![e_1]\!] \rho \text{ in} \\ \langle \{ -v_1 \mid v_1 \in V_1 \}, O_1 \rangle$$

$$E[\![e_1 \diamond_{\omega} e_2]\!] \rho \stackrel{\text{def}}{=} \text{let } \forall i \in \{1, 2\}: \langle V_i, O_i \rangle = E[\![e_i]\!] \rho \text{ in} \\ \langle \{ v_1 \diamond v_2 \mid v_i \in V_i, \diamond \neq / \vee v_2 \neq 0 \}, \\ O_1 \cup O_2 \cup \{ \omega \text{ if } \diamond = / \wedge 0 \in V_2 \} \rangle$$

- defined by structural induction on the syntax
- evaluates in an environment ρ to a **set of values**
- also returns a set of **accumulated errors** (divisions by zero)

Reminders: semantics in equational form

Principle: (without handling errors in Ω)

- see lfp f as the least solution of an equation $x = f(x)$
- partition states by control: $\mathcal{P}(\mathcal{L} \times \mathcal{E}) \simeq \mathcal{L} \rightarrow \mathcal{P}(\mathcal{E})$

$\mathcal{X}_\ell \in \mathcal{P}(\mathcal{E})$: invariants at $\ell \in \mathcal{L}$

$\forall \ell \in \mathcal{L}: \mathcal{X}_\ell \stackrel{\text{def}}{=} \{m \in \mathcal{E} \mid (\ell, m) \in \mathcal{R}\}$

\implies set of (recursive) equations on \mathcal{X}_ℓ

Example:

ℓ_1 $i \leftarrow 2;$

ℓ_2 $n \leftarrow [-\infty, +\infty];$

ℓ_3 **while** ℓ_4 $i < n$ **do**

ℓ_5 **if** $[0, 1] = 0$ **then**

ℓ_6 $i \leftarrow i + 1$

ℓ_7

ℓ_8

$\mathcal{X}_1 = \mathcal{I}$

$\mathcal{X}_2 = \mathcal{C}[[i \leftarrow 2]] \mathcal{X}_1$

$\mathcal{X}_3 = \mathcal{C}[[n \leftarrow [-\infty, +\infty]]] \mathcal{X}_2$

$\mathcal{X}_4 = \mathcal{X}_3 \cup \mathcal{X}_7$

$\mathcal{X}_5 = \mathcal{C}[[i < n]] \mathcal{X}_4$

$\mathcal{X}_6 = \mathcal{X}_5$

$\mathcal{X}_7 = \mathcal{X}_5 \cup \mathcal{C}[[i \leftarrow i + 1]] \mathcal{X}_6$

$\mathcal{X}_8 = \mathcal{C}[[i \geq n]] \mathcal{X}_4$

Semantics in denotational form

Input-output function $C[\textit{stat}]$

$$\underline{C[\textit{stat}] : (\mathcal{P}(\mathcal{E}) \times \mathcal{P}(\Omega)) \rightarrow (\mathcal{P}(\mathcal{E}) \times \mathcal{P}(\Omega))}$$

$$C[X \leftarrow e] \langle R, O \rangle \stackrel{\text{def}}{=} \langle \emptyset, O \rangle \sqcup \bigsqcup_{\rho \in R} \langle \{ \rho[X \mapsto v] \mid v \in V_\rho \}, O_\rho \rangle$$

$$C[e \bowtie 0?] \langle R, O \rangle \stackrel{\text{def}}{=} \langle \emptyset, O \rangle \sqcup \bigsqcup_{\rho \in R} \langle \{ \rho \mid \exists v \in V_\rho : v \bowtie 0 \}, O_\rho \rangle$$

$$\text{where } \langle V_\rho, O_\rho \rangle \stackrel{\text{def}}{=} E[e] \rho$$

$$C[\text{if } e \bowtie 0 \text{ then } s] X \stackrel{\text{def}}{=} (C[s] \circ C[e \bowtie 0?]) X \sqcup C[e \nabla 0?] X$$

$$C[\text{while } e \bowtie 0 \text{ do } s] X \stackrel{\text{def}}{=} \\ C[e \nabla 0?] (\text{Ifp } \lambda Y. X \sqcup (C[s] \circ C[e \bowtie 0?]) Y)$$

$$C[s_1; s_2] \stackrel{\text{def}}{=} C[s_2] \circ C[s_1]$$

- mutate memory states in \mathcal{E} , accumulate errors in Ω
(\sqcup is the element-wise \cup in $\mathcal{P}(\mathcal{E}) \times \mathcal{P}(\Omega)$)
- structured: nested loops yield nested fixpoints
- big-step: forget information on intermediate locations ℓ

Abstract semantics in denotational form

Extend a numeric abstract domain $\mathcal{E}^\#$ abstracting $\mathcal{P}(\mathcal{E})$
to $\mathcal{D}^\# \stackrel{\text{def}}{=} \mathcal{E}^\# \times \mathcal{P}(\Omega)$.

$$\underline{C^\#[\textit{stat}]} : \mathcal{D}^\# \rightarrow \mathcal{D}^\#$$

$C^\#[X \leftarrow e] \langle R^\#, O \rangle$ and $C^\#[e \bowtie 0?] \langle R^\#, O \rangle$ are given

$$C^\#[\textit{if } e \bowtie 0 \textit{ then } s] X^\# \stackrel{\text{def}}{=} (C^\#[s] \circ C^\#[e \bowtie 0?]) X^\# \sqcup^\# C^\#[e \nabla 0?] X^\#$$

$$C^\#[\textit{while } e \bowtie 0 \textit{ do } s] X^\# \stackrel{\text{def}}{=} C^\#[e \nabla 0?] (\textit{lim} \lambda Y^\#. Y^\# \nabla (X^\# \sqcup^\# (C^\#[s] \circ C^\#[e \bowtie 0?]) Y^\#))$$

$$C^\#[s_1; s_2] \stackrel{\text{def}}{=} C^\#[s_2] \circ C^\#[s_1]$$

$$C^\#[s_1; s_2] \stackrel{\text{def}}{=} C^\#[s_2] \circ C^\#[s_1]$$

$$C^\#[s_1; s_2] \stackrel{\text{def}}{=} C^\#[s_2] \circ C^\#[s_1]$$

- the abstract interpreter mimicks an actual interpreter
- efficient in memory (intermediate invariants are not kept)
- less flexibility in the iteration scheme (iteration order, widening points, etc.)

Concurrent program semantics

Labelled transition systems

Labelled transition system: $(\Sigma, \mathcal{A}, \tau, \mathcal{I})$

- Σ : set of program states
- \mathcal{A} : set of actions
- $\tau \subseteq \Sigma \times \mathcal{A} \times \Sigma$: transition relation
we note $(\sigma, a, \sigma') \in \tau$ as $\sigma \xrightarrow{a}_{\tau} \sigma'$
- $\mathcal{I} \subseteq \Sigma$: set of initial states

Labelled traces: sequences of states interspersed with actions

denoted as $\sigma_0 \xrightarrow{a_0} \sigma_1 \xrightarrow{a_1} \cdots \sigma_n \xrightarrow{a_n} \sigma_{n+1}$

From concurrent programs to labelled transition systems

Notations:

- concurrent program: $prog ::= \ell_1^i stat_1 \ell_1^x \parallel \dots \parallel \ell_n^i stat_n \ell_n^x$
- threads are identified by number in $\mathbb{T} \stackrel{\text{def}}{=} \{1, \dots, n\}$

Program states: $\Sigma \stackrel{\text{def}}{=} ((\mathbb{T} \rightarrow \mathcal{L}) \times \mathcal{E}) \cup \Omega$

- control state $L(t) \in \mathcal{L}$ for each thread $t \in \mathbb{T}$, and
- single shared memory state $\rho \in \mathcal{E}$
- or error state in $\omega \in \Omega$

Initial states:

threads start at their first control point ℓ_t^i , variables are set to 0:

$$\mathcal{I} \stackrel{\text{def}}{=} \{(\lambda t. \ell_t^i, \lambda V. 0)\}$$

Actions: thread identifiers: $\mathcal{A} \stackrel{\text{def}}{=} \mathbb{T}$

From concurrent programs to labelled transition systems

Transition relation: $\tau \subseteq \Sigma \times \mathcal{A} \times \Sigma$

$$(L, \rho) \xrightarrow{t}_{\tau} (L', \rho') \iff (L(t), \rho) \rightarrow_{\tau[\text{stat}_t]} (L'(t), \rho') \wedge \forall u \neq t: L(u) = L'(u)$$

$$(L, \rho) \xrightarrow{t}_{\tau} \omega \iff (L(t), \rho) \rightarrow_{\tau[\text{stat}_t]} \omega$$

- based on the transition relation of individual threads seen as sequential processes stat_t :¹

$$\tau[\text{stat}] \subseteq (\mathcal{L} \times \mathcal{E}) \times ((\mathcal{L} \times \mathcal{E}) \cup \Omega)$$

- choose a thread t to run
 - update ρ and $L(t)$
 - leave $L(u)$ intact for $u \neq t$
- each $\sigma \rightarrow \sigma'$ in $\tau[\text{stat}_t]$ leads to many transitions in τ !

