

An Observational Theory of Imperative Concurrent Data Structures in the π -Calculus

Luca Fossati Kohei Honda

Electronics, Engineering and Computer Science
Queen Mary University of London

Lisboa, 19/04/2011

Intro

Traditional global progress properties of concurrent programs :

- *Deadlock-Freedom*
- *Starvation-Freedom*

Intro

Traditional global progress properties of concurrent programs :

- *Deadlock-Freedom*
- *Starvation-Freedom*

⇓ Critical section

→ Lock-based only

Intro

A more general approach :

- *Non-Blockingness*
- *Wait-Freedom*

Intro

A more general approach :

- *Non-Blockingness*

- *Wait-Freedom*

↑ Abstraction

↑ Extensionality

Intro

A more general approach :

- *Non-Blockingness*

- *Wait-Freedom*

↑ Abstraction

↑ Extensionality

↓ Lack of rigorous semantic basis

Intro

Non-Blockingness

“A data structure is *non-blocking* if it guarantees that *some* process will always be able to complete its pending operation in a finite number of its own steps, regardless of the execution speed of other processes.” [Taubenfeld, '06]

Intro

How to formalise:

- “a finite number of its own steps”
- “regardless of the execution speed of other processes”

?

Intro

How to formalise:

- “a finite number of its own steps”
- “regardless of the execution speed of other processes”

\Rightarrow *Fairness*

\Rightarrow *Partial Failures*

Intro

How to formalise:

- “a finite number of its own steps”
- “regardless of the execution speed of other processes”

\Rightarrow *Fairness*

\Rightarrow *Partial Failures*

\Rightarrow *π -calculus*

Intro

AIM :

To provide an **extensional** theory which is **general** enough to cover all the concurrent data structures whose behaviours are representable in the π -calculus.

Index

- π -calculus
- Linear/Affine Types
- Asynchronous fair LTS + partial failures
- Global Progress :
 - Non-blockingness
 - Wait-Freedom
- Case study : Queues
 - Correctness (state space)
 - Behavioural Classification

The Calculus

$$\begin{array}{l} P ::= u \&_{i \in I} \{ \mathbf{1}_i(\vec{x}_i).P_i \} \\ | \bar{u} \oplus \mathbf{1} \langle \vec{e} \rangle \\ | \text{if } e \text{ then } P \text{ else } Q \\ | P|Q \\ | (\nu u)P \\ | (\mu X(\vec{x}).P) \langle \vec{e} \rangle \\ | X \langle \vec{x} \rangle \\ | \mathbf{0} \end{array}$$

Reductions

One rule:

$$u \&_{i \in I} \{ \mathbf{1}_i(\vec{x}_i).P_i \} \mid \bar{u} \oplus \mathbf{1}_j \langle \vec{e} \rangle \longrightarrow P_j \{ \vec{e} / \vec{x}_j \} \quad (j \in I)$$

Closed under the standard structural congruence, \equiv .

Note in particular:

$$(\mu X(\vec{x}).P) \langle \vec{e} \rangle \equiv P \{ (\mu X(\vec{x}).P) / X \} \{ \vec{e} / \vec{x} \}$$

Example

Some simple concurrent data structures :

$$\text{Ref}\langle u, v \rangle \stackrel{\text{def}}{=} u \& \left\{ \begin{array}{l} \text{read}(z) : \bar{z}\langle v \rangle \mid \text{Ref}\langle u, v \rangle, \\ \text{write}(y, z) : \bar{z} \mid \text{Ref}\langle u, y \rangle \end{array} \right\}$$

$$\text{Ref}^{\text{cas}}\langle u, v \rangle \stackrel{\text{def}}{=} u \& \left\{ \begin{array}{l} \text{read}(z) : \bar{z}\langle v \rangle \mid \text{Ref}^{\text{cas}}\langle u, v \rangle, \\ \text{write}(y, z) : \bar{z} \mid \text{Ref}^{\text{cas}}\langle u, y \rangle, \\ \text{cas}(x, y, z) : \text{if } x = v \text{ then } \bar{z}\langle \text{tt} \rangle \mid \text{Ref}^{\text{cas}}\langle u, y \rangle \\ \quad \quad \quad \text{else } \bar{z}\langle \text{ff} \rangle \mid \text{Ref}^{\text{cas}}\langle u, v \rangle \end{array} \right\}$$