¹See lesson 02-B for the full definition of $\tau[\text{stat}]$.

Interleaved trace semantics

Maximal and finite prefix trace semantics are as before:

Blocking states: $\mathcal{B} \stackrel{\text{def}}{=} \{ \sigma \mid \forall \sigma', t: \sigma \not\rightarrow_{\tau}^t \sigma' \}$

Maximal traces: \mathcal{M}_{∞} (finite or infinite)

$\mathcal{M}_{\infty} \stackrel{\text{def}}{=} \{ \sigma_0 \xrightarrow{t_0} \dots \xrightarrow{t_{n-1}} \sigma_n \mid \sigma_0 \in \mathcal{I} \wedge \sigma_n \in \mathcal{B} \wedge \forall i < n: \sigma_i \xrightarrow{t_i} \sigma_{i+1} \} \cup$
 $\{ \sigma_0 \xrightarrow{t_0} \sigma_1 \dots \mid \sigma_0 \in \mathcal{I} \wedge \forall i: \sigma_i \xrightarrow{t_i} \sigma_{i+1} \}$

Finite prefix traces: \mathcal{T}_p

$\mathcal{T}_p \stackrel{\text{def}}{=} \{ \sigma_0 \xrightarrow{t_0} \dots \xrightarrow{t_{n-1}} \sigma_n \mid \sigma_0 \in \mathcal{I} \wedge \forall i < n: \sigma_i \xrightarrow{t_i} \sigma_{i+1} \}$

Fixpoint form: $\mathcal{T}_p = \text{lfp } F_p$ where

$F_p(X) = \mathcal{I} \cup \{ \sigma_0 \xrightarrow{t_0} \dots \xrightarrow{t_n} \sigma_{n+1} \mid \sigma_0 \xrightarrow{t_0} \dots \xrightarrow{t_{n-1}} \sigma_n \in X \wedge \sigma_n \xrightarrow{t_n} \sigma_{n+1} \}$

Abstraction: $\mathcal{T}_p = \alpha_{* \preceq}(\mathcal{M}_{\infty})$

Fairness

Fairness conditions: avoid threads being denied to run

Given $enabled(\sigma, t) \stackrel{\text{def}}{\iff} \exists \sigma' \in \Sigma: \sigma \xrightarrow{t}_\tau \sigma'$,

an infinite trace $\sigma_0 \xrightarrow{t_0} \dots \sigma_n \xrightarrow{t_n} \dots$ is:

- **weakly fair** if $\forall t \in \mathbb{T}$:
 $(\exists i: \forall j \geq i: enabled(\sigma_j, t)) \implies (\forall i: \exists j \geq i: a_j = t)$
 (no thread can be continuously enabled without running)
- **strongly fair** if $\forall t \in \mathbb{T}$:
 $(\forall i: \exists j \geq i: enabled(\sigma_j, t)) \implies (\forall i: \exists j \geq i: a_j = t)$
 (no thread can be infinitely often enabled without running)

Proofs under fairness conditions given:

- the maximal traces \mathcal{M}_∞ of a program
- a property X to prove (as a set of traces)
- the set F of all (weakly or strongly) fair and of finite traces
 \implies prove $\mathcal{M}_\infty \cap F \subseteq X$ instead of $\mathcal{M}_\infty \subseteq X$

Fairness (cont.)

Example: **while** $x \geq 0$ **do** $x \leftarrow x + 1$ **||** $x \leftarrow -1$

- **may not** terminate **without fairness**
- **always** terminates under **weak** and **strong fairness**

Finite prefix traces

$\mathcal{M}_\infty \cap F \subseteq X$ reduces to $\alpha_{*\underline{\gamma}}(\mathcal{M}_\infty \cap F) \subseteq \alpha_{*\underline{\gamma}}(X)$

for all fairness conditions F , $\alpha_{*\underline{\gamma}}(\mathcal{M}_\infty \cap F) = \alpha_{*\underline{\gamma}}(\mathcal{M}_\infty) = \mathcal{T}_p$

\implies fairness-dependent properties cannot be proved with finite prefixes

In the following, we ignore fairness conditions.

(see [Cous85])

Equational state semantics

State abstraction \mathcal{R} : as before

- $\mathcal{R} \stackrel{\text{def}}{=} \{ \sigma \mid \exists \sigma_0 \xrightarrow{t_0} \dots \sigma_n : \sigma_0 \in \mathcal{I} \forall i < n : \sigma_i \xrightarrow{t_i} \sigma_{i+1} \wedge \sigma = \sigma_n \}$
- $\mathcal{R} = \alpha_p(\mathcal{T}_p)$ where $\alpha_p(X) \stackrel{\text{def}}{=} \{ \sigma \mid \exists \sigma_0 \xrightarrow{t_0} \dots \sigma_n \in X : \sigma = \sigma_n \}$
- $\mathcal{R} = \text{lfp } F_{\mathcal{R}}$ where $F_{\mathcal{R}}(X) = \mathcal{I} \cup \{ \sigma \mid \exists \sigma' \in X, t \in \mathbb{T} : \sigma' \xrightarrow{t} \sigma \}$

Equational form: (without handling errors in Ω)

- for each $L \in \mathbb{T} \rightarrow \mathcal{L}$, a variable \mathcal{X}_L with value in \mathcal{E}
- equations are derived from thread equations $eq(stat_t)$ as:²

$$\mathcal{X}_{L_1} = \bigcup_{t \in \mathbb{T}} \{ F(\mathcal{X}_{L_2}, \dots, \mathcal{X}_{L_N}) \mid$$

$$\exists (\mathcal{X}_{\ell_1} = F(\mathcal{X}_{\ell_2}, \dots, \mathcal{X}_{\ell_N})) \in eq(stat_t) :$$

$$\forall i \leq N : L_i(t) = \ell_i, \forall u \neq t : L_i(u) = L_1(u) \}$$

(join with \cup equations updating a single thread)

²See lesson 02-B for the definition of $eq(stat)$.

Equational state semantics (example)

Example: inferring $0 \leq x \leq y \leq 10$

t_1	t_2
while ^{ℓ_1} $0 = 0$ do ^{ℓ_2} if $x < y$ then ^{ℓ_3} $x \leftarrow x + 1$	while ^{ℓ_4} $0 = 10$ do ^{ℓ_5} if $y < 10$ then ^{ℓ_6} $y \leftarrow y + 1$

Equational state semantics (example)

Example: inferring $0 \leq x \leq y \leq 10$

t_1	t_2
while ℓ^1 $0 = 0$ do ℓ^2 if $x < y$ then ℓ^3 $x \leftarrow x + 1$	while ℓ^4 $0 = 0$ do ℓ^5 if $y < 10$ then ℓ^6 $y \leftarrow y + 1$

(Simplified) equation system:

$$\begin{aligned}
 \mathcal{X}_{1,4} &= \mathcal{I} \cup \mathcal{C}[x \leftarrow x + 1] \mathcal{X}_{3,4} \cup \mathcal{C}[x \geq y] \mathcal{X}_{2,4} \\
 &\quad \cup \mathcal{C}[y \leftarrow y + 1] \mathcal{X}_{1,6} \cup \mathcal{C}[y \geq 10] \mathcal{X}_{1,5} \\
 \mathcal{X}_{2,4} &= \mathcal{X}_{1,4} \cup \mathcal{C}[y \leftarrow y + 1] \mathcal{X}_{2,6} \cup \mathcal{C}[y \geq 10] \mathcal{X}_{2,5} \\
 \mathcal{X}_{3,4} &= \mathcal{C}[x < y] \mathcal{X}_{2,4} \cup \mathcal{C}[y \leftarrow y + 1] \mathcal{X}_{3,6} \cup \mathcal{C}[y \geq 10] \mathcal{X}_{3,5} \\
 \mathcal{X}_{1,5} &= \mathcal{C}[x \leftarrow x + 1] \mathcal{X}_{3,5} \cup \mathcal{C}[x \geq y] \mathcal{X}_{2,5} \cup \mathcal{X}_{1,4} \\
 \mathcal{X}_{2,5} &= \mathcal{X}_{1,5} \cup \mathcal{X}_{2,4} \\
 \mathcal{X}_{3,5} &= \mathcal{C}[x < y] \mathcal{X}_{2,5} \cup \mathcal{X}_{3,4} \\
 \mathcal{X}_{1,6} &= \mathcal{C}[x \leftarrow x + 1] \mathcal{X}_{3,6} \cup \mathcal{C}[x \geq y] \mathcal{X}_{2,6} \cup \mathcal{C}[y < 10] \mathcal{X}_{1,5} \\
 \mathcal{X}_{2,6} &= \mathcal{X}_{1,6} \cup \mathcal{C}[y < 10] \mathcal{X}_{2,5} \\
 \mathcal{X}_{3,6} &= \mathcal{C}[x < y] \mathcal{X}_{2,6} \cup \mathcal{C}[y < 10] \mathcal{X}_{3,5}
 \end{aligned}$$

Equational state semantics (example)

Example: inferring $0 \leq x \leq y \leq 10$

t_1	t_2
while ℓ^1 $0 = 0$ do ℓ^2 if $x < y$ then ℓ^3 $x \leftarrow x + 1$	while ℓ^4 $0 = 0$ do ℓ^5 if $y < 10$ then ℓ^6 $y \leftarrow y + 1$

Pros:

- easy to construct
- easy to further abstract in an abstract domain \mathcal{E}^\sharp

Cons:

- explosion of the number of variables and equations
- explosion of the size of equations
 \implies efficiency issues
- the equation system does *not* reflect the program structure
 (not defined by induction on the concurrent program)

Wish-list

We would like to:

- keep information attached to syntactic program locations
(control points in \mathcal{L} , not control point tuples in $\mathbb{T} \rightarrow \mathcal{L}$)
- be able to abstract away control information
(precision/cost trade-off control)
- avoid duplicating thread instructions
- have a computation structure based on the program syntax
(denotational style)

Ideally:

thread-modular denotational-style semantics

(analyze each thread independently by induction on its syntax)

Detour through proof methods

Floyd–Hoare logic

Logic to prove properties about **sequential** programs [Hoar69].

Hoare triples: $\{P\} \text{stat} \{Q\}$

- annotate programs with **logic assertions** $\{P\} \text{stat} \{Q\}$
(if P holds before stat , then Q holds after stat)
- check that $\{P\} \text{stat} \{Q\}$ is derivable with the following rules
(the assertions are program invariants)

$$\frac{}{\{P[e/x]\} X \leftarrow e \{P\}} \qquad \frac{\{P \wedge e \bowtie 0\} s \{Q\} \quad P \wedge e \not\bowtie 0 \Rightarrow Q}{\{P\} \mathbf{if} \ e \bowtie 0 \ \mathbf{then} \ s \{Q\}}$$

$$\frac{\{P\} s_1 \{Q\} \quad \{Q\} s_2 \{R\}}{\{P\} s_1; s_2 \{R\}} \qquad \frac{\{P \wedge e \bowtie 0\} s \{P\}}{\{P\} \mathbf{while} \ e \bowtie 0 \ \mathbf{do} \ s \{P \wedge e \not\bowtie 0\}}$$

$$\frac{\{P'\} s \{Q'\} \quad P \Rightarrow P' \quad Q' \Rightarrow Q}{\{P\} s \{Q\}}$$

Floyd–Hoare logic as abstract interpretation

Link with the equational state semantics:

Correspondence between $\ell \text{ stat } \ell'$ and $\{P\} \text{ stat } \{Q\}$:

- if P (resp. Q) models exactly the points in \mathcal{X}_ℓ (resp. $\mathcal{X}_{\ell'}$) then $\{P\} \text{ stat } \{Q\}$ is a derivable Hoare triple
- if $\{P\} \text{ stat } \{Q\}$ is derivable, then $\mathcal{X}_\ell \models P$ and $\mathcal{X}_{\ell'} \models Q$ (all the points in \mathcal{X}_ℓ (resp. $\mathcal{X}_{\ell'}$) satisfy P (resp. Q))

$\implies \mathcal{X}_\ell$ provide the most precise Hoare assertions in a **constructive form**

- $\gamma(\mathcal{X}^\#)$ provide (less precise) Hoare assertions in a **computable form**

Link with the denotational semantics:

both $C \llbracket \text{stat} \rrbracket$ and the proof tree for $\{P\} \text{ stat } \{Q\}$ reflect the syntactic structure of stat (compositional methods)

Owicki–Gries proof method

Extension of Floyd–Hoare to **concurrent** programs [Owic76].

Principle: add a new rule, for \parallel

$$\frac{\{P_1\} s_1 \{Q_1\} \quad \{P_2\} s_2 \{Q_2\}}{\{P_1 \wedge P_2\} s_1 \parallel s_2 \{Q_1 \wedge Q_2\}}$$

Owicki–Gries proof method

Extension of Floyd–Hoare to **concurrent** programs [Owic76].

Principle: add a new rule, for \parallel

$$\frac{\{P_1\} s_1 \{Q_1\} \quad \{P_2\} s_2 \{Q_2\}}{\{P_1 \wedge P_2\} s_1 \parallel s_2 \{Q_1 \wedge Q_2\}}$$

This rule is **not always sound!**

e.g., we have $\{X = 0, Y = 0\} X \leftarrow 1 \{X = 1, Y = 0\}$
 and $\{X = 0, Y = 0\} \text{if } X = 0 \text{ then } Y \leftarrow 1 \{X = 0, Y = 1\}$
 but not $\{X = 0, Y = 0\} X \leftarrow 1 \parallel \text{if } X = 0 \text{ then } Y \leftarrow 1 \{false\}$

\implies we need a side-condition to the rule:

$\{P_1\} s_1 \{Q_1\}$ and $\{P_2\} s_2 \{Q_2\}$ **must not interfere**

Owicki–Gries proof method (cont.)

interference freedom

given proofs Δ_1 and Δ_2 of $\{P_1\} s_1 \{Q_1\}$ and $\{P_2\} s_2 \{Q_2\}$

Δ_1 does not interfere with Δ_2 if:

for any Φ appearing before a statement in Δ_1

for any $\{P'_2\} s'_2 \{Q'_2\}$ appearing in Δ_2

$\{\Phi \wedge P'_2\} s'_2 \{\Phi\}$ holds

and moreover $\{Q_1 \wedge P'_2\} s'_2 \{Q_1\}$

i.e.: the assertions used to prove $\{P_1\} s_1 \{Q_1\}$ are stable by s_2

e.g., $\{X = 0, Y \in [0, 1]\} X \leftarrow 1 \{X = 1, Y \in [0, 1]\}$

$\{X \in [0, 1], Y = 0\}$ if $X = 0$ then $Y \leftarrow 1 \{X \in [0, 1], Y \in [0, 1]\}$

$\implies \{X = 0, Y = 0\} X \leftarrow 1 \parallel \text{if } X = 0 \text{ then } Y \leftarrow 1 \{X = 1, Y \in [0, 1]\}$

Summary:

- pros: the invariants are local to threads
- cons: the proof is **not compositional**

(proving one thread requires delving into the proof of other threads)

\implies not satisfactory

Jones' rely-guarantee proof method

Idea: **explicit interferences** with (more) annotations [Jone81].

Rely-guarantee “quintuples”: $R, G \vdash \{P\} \textit{stat} \{Q\}$

- if P is true before \textit{stat} is executed
- **and the effect of other threads is included in R** (rely)
- then Q is true after \textit{stat}
- **and the effect of \textit{stat} is included in G** (guarantee)

where:

- P and Q are assertions on states (in $\mathcal{P}(\Sigma)$)
- R and G are assertions on transitions (in $\mathcal{P}(\Sigma \times \mathcal{A} \times \Sigma)$)

The parallel composition rule becomes:

$$\frac{R \vee G_2, G_1 \vdash \{P_1\} s_1 \{Q_1\} \quad R \vee G_1, G_2 \vdash \{P_2\} s_2 \{Q_2\}}{R, G_1 \vee G_2 \vdash \{P_1 \wedge P_2\} s_1 \parallel s_2 \{Q_1 \wedge Q_2\}}$$

Rely-guarantee example

Example: proving $0 \leq x \leq y \leq 10$

checking t_1

```

 $\ell^1$  while 0 = 0 do
   $\ell^2$  if  $x < y$  then
     $\ell^3$   $x \leftarrow x + 1$ 
  
```

at ℓ^1, ℓ^2 : $0 \leq x \leq y \leq 10$

at ℓ^3 : $0 \leq x < y \leq 10$

checking t_2

```

 $x \leq y$ 
  |
   $\ell^4$  while 0 = 0 do
     $\ell^5$  if  $y < 10$  then
       $\ell^6$   $y \leftarrow y + 1$ 
  
```

at ℓ^4, ℓ^5 : $0 \leq x \leq y \leq 10$

at ℓ^6 : $0 \leq x \leq y < 10$

Rely-guarantee example

Example: proving $0 \leq x \leq y \leq 10$

checking t_1

ℓ^1 while $0 = 0$ do	x unchanged
ℓ^2 if $x < y$ then	y incremented
ℓ^3 $x \leftarrow x + 1$	$y \leq 10$

at ℓ^1, ℓ^2 : $0 \leq x \leq y \leq 10$

at ℓ^3 : $0 \leq x < y \leq 10$

checking t_2

y unchanged	ℓ^4 while $0 = 0$ do
$x \leq y$	ℓ^5 if $y < 10$ then
	ℓ^6 $y \leftarrow y + 1$

at ℓ^4, ℓ^5 : $0 \leq x \leq y \leq 10$

at ℓ^6 : $0 \leq x \leq y < 10$

In this example:

- guarantee exactly what is relied on ($R_1 = G_1$ and $R_2 = G_2$)
- rely and guarantee are global assertions

Benefits of rely-guarantee:

- invariants are still local to threads
- checking a thread does not require looking at other threads, only at an **abstraction of their semantics**

Auxiliary variables

Example

t_1	t_2
$\ell_1 \ x \leftarrow x + 1 \ \ell_2$	$\ell_3 \ x \leftarrow x + 1 \ \ell_4$

Goal: prove $\{x = 0\} t_1 \parallel t_2 \{x = 2\}$.