Example

Reduction :

$$\begin{aligned}
 & \text{Ref}^{\text{cas}}\langle a, 0 \rangle | (\nu c)(\bar{a} \oplus \text{cas}\langle 0, 1, c \rangle | c(x).P) \\
 \longrightarrow & (\nu c)((\text{if } 0 = 0 \text{ then } \bar{c}\langle \text{tt} \rangle | \text{Ref}^{\text{cas}}\langle a, 1 \rangle \text{ else } \bar{c}\langle \text{ff} \rangle | \text{Ref}^{\text{cas}}\langle a, 0 \rangle) | c(x).P) \\
 \longrightarrow & (\nu c)((\text{if } \text{tt} \text{ then } \bar{c}\langle \text{tt} \rangle | \text{Ref}^{\text{cas}}\langle a, 1 \rangle \text{ else } \bar{c}\langle \text{ff} \rangle | \text{Ref}^{\text{cas}}\langle a, 0 \rangle) | c(x).P) \\
 \longrightarrow^* & \text{Ref}^{\text{cas}}\langle a, 1 \rangle | P\{\text{tt}/x\}
 \end{aligned}$$

Example

Two different mutex agents:

$$\begin{aligned} \text{Mtx}\langle u \rangle &\stackrel{\text{def}}{=} u(x).\bar{x}(h)h.\text{Mtx}\langle u \rangle \\ \text{Mtx}^{\text{spin}}\langle u \rangle &\stackrel{\text{def}}{=} (\nu c)(!u(x).\mu X. \\ &\quad (\text{if cas}(c, 0, 1) \text{ then } \bar{x}(h)h.\text{CAS}(c, 1, 0) \text{ else }) \mid \\ &\quad \text{Ref}^{\text{cas}}\langle c, 0 \rangle) \end{aligned}$$

where

$$\text{if cas}(u, v, w) \text{ then } P \text{ else } Q \stackrel{\text{def}}{=} (\nu c)(\bar{u} \oplus \text{cas}\langle v, w, c \rangle \mid c(x).\text{if } x \text{ then } P \text{ else } Q)$$

and

$$\text{CAS}(u, v, w) \stackrel{\text{def}}{=} \text{if cas}(u, v, w) \text{ then } \mathbf{0} \text{ else } \mathbf{0}$$

Types

$$\tau ::= \&_{i \in I}^M l_i(\vec{\tau}_i) \mid \oplus_{i \in I}^M l_i(\vec{\tau}_i) \mid \text{int} \mid \text{bool} \mid \perp$$

Modalities (as in *Linear Logic*, *Games*, ...):

L channel can be used “exactly once” (*linear*)

A channel can be used “at most once” (*affine*)

L* input end always available and shared by unboundedly many outputs (*unbounded l.*)

A* input end as above but may be unavailable (*unbounded a.*)

Example

Typings for the previously introduced examples :

1.

$$u : \&^{L*} \{ read(\uparrow^L(\text{nat})), write(\text{nat } \uparrow^L()) \} \vdash \text{Ref}\langle u, 3 \rangle$$

2.

$$u : \&^{L*} \{ read(\uparrow^L(\text{nat})), write(\text{nat } \uparrow^L()), cas(\text{natnat } \uparrow^L(\text{bool})), \} \vdash \text{Ref}^{\text{cas}}\langle u, 0 \rangle$$

3.

$$u : \downarrow^{A*}(\uparrow^A(\downarrow^A())) \vdash P \quad (P \in \{ \text{Mtx}\langle u \rangle, \text{Mtx}^{\text{spin}}\langle u \rangle \})$$