Auxiliary variables

Example

t_1	t_2
$\ell 1 \quad x \leftarrow x + 1 \quad \ell 2$	$\ell 3 \quad x \leftarrow x + 1 \quad \ell 4$

Goal: prove $\{x = 0\} t_1 \parallel t_2 \{x = 2\}$.

we must rely on and guarantee that
each thread increments x exactly once!

Solution: auxiliary variables

do not change the semantics but store extra information:

- past values of variables (history of the computation)
- program counter of other threads (pc_t)

Example: for t_1 : $\{(pc_2 = \ell 3 \wedge x = 0) \vee (pc_2 = \ell 4 \wedge x = 1)\}$
 $x \leftarrow x + 1$
 $\{(pc_2 = \ell 3 \wedge x = 1) \vee (pc_2 = \ell 4 \wedge x = 2)\}$

Rely-guarantee as abstract interpretation

Local invariants

State projection: on a thread $t \in \mathbb{T}$

- add auxiliary variables $\mathbb{V}_t \stackrel{\text{def}}{=} \mathbb{V} \cup \{pc_u \mid u \in \mathbb{T}, u \neq t\}$
- enriched environments for t : $\mathcal{E}_t \stackrel{\text{def}}{=} \mathbb{V}_t \rightarrow \mathbb{R}$
(for simplicity, pc_u are numeric variables, i.e., $\mathcal{L} \subseteq \mathbb{R}$)
- local states: $\Sigma_t \stackrel{\text{def}}{=} (\mathcal{L} \times \mathcal{E}_t) \cup \Omega$
(recall that $\Sigma \stackrel{\text{def}}{=} ((\mathbb{T} \rightarrow \mathcal{L}) \times \mathcal{E}) \cup \Omega$)
- **projection**: $\pi_t(L, \rho) \stackrel{\text{def}}{=} (L(t), \rho[\forall u \neq t: pc_u \mapsto L(u)])$
extended naturally to $\pi_t: \mathcal{P}(\Sigma) \rightarrow \mathcal{P}(\Sigma_t)$

Local invariants on t : $RI(t) \stackrel{\text{def}}{=} \pi_t(\mathcal{R})$

(where \mathcal{R} is the reachable state abstraction)

Note: π_t is a bijection, no information is lost

Interferences

Interference: caused by a thread $t \in \mathbb{T}$

$$A \in \mathbb{T} \rightarrow \mathcal{P}(\Sigma \times \Sigma)$$

$$A(t) \stackrel{\text{def}}{=} \alpha^{itf}(\mathcal{T}_p)(t)$$

$$\text{where } \alpha^{itf}(X)(t) \stackrel{\text{def}}{=} \{(\sigma, \sigma') \mid \exists \dots \sigma \xrightarrow{t} \sigma' \dots \in X\}$$

subset of the transition system τ

- spawned by t and
- actually observed in some execution trace
(recall that \mathcal{T}_p is the prefix trace abstraction)

Fixpoint form

Local state fixpoint:

- we express $\mathcal{R}l(t)$ as a function of A and thread $t \in \mathbb{T}$:

$\mathcal{R}l(t) = \text{lfp } R_t(A)$ where

$R_t : (\mathbb{T} \rightarrow \mathcal{P}(\Sigma \times \Sigma)) \rightarrow \mathcal{P}(\Sigma_t) \rightarrow \mathcal{P}(\Sigma_t)$

$R_t(Y)(X) \stackrel{\text{def}}{=} \pi_t(\mathcal{I}) \cup$

$\{ \pi_t(\sigma') \mid \exists \pi_t(\sigma) \in X : \sigma \xrightarrow{t}_\tau \sigma' \vee \exists u \neq t : (\sigma, \sigma') \in Y(u) \}$

A state is reachable if it is initial,
or reachable by transitions from t or from the environment A .

R_t only looks into the syntax of thread t .

R_t is parameterized by the interferences from other threads Y .

Fixpoint form (cont.)

Interferences:

- we express $A(t)$ as a function of $\mathcal{R}I$ and thread $t \in \mathbb{T}$:

$A(t) = B(\mathcal{R}I)(t)$ where

$$B : (\prod_{t \in \mathbb{T}} \{t\} \rightarrow \mathcal{P}(\Sigma_t)) \rightarrow \mathbb{T} \rightarrow \mathcal{P}(\Sigma \times \Sigma)$$

$$B(Z)(t) \stackrel{\text{def}}{=} \{(\sigma, \sigma') \mid \pi_t(\sigma) \in Z(t) \wedge \sigma \xrightarrow{t}_\tau \sigma'\}$$

Collect transitions starting from reachable states.

No fixpoint needed.

Fixpoint form (cont.)

Nested fixpoint characterization:

- 1 $\mathcal{R}I(t) = \text{lfp } R_t(A)$
- 2 $A(t) = B(\mathcal{R}I)(t)$
- 3 mutual dependency between $\mathcal{R}I$ and A

Fixpoint form (cont.)

Nested fixpoint characterization:

- ① $\mathcal{R}I(t) = \text{lfp } R_t(A)$
- ② $A(t) = B(\mathcal{R}I)(t)$
- ③ mutual dependency between $\mathcal{R}I$ and A
 \implies solved using a **fixpoint**:

$\mathcal{R}I = \text{lfp } H$ where

$$H : (\prod_{t \in \mathbb{T}} \{t\} \rightarrow \mathcal{P}(\Sigma_t)) \rightarrow (\prod_{t \in \mathbb{T}} \{t\} \rightarrow \mathcal{P}(\Sigma_t))$$

$$H(Z)(t) \stackrel{\text{def}}{=} \text{lfp } R_t(B(Z))$$

Fixpoint form (cont.)

Constructive fixpoint form:

Use Kleene's iteration to construct fixpoints:

- $\mathcal{R}I = \text{lfp } H = \bigsqcup_{n \in \mathbb{N}} H^n(\lambda t. \emptyset)$
in the pointwise powerset lattice $\prod_{t \in \mathbb{T}} \{t\} \rightarrow \mathcal{P}(\Sigma_t)$
- $H(Z)(t) = \text{lfp } R_t(B(Z)) = \bigcup_{n \in \mathbb{N}} (R_t(B(Z)))^n(\emptyset)$
in the powerset lattice $\mathcal{P}(\Sigma_t)$
(similar to the sequential semantics of thread t in isolation)

\implies nested iterations

Abstract rely-guarantee

Suggested algorithm: nested iterations with acceleration

once abstract domains for states and interferences are chosen

- start from $\mathcal{R}I_0^\# \stackrel{\text{def}}{=} A_0^\# \stackrel{\text{def}}{=} \lambda t. \perp^\#$
- while $A_n^\#$ is not stable
 - compute $\forall t \in \mathbb{T}: \mathcal{R}I_{n+1}^\#(t) \stackrel{\text{def}}{=} \text{lfp } R_t^\#(A_n^\#)$
by iteration with widening ∇
(\simeq separate analysis of each thread)
 - compute $A_{n+1}^\# \stackrel{\text{def}}{=} A_n^\# \nabla B^\#(\mathcal{R}I_{n+1}^\#)$
- when $A_n^\# = A_{n+1}^\#$, return $\mathcal{R}I_n^\#$

\implies thread-modular analysis
parameterized by abstract domains
able to easily reuse existing sequential analyses

Flow-insensitive abstraction

Idea:

- reduce as much control information as possible
- but keep flow-sensitivity on each thread's control location

Local state abstraction: remove **auxiliary** variables

$$\alpha_{\mathcal{R}}^{nf} : \mathcal{P}(\Sigma_t) \rightarrow \mathcal{P}((\mathcal{L} \times \mathcal{E}) \cup \Omega)$$

$$\alpha_{\mathcal{R}}^{nf}(X) \stackrel{\text{def}}{=} \{ (l, \rho|_{\mathcal{V}}) \mid (l, \rho) \in X \} \cup (X \cap \Omega)$$

Interference abstraction: remove **all** control state

$$\alpha_A^{nf} : \mathcal{P}(\Sigma \times \Sigma) \rightarrow \mathcal{P}(\mathcal{E} \times \mathcal{E})$$

$$\alpha_A^{nf}(Y) \stackrel{\text{def}}{=} \{ (\rho, \rho') \mid \exists L, L' \in \mathbb{T} \rightarrow \mathcal{L} : ((L, \rho), (L', \rho')) \in Y \}$$

Flow-insensitive abstraction (cont.)

Flow-insensitive fixpoint semantics: (omitting errors Ω)

We apply $\alpha_{\mathcal{R}}^{nf}$ and α_A^{nf} to the nested fixpoint semantics.

$\mathcal{R}^{nf} \stackrel{\text{def}}{=} \text{lfp } \lambda Z. \lambda t. \text{lfp } R^{nf}_t(B^{nf}(Z))$, where

$B^{nf}(Z)(t) \stackrel{\text{def}}{=} \{(\rho, \rho') \mid \exists \ell, \ell' \in \mathcal{L}: (\ell, \rho) \in Z(t) \wedge (\ell, \rho) \rightarrow_t (\ell', \rho')\}$

$R^{nf}_t(Y)(X) \stackrel{\text{def}}{=} R_t^{loc}(X) \cup A_t^{nf}(Y)(X)$

$R_t^{loc}(X) \stackrel{\text{def}}{=} \{(\ell_t^i, \lambda V.0)\} \cup \{(\ell', \rho') \mid \exists (\ell, \rho) \in X: (\ell, \rho) \rightarrow_t (\ell', \rho')\}$

$A_t^{nf}(Y)(X) \stackrel{\text{def}}{=} \{(\ell, \rho') \mid \exists \rho, u \neq t: (\ell, \rho) \in X \wedge (\rho, \rho') \in Y(u)\}$

where \rightarrow_t is the transition relation for thread t alone: $\tau[\text{stat}_t]$

Cost/precision trade-off:

- less variables
 - \implies subsequent numeric abstractions are more efficient
- sufficient to analyze our first example (p. 34)
- insufficient to analyze $x \leftarrow x + 1 \parallel x \leftarrow x + 1$

Non-relational interference abstraction

Idea: simplify further flow-insensitive interferences

- numeric relations are more costly than numeric sets
 \implies remove input sensitivity
- relational domains are more costly than non-relational
 \implies abstract the interference on each variable separately

Non-relational interference abstraction:

$$\alpha_A^{nr} : \mathcal{P}(\mathcal{E} \times \mathcal{E}) \rightarrow (\mathbb{V} \rightarrow \mathcal{P}(\mathbb{R}))$$

$$\alpha_A^{nr}(Y) \stackrel{\text{def}}{=} \lambda V. \{x \in \mathbb{V} \mid \exists (\rho, \rho') \in Y : \rho(V) \neq x \wedge \rho'(V) = x\}$$

(remember which variables are modified and their new values)

To apply interferences, we get, in the nested fixpoint form:

$$A_t^{nr}(Y)(X) \stackrel{\text{def}}{=} \{(\ell, \rho[V \mapsto v]) \mid (\ell, \rho) \in X, V \in \mathbb{V}, \exists u \neq t : v \in Y(u)(V)\}$$

A note on unbounded threads

Extension: relax the finiteness constraint on \mathbb{T}

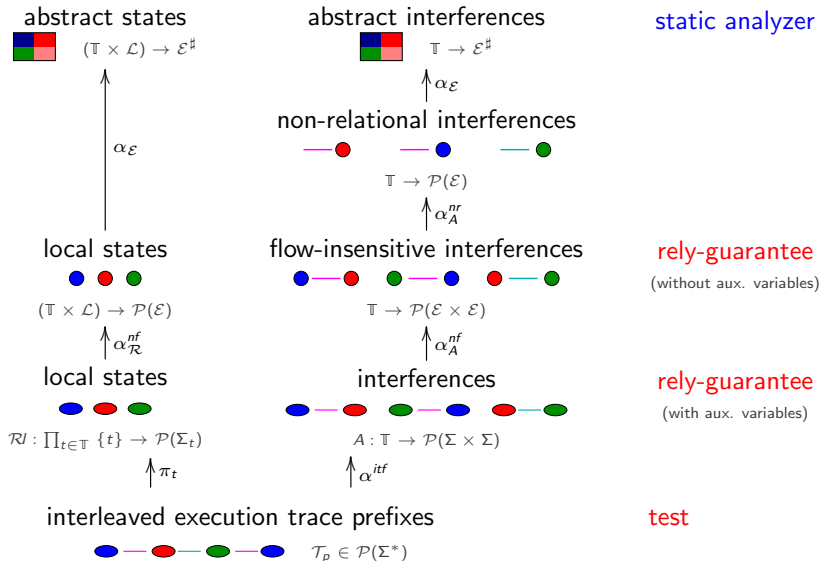
- there is still a **finite syntactic set** of threads \mathbb{T}_s
- some threads $\mathbb{T}_\infty \subseteq \mathbb{T}_s$ can have several instances
(possibly an unbounded number)

Flow-insensitive analysis:

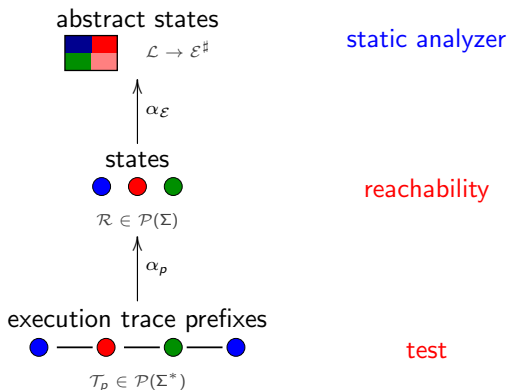
- local state and interference domains have finite dimensions
(\mathcal{E}_t and $(\mathcal{L} \times \mathcal{E}) \times (\mathcal{L} \times \mathcal{E})$, as opposed to \mathcal{E} and $\mathcal{E} \times \mathcal{E}$)
- all instances of a thread $t \in \mathbb{T}_s$ are isomorphic
 \implies iterate the analysis on the finite set \mathbb{T}_s (instead of \mathbb{T})
- we must handle **self-interferences** for threads in \mathbb{T}_∞ :

$$A_t^{nf}(Y)(X) \stackrel{\text{def}}{=} \{ (\ell, \rho') \mid \exists \rho, u: (u \neq t \vee t \in \mathbb{T}_\infty) \wedge (\ell, \rho) \in X \wedge (\rho, \rho') \in Y(u) \}$$

From traces to thread-modular analyses



Compare with sequential analyses. . .



Construction of an interference-based analysis

Reminder: sequential analysis in denotational form

Expression semantics: $E[\text{expr}] : \mathcal{E} \rightarrow (\mathcal{P}(\mathbb{R}) \times \mathcal{P}(\Omega))$

$$E[X] \rho \stackrel{\text{def}}{=} \langle \{\rho(X)\}, \emptyset \rangle$$

$$E[[c_1, c_2]] \rho \stackrel{\text{def}}{=} \langle \{x \in \mathbb{R} \mid c_1 \leq x \leq c_2\}, \emptyset \rangle$$

$$E[-e_1] \rho \stackrel{\text{def}}{=} \text{let } \langle V_1, O_1 \rangle = E[e_1] \rho \text{ in } \langle \{-v_1 \mid v_1 \in V_1\}, O_1 \rangle$$

$$E[e_1 \diamond_{\omega} e_2] \rho \stackrel{\text{def}}{=} \text{let } \forall i \in \{1, 2\}: \langle V_i, O_i \rangle = E[e_i] \rho \text{ in} \\ \langle \{v_1 \diamond v_2 \mid v_i \in V_i, \diamond \neq / \vee v_2 \neq 0\}, O_1 \cup O_2 \cup \{\omega \text{ if } \diamond = / \wedge 0 \in V_2\} \rangle$$

Statement semantics: $C[\text{stat}] : (\mathcal{P}(\mathcal{E}) \times \mathcal{P}(\Omega)) \rightarrow (\mathcal{P}(\mathcal{E}) \times \mathcal{P}(\Omega))$

$$C[X \leftarrow e] \langle R, O \rangle \stackrel{\text{def}}{=} \langle \emptyset, O \rangle \sqcup \bigsqcup_{\rho \in R} \langle \{\rho[X \mapsto v] \mid v \in V_{\rho}\}, O_{\rho} \rangle$$

$$C[e \bowtie 0?] \langle R, O \rangle \stackrel{\text{def}}{=} \langle \emptyset, O \rangle \sqcup \bigsqcup_{\rho \in R} \langle \{\rho \mid \exists v \in V_{\rho}: v \bowtie 0\}, O_{\rho} \rangle$$

$$C[\text{if } e \bowtie 0 \text{ then } s] X \stackrel{\text{def}}{=} (C[s] \circ C[e \bowtie 0?]) X \sqcup C[e \bowtie 0?] X$$

$$C[\text{while } e \bowtie 0 \text{ do } s] X \stackrel{\text{def}}{=} \\ C[e \bowtie 0?] (\text{lfp } \lambda Y. X \sqcup (C[s] \circ C[e \bowtie 0?]) Y)$$

$$C[s_1; s_2] \stackrel{\text{def}}{=} C[s_2] \circ C[s_1]$$

$$\text{where } \langle V_{\rho}, O_{\rho} \rangle \stackrel{\text{def}}{=} E[e] \rho$$

Denotational semantics with interferences

Interferences in $\mathbb{I} \stackrel{\text{def}}{=} \mathbb{T} \times \mathbb{V} \times \mathbb{R}$

$\langle t, X, v \rangle$ means: t can store the value v into the variable X

We define the analysis of a thread t
with respect to a set of interferences $I \subseteq \mathbb{I}$.

Expressions with interference: for thread t

$E_t \llbracket expr \rrbracket : (\mathcal{E} \times \mathcal{P}(\mathbb{I})) \rightarrow (\mathcal{P}(\mathbb{R}) \times \mathcal{P}(\Omega))$

- Apply interferences to read variables:

$$E_t \llbracket X \rrbracket \langle \rho, I \rangle \stackrel{\text{def}}{=} \langle \{ \rho(X) \} \cup \{ v \mid \exists u \neq t: \langle u, X, v \rangle \in I \}, \emptyset \rangle$$

- Pass recursively I down to sub-expressions:

$$E_t \llbracket -e_1 \rrbracket \langle \rho, I \rangle \stackrel{\text{def}}{=} \text{let } \langle V_1, O_1 \rangle = E_t \llbracket e_1 \rrbracket \langle \rho, I \rangle \text{ in } \langle \{ -v_1 \mid v_1 \in V_1 \}, O_1 \rangle$$

...

Denotational semantics with interferences (cont.)

Statements with interference: for thread t

$$C_t[\text{stat}] : (\mathcal{P}(\mathcal{E}) \times \mathcal{P}(\Omega) \times \mathcal{P}(\mathbb{I})) \rightarrow (\mathcal{P}(\mathcal{E}) \times \mathcal{P}(\Omega) \times \mathcal{P}(\mathbb{I}))$$

- pass interferences to expressions
- collect new interferences due to assignments
- accumulate interferences from inner statements

$$C_t[X \leftarrow e] \langle R, O, I \rangle \stackrel{\text{def}}{=} \langle \emptyset, O, I \rangle \sqcup \bigsqcup_{\rho \in R} \langle \{ \rho[X \mapsto v] \mid v \in V_\rho \}, O_\rho, \{ \langle t, X, v \rangle \mid v \in V_\rho \} \rangle$$

$$C_t[s_1; s_2] \stackrel{\text{def}}{=} C_t[s_2] \circ C_t[s_1]$$

...

$$\text{(noting } \langle V_\rho, O_\rho \rangle \stackrel{\text{def}}{=} E_t[e] \langle \rho, I \rangle \text{)}$$

(\sqcup is now the element-wise \cup in $\mathcal{P}(\mathcal{E}) \times \mathcal{P}(\Omega) \times \mathcal{P}(\mathbb{I})$)

Denotational semantics with interferences (cont.)

Program semantics: $P[[prog]] \subseteq \Omega$

Given $prog ::= stat_1 \parallel \dots \parallel stat_n$, we compute:

$$P[[prog]] \stackrel{\text{def}}{=} \left[\text{Ifp } \lambda \langle O, I \rangle. \bigsqcup_{t \in \mathbb{T}} [C_t[[stat_t]] \langle \mathcal{E}_0, \emptyset, I \rangle]_{\Omega, \perp} \right]_{\Omega}$$

- each thread analysis starts in an initial environment set $\mathcal{E}_0 \stackrel{\text{def}}{=} \{ \lambda V.0 \}$
- $[X]_{\Omega, \perp}$ projects $X \in \mathcal{P}(\mathcal{E}) \times \mathcal{P}(\Omega) \times \mathcal{P}(\perp)$ on $\mathcal{P}(\Omega) \times \mathcal{P}(\perp)$ and interferences and errors from all threads are joined (the output environments are ignored)
- $P[[prog]]$ only outputs the set of possible run-time errors

Example

Example	
t_1	t_2
while ℓ_1 $0 = 0$ do ℓ_2 if $x < y$ then ℓ_3 $x \leftarrow x + 1$	while ℓ_4 $0 = 0$ do ℓ_5 if $y < 10$ then ℓ_6 $y \leftarrow y + 1$

Concrete interference semantics:

iteration 1

$I = \emptyset$

ℓ_1 : $x = 0, y = 0$

ℓ_4 : $x = 0, y \in [0, 10]$

new $I = \{ \langle t_2, y, 1 \rangle, \dots, \langle t_2, y, 10 \rangle \}$

Example

Example	
t_1	t_2
while ℓ_1 $0 = 0$ do ℓ_2 if $x < y$ then ℓ_3 $x \leftarrow x + 1$	while ℓ_4 $0 = 0$ do ℓ_5 if $y < 10$ then ℓ_6 $y \leftarrow y + 1$

Concrete interference semantics:

iteration 2

$$I = \{ \langle t_2, y, 1 \rangle, \dots, \langle t_2, y, 10 \rangle \}$$

$$\ell_1 : x \in [0, 10], y = 0$$

$$\ell_4 : x = 0, y \in [0, 10]$$

$$\text{new } I = \{ \langle t_1, x, 1 \rangle, \dots, \langle t_1, x, 10 \rangle, \langle t_2, y, 1 \rangle, \dots, \langle t_2, y, 10 \rangle \}$$

Example

Example	
t_1	t_2
while ℓ^1 $0 = 0$ do ℓ^2 if $x < y$ then ℓ^3 $x \leftarrow x + 1$	while ℓ^4 $0 = 0$ do ℓ^5 if $y < 10$ then ℓ^6 $y \leftarrow y + 1$

Concrete interference semantics:

iteration 3

$$I = \{ \langle t_1, x, 1 \rangle, \dots, \langle t_1, x, 10 \rangle, \langle t_2, y, 1 \rangle, \dots, \langle t_2, y, 10 \rangle \}$$

$$\ell^1 : x \in [0, 10], y = 0$$

$$\ell^4 : x = 0, y \in [0, 10]$$

$$\text{new } I = \{ \langle t_1, x, 1 \rangle, \dots, \langle t_1, x, 10 \rangle, \langle t_2, y, 1 \rangle, \dots, \langle t_2, y, 10 \rangle \}$$

Example

Example	
t_1	t_2
while ℓ_1 $0 = 0$ do ℓ_2 if $x < y$ then ℓ_3 $x \leftarrow x + 1$	while ℓ_4 $0 = 0$ do ℓ_5 if $y < 10$ then ℓ_6 $y \leftarrow y + 1$

Concrete interference semantics:

iteration 3

$$I = \{ \langle t_1, x, 1 \rangle, \dots, \langle t_1, x, 10 \rangle, \langle t_2, y, 1 \rangle, \dots, \langle t_2, y, 10 \rangle \}$$

$$\ell_1 : x \in [0, 10], y = 0$$

$$\ell_4 : x = 0, y \in [0, 10]$$

$$\text{new } I = \{ \langle t_1, x, 1 \rangle, \dots, \langle t_1, x, 10 \rangle, \langle t_2, y, 1 \rangle, \dots, \langle t_2, y, 10 \rangle \}$$

Note: we don't get that $x \leq y$ at ℓ_1 , only that $x, y \in [0, 10]$

Interference abstraction

Abstract interferences \mathbb{I}^\sharp

$\mathcal{P}(\mathbb{I}) \stackrel{\text{def}}{=} \mathcal{P}(\mathbb{T} \times \mathbb{V} \times \mathbb{R})$ is abstracted as $\mathbb{I}^\sharp \stackrel{\text{def}}{=} (\mathbb{T} \times \mathbb{V}) \rightarrow \mathcal{R}^\sharp$
 where \mathcal{R}^\sharp abstracts $\mathcal{P}(\mathbb{R})$ (e.g. intervals)

Abstract semantics with interferences $C_t^\sharp \llbracket s \rrbracket$

derived from $C^\sharp \llbracket s \rrbracket$ in a generic way:

Example: $C_t^\sharp \llbracket X \leftarrow e \rrbracket \langle R^\sharp, \Omega, I^\sharp \rangle$

- for each Y in e , get its interference $Y_{\mathcal{R}}^\sharp = \bigsqcup_{\mathcal{R}}^\sharp \{ I^\sharp \langle u, Y \rangle \mid u \neq t \}$
- if $Y_{\mathcal{R}}^\sharp \neq \perp_{\mathcal{R}}^\sharp$, replace Y in e with $get \langle Y, R^\sharp \rangle \sqcup_{\mathcal{R}}^\sharp Y_{\mathcal{R}}^\sharp$
 (where $get \langle Y, R^\sharp \rangle$ extracts the abstract values in \mathcal{R}^\sharp of a variable Y from $R^\sharp \in \mathcal{E}^\sharp$)
- compute $\langle R^{\sharp'}, O' \rangle = C^\sharp \llbracket e \rrbracket \langle R^\sharp, O \rangle$
- enrich $I^\sharp \langle t, X \rangle$ with $get \langle X, R^{\sharp'} \rangle$

Static analysis with interferences

Abstract analysis

$$P^\# \llbracket prog \rrbracket \stackrel{\text{def}}{=} \left[\lim \lambda \langle O, I^\# \rangle. \langle O, I^\# \rangle \nabla \bigsqcup_{t \in \mathbb{T}}^\# \left[C_t^\# \llbracket stat_t \rrbracket \langle \mathcal{E}_0^\#, \emptyset, I^\# \rangle \right]_{\Omega, I^\#} \right]_\Omega$$

- effective analysis by structural induction
- termination ensured by a widening
- parametrized by a choice of abstract domains $\mathcal{R}^\#, \mathcal{E}^\#$
- interferences are flow-insensitive and non-relational in $\mathcal{R}^\#$
- thread analysis remains flow-sensitive and relational in $\mathcal{E}^\#$

[Miné12]

Path-based semantics

Control paths

$atomic ::= X \leftarrow expr \mid expr \bowtie 0?$

Control paths

$\pi : stat \rightarrow \mathcal{P}(atomic^*)$

$\pi(X \leftarrow e) \stackrel{\text{def}}{=} \{X \leftarrow e\}$

$\pi(\text{if } e \bowtie 0 \text{ then } s) \stackrel{\text{def}}{=} (\{e \bowtie 0?\} \cdot \pi(s)) \cup \{e \not\bowtie 0?\}$

$\pi(\text{while } e \bowtie 0 \text{ do } s) \stackrel{\text{def}}{=} \left(\bigcup_{i \geq 0} (\{e \bowtie 0?\} \cdot \pi(s))^i \right) \cdot \{e \not\bowtie 0?\}$

$\pi(s_1; s_2) \stackrel{\text{def}}{=} \pi(s_1) \cdot \pi(s_2)$

$\pi(stat)$ is a (generally infinite) set of finite control paths

Path-based concrete semantics of sequential programs

Join-over-all-path semantics

$$\sqcap \llbracket P \rrbracket : (\mathcal{P}(\mathcal{E}) \times \mathcal{P}(\Omega)) \rightarrow (\mathcal{P}(\mathcal{E}) \times \mathcal{P}(\Omega)) \quad P \subseteq \text{atomic}^*$$

$$\sqcap \llbracket P \rrbracket \langle R, O \rangle \stackrel{\text{def}}{=} \bigsqcup_{s_1 \dots s_n \in P} (C \llbracket s_n \rrbracket \circ \dots \circ C \llbracket s_1 \rrbracket) \langle R, O \rangle$$

Semantic equivalence

$$C \llbracket \text{stat} \rrbracket = \sqcap \llbracket \pi(\text{stat}) \rrbracket$$

(not true in the abstract)

Advantages:

- easily extended to concurrent programs (path interleavings)
- able to model program transformations (weak memory models)

Path-based concrete semantics of concurrent programs

Concurrent control paths

$$\begin{aligned} \pi_* &\stackrel{\text{def}}{=} \{ \text{interleavings of } \pi(\text{stat}_t), t \in \mathbb{T} \} \\ &= \{ p \in \text{atomic}^* \mid \forall t \in \mathbb{T}, \text{proj}_t(p) \in \pi(\text{stat}_t) \} \end{aligned}$$

Interleaving program semantics

$$P_*[[\text{prog}]] \stackrel{\text{def}}{=} [\sqcap[\pi_*]] \langle \mathcal{E}_0, \emptyset \rangle]_\Omega$$

$(\text{proj}_t(p))$ keeps only the atomic statement in p coming from thread t)

Soundness of the interference semantics

Soundness theorem

$$P_*[[prog]] \subseteq P[[prog]]$$

Proof sketch:

- define $\sqcap_t[[P]]X \stackrel{\text{def}}{=} \bigsqcup \{ C_t[[s_1; \dots; s_n]]X \mid s_1 \cdot \dots \cdot s_n \in P \}$,
then $\sqcap_t[[\pi(s)]] = C_t[[s]]$;
- given the interference fixpoint $I \subseteq \mathbb{I}$ from $P[[prog]]$,
prove by recurrence on the length of $p \in \pi_*$ that:
 - $\forall t \in \mathbb{T}, \forall \rho \in [\sqcap[[p]]\langle \mathcal{E}_0, \emptyset \rangle]_{\mathcal{E}}$,
 $\exists \rho' \in [\sqcap_t[[proj_t(p)]]\langle \mathcal{E}_0, \emptyset, I \rangle]_{\mathcal{E}}$ such that
 $\forall X \in \mathbb{V}, \rho(X) = \rho'(X)$ or $\langle u, X, \rho(X) \rangle \in I$ for some $u \neq t$.
 - $[\sqcap[[p]]\langle \mathcal{E}_0, \emptyset \rangle]_{\Omega} \subseteq \bigcup_{t \in \mathbb{T}} [\sqcap_t[[proj_t(p)]]\langle \mathcal{E}_0, \emptyset, I \rangle]_{\Omega}$

Note: sound but not complete

Weakly consistent memories

Issues with weak consistency

program written

$F_1 \leftarrow 1;$	$F_2 \leftarrow 1;$
if $F_2 = 0$ then	if $F_1 = 0$ then
S_1	S_2

(simplified Dekker mutual exclusion algorithm)

S_1 and S_2 **cannot** execute simultaneously.

Issues with weak consistency

program written

$F_1 \leftarrow 1;$	$F_2 \leftarrow 1;$
if $F_2 = 0$ then	if $F_1 = 0$ then
S_1	S_2

→

program executed

if $F_2 = 0$ then	if $F_1 = 0$ then
$F_1 \leftarrow 1;$	$F_2 \leftarrow 1;$
S_1	S_2

(simplified Dekker mutual exclusion algorithm)

S_1 and S_2 can execute simultaneously.

Not a sequentially consistent behavior!

Caused by:

- write FIFOs, caches, distributed memory
- hardware or compiler optimizations, transformations
- ...

behavior accepted by Java [Mans05]

Out of thin air principle

original program

$$\begin{array}{l|l} R_1 \leftarrow X; & R_2 \leftarrow Y; \\ Y \leftarrow R_1 & X \leftarrow R_2 \end{array}$$

(example from causality test case #4 for Java by Pugh et al.)

We should not have $R_1 = 42$.

Out of thin air principle

original program

$$R_1 \leftarrow X; \quad R_2 \leftarrow Y;$$

$$Y \leftarrow R_1 \quad | \quad X \leftarrow R_2$$
 \longrightarrow

“optimized” program

$$Y \leftarrow 42; \quad R_2 \leftarrow Y;$$

$$R_1 \leftarrow X; \quad | \quad X \leftarrow R_2$$

$$Y \leftarrow R_1 \quad | \quad X \leftarrow R_2$$

(example from causality test case #4 for Java by Pugh et al.)

We should not have $R_1 = 42$.

Possible if we allow speculative writes!

\implies we **disallow** this kind of program transformations.

(also forbidden in Java)

Atomicity and granularity

original program

$X \leftarrow X + 1 \mid X \leftarrow X + 1$

We assumed that assignments are atomic...

Atomicity and granularity

original program

$$X \leftarrow X + 1 \mid X \leftarrow X + 1$$


executed program

$$\begin{array}{l|l} r_1 \leftarrow X + 1 & r_2 \leftarrow X + 1 \\ X \leftarrow r_1 & X \leftarrow r_2 \end{array}$$

We assumed that assignments are atomic...
but that may not be the case

The second program admits more behaviors
e.g.: $X = 1$ at the end of the program

[Reyn04]

Path-based definition of weak consistency

Acceptable control path transformations: $p \rightsquigarrow q$

only reduce interferences and errors

- **Reordering:** $X_1 \leftarrow e_1 \cdot X_2 \leftarrow e_2 \rightsquigarrow X_2 \leftarrow e_2 \cdot X_1 \leftarrow e_1$
(if $X_1 \notin \text{var}(e_2)$, $X_2 \notin \text{var}(e_1)$, and e_1 does not stop the program)
- **Propagation:** $X \leftarrow e \cdot s \rightsquigarrow X \leftarrow e \cdot s[e/X]$
(if $X \notin \text{var}(e)$, $\text{var}(e)$ are thread-local, and e is deterministic)
- **Factorization:** $s_1 \cdot \dots \cdot s_n \rightsquigarrow X \leftarrow e \cdot s_1[X/e] \cdot \dots \cdot s_n[X/e]$
(if X is fresh, $\forall i, \text{var}(e) \cap \text{lval}(s_i) = \emptyset$, and e has no error)
- **Decomposition:** $X \leftarrow e_1 + e_2 \rightsquigarrow T \leftarrow e_1 \cdot X \leftarrow T + e_2$
(change of granularity)
- ...

but **NOT:**

- “out-of-thin-air” writes: $X \leftarrow e \rightsquigarrow X \leftarrow 42 \cdot X \leftarrow e$

Soundness of the interference semantics

Interleaving semantics of transformed programs $P'_*[[prog]]$

- $\pi'(s) \stackrel{\text{def}}{=} \{p \mid \exists p' \in \pi(s): p' \rightsquigarrow^* p\}$
- $\pi'_* \stackrel{\text{def}}{=} \{\text{interleavings of } \pi'(stat_t), t \in \mathbb{T}\}$
- $P'_*[[prog]] \stackrel{\text{def}}{=} [\sqcap[\pi'_*]]\langle \mathcal{E}_0, \emptyset \rangle_\Omega$

Soundness theorem

$$P'_*[[prog]] \subseteq P[[prog]]$$

\implies the interference semantics is sound
wrt. weakly consistent memories and changes of granularity

Synchronisation

Scheduling

Synchronization primitives


```
stat ::= lock(m)  
      | unlock(m)
```

$m \in \mathbb{M}$: finite set of non-recursive mutexes

Scheduling

- mutexes ensure **mutual exclusion**
a each time, each mutex can be locked by a single thread
- mutexes enforce **memory consistency** and atomicity
no optimization across **lock** and **unlock** instructions
memory caches and buffer are flushed

Mutual exclusion



Interleaving semantics $P_*[[prog]]$:

restrict interleavings of control paths

Interference semantics $P[[prog]]$, $P^\sharp[[prog]]$:


partition wrt. an abstract local view of the scheduler \mathbb{C}

$$\bullet \mathcal{E} \rightsquigarrow \mathcal{E} \times \mathbb{C}, \quad \mathcal{E}^\sharp \rightsquigarrow \mathbb{C} \rightarrow \mathcal{E}^\sharp$$

$$\bullet \mathbb{I} \stackrel{\text{def}}{=} \mathbb{T} \times \mathbb{V} \times \mathbb{R} \rightsquigarrow \mathbb{I} \stackrel{\text{def}}{=} \mathbb{T} \times \mathbb{C} \times \mathbb{V} \times \mathbb{R},$$

$$\mathbb{I}^\sharp \stackrel{\text{def}}{=} (\mathbb{T} \times \mathbb{V}) \rightarrow \mathcal{R}^\sharp \rightsquigarrow \mathbb{I}^\sharp \stackrel{\text{def}}{=} (\mathbb{T} \times \mathbb{C} \times \mathbb{V}) \rightarrow \mathcal{R}^\sharp$$

Mutual exclusion




Data-race effects

Partition wrt. **mutexes** $M \subseteq \mathbb{M}$ held by the current thread t

- $C_t[[X \leftarrow e]] \langle \rho, M, I \rangle$ adds $\{ \langle t, M, X, v \rangle \mid v \in E_t[[X]] \langle \rho, M, I \rangle \}$ to I
- $E_t[[X]] \langle \rho, M, I \rangle = \{ \rho(X) \} \cup \{ v \mid \langle t', M', X, v \rangle \in I, t \neq t', M \cap M' = \emptyset \}$
- flow-insensitive, subject to weak memory consistency

Mutual exclusion



Well-synchronized effects

- last write before unlock affects first read after lock
- partition interferences wrt. a protecting mutex m (and M)
- $C_t[\text{unlock}(m)] \langle \rho, M, I \rangle$ stores $\rho(X)$ into I
- $C_t[\text{lock}(m)] \langle \rho, M, I \rangle$ imports values form I into ρ
- **imprecision**: non-relational, largely flow-insensitive

Example analysis

abstract consumer/producer

t_1	t_2
<pre> while 0=0 do lock(m);^{ℓ1} if X > 0 then ^{ℓ2} X ← X - 1; unlock(m); ^{ℓ3} Y ← X </pre>	<pre> while 0=0 do lock(m); X ← X + 1; if X > 10 then X ← 10; unlock(m) </pre>

- at $\ell 1$, the **unlock** – **lock** effect from t_2 imports $\{X\} \times [1, 10]$
- at $\ell 2$, $X \in [1, 10]$, no effect from t_2 : $X \leftarrow X - 1$ is safe
- at $\ell 3$, $X \in [0, 9]$, and t_2 has the effects $\{X\} \times [1, 10]$
so, $Y \in [0, 10]$

Limitations of the interference abstraction

Lack of relational lock invariants

a difficult example

$$\mathcal{E}_0 : X = Y = 5$$

while 1 do

lock(m);

if $X > 0$ then

$X \leftarrow X - 1;$

$Y \leftarrow Y - 1;$

unlock(m)

while 1 do

lock(m);

if $X < 10$ then

$X \leftarrow X + 1;$

$Y \leftarrow Y + 1;$

unlock(m)

Our analysis finds $X \in [0, 10]$, but **no bound** on Y .

Actually $Y \in [0, 10]$.

To prove this, we would need to infer the **relational invariant** $X = Y$ at lock boundaries.

Lack of inter-process flow-sensitivity

a more difficult example

<pre> while 1 do lock(m); $X \leftarrow X + 1;$ unlock(m); lock(m); $X \leftarrow X - 1;$ unlock(m) </pre>	<pre> while 1 do lock(m); $X \leftarrow X + 1;$ unlock(m); lock(m); $X \leftarrow X - 1;$ unlock(m) </pre>
---	---

Our analysis finds **no bound** on X .

Actually $X \in [-2, 2]$ at all program points.

To prove this we need to infer an **invariant on the history of interleaved executions**:

no more than two incrementation (resp. decrementation) can occur without a decrementation (resp. incrementation).

Summary

Conclusion

We presented a static analysis that is:

- inspired from **thread-modular** proof methods
- **sound** for all **interleavings**
- **sound** for **weakly consistent memory** semantics
- aware of **scheduling** and **synchronization**
- **parametrized** by abstract domains

Future work: leverage the connection with rely-guarantee

- **relational interferences**
(especially for synchronized program parts)
- **flow-sensitive** interferences and invariants

Bibliography

Bibliography

[Bour93] **F. Bourdoncle**. *Efficient chaotic iteration strategies with widenings*. In Proc. FMPA'93, LNCS vol. 735, pp. 128–141, Springer, 1993.

[Carr09] **J.-L. Carré & C. Hymans**. *From single-thread to multithreaded: An efficient static analysis algorithm*. In arXiv:0910.5833v1, EADS, 2009.

[Cous84] **P. Cousot & R. Cousot**. *Invariance proof methods and analysis techniques for parallel programs*. In Automatic Program Construction Techniques, chap. 12, pp. 243–271, Macmillan, 1984.

[Cous85] **R. Cousot**. *Fondements des méthodes de preuve d'invariance et de fatalité de programmes parallèles*. In Thèse d'Etat es sc. math., INP Lorraine, Nancy, 1985.

[Hoar69] **C. A. R. Hoare**. *An axiomatic basis for computer programming*. In Com. ACM, 12(10):576–580, 1969.

Bibliography (cont.)

[Jone81] **C. B. Jones**. *Development methods for computer programs including a notion of interference*. In PhD thesis, Oxford University, 1981.

[Lamp77] **L. Lamport**. *Proving the correctness of multiprocess programs*. In IEEE Trans. on Software Engineering, 3(2):125–143, 1977.

[Lamp78] **L. Lamport**. *Time, clocks, and the ordering of events in a distributed system*. In Comm. ACM, 21(7):558–565, 1978.

[Mans05] **J. Manson, B. Pugh & S. V. Adve**. *The Java memory model*. In Proc. POPL'05, pp. 378–391, ACM, 2005.

[Miné12] **A. Miné**. *Static analysis of run-time errors in embedded real-time parallel C programs*. In LMCS 8(1:26), 63 p., arXiv, 2012.

[Owic76] **S. Owicki & D. Gries**. *An axiomatic proof technique for parallel programs I*. In Acta Informatica, 6(4):319–340, 1976.

Bibliography (cont.)

[Reyn04] **J. C. Reynolds.** *Toward a grainless semantics for shared-variable concurrency.* In Proc. FSTTCS'04, LNCS vol. 3328, pp. 35–48, Springer, 2004.

[Sara07] **V. A. Saraswat, R. Jagadeesan, M. M. Michael & C. von Praun.** *A theory of memory models.* In Proc. PPOPP'07, pp. 161–172, ACM, 2007.