Labelled Transition System

Labels :

$$\ell ::= \tau \mid (\nu \vec{c})a\&l(\vec{v}) \mid (\nu \vec{c})a \oplus l\langle \vec{v} \rangle$$

Untyped transitions :

(Bra)

$$P \xrightarrow{(\nu \vec{c})a\&l\langle \vec{v} \rangle} P|\bar{a} \oplus \mathbf{1}\langle \vec{v} \rangle$$

(Sel)

$$(\nu \vec{c})(P|\bar{a} \oplus \mathbf{1}\langle \vec{v} \rangle) \xrightarrow{(\nu \vec{c})\bar{a} \oplus l\langle \vec{v} \rangle} P$$

Labelled Transition System

Environment transitions :

$$\begin{aligned}
 & \Gamma, a : \&^{L^*, A^*} \{l_i(\vec{\tau}_i)\}_{i \in I} \xrightarrow{(\nu \vec{c})a \& l_j \langle \vec{v}_j \rangle} \Gamma \odot \vec{v} : \vec{\tau}_j, a : \&^{L^*, A^*} \{l_i(\vec{\tau}_i)\}_{i \in I} \\
 & (\Gamma \odot \vec{v} : \vec{\tau}_j) / \vec{c}, a : \oplus^{L^*, A^*} \{l_i(\vec{\tau}_i)\}_{i \in I} \xrightarrow{(\nu \vec{c})a \oplus l_j \langle \vec{v}_j \rangle} \Gamma, a : \oplus^{L^*, A^*} \{l_i(\vec{\tau}_i)\}_{i \in I} \\
 & \Gamma, a : \&^{L, A} \{l_i(\vec{\tau}_i)\}_{i \in I} \xrightarrow{(\nu \vec{c})a \& l_j \langle \vec{v}_j \rangle} \Gamma \odot \vec{v} : \vec{\tau}_j, a : \perp \\
 & (\Gamma \odot \vec{v} : \vec{\tau}_j) / \vec{c}, a : \oplus^{L, A} \{l_i(\vec{\tau}_i)\}_{i \in I} \xrightarrow{(\nu \vec{c})a \oplus l_j \langle \vec{v}_j \rangle} \Gamma
 \end{aligned}$$

Typed transitions :

$$\Gamma \vdash P \xrightarrow{\ell} \Gamma' \vdash P' \quad \stackrel{\text{def}}{\Leftrightarrow} \quad P \xrightarrow{\ell} P' \wedge \Gamma \xrightarrow{\ell} \Gamma'$$

Bisimilarity

Definition 3.3 (bisimilarity) A typed relation \mathcal{R} is a *weak bisimulation* or often *bisimulation* when, for each $\Gamma \vdash PRQ$, we have: $P \xrightarrow{\ell} P'$ implies $Q \xRightarrow{\hat{\ell}} Q'$ s.t. $P'\mathcal{R}Q'$, and the symmetric case. The maximum bisimulation is written \approx .

Proposition 3.4 \approx is a typed congruence.

Fairness

Definition 3.5 (Fairness) A maximal transition sequence Φ from closed $\Gamma \vdash P$ is *fair* if no subject is infinitely often enabled in Φ .

Fairness

Definition 3.5 (Fairness) A maximal transition sequence Φ from closed $\Gamma \vdash P$ is *fair* if no subject is infinitely often enabled in Φ .

Let $P = !a.(\bar{b}|\bar{a})|\bar{a}$ and $Q = \text{Ref}\langle r, 3 \rangle|\bar{r} \oplus \text{read}\langle c \rangle$.

Then $P|Q$ admits an infinite unfair transition sequence.

Fairness

Definition 3.5 (Fairness) A maximal transition sequence Φ from closed $\Gamma \vdash P$ is *fair* if no subject is infinitely often enabled in Φ .

Let $P = !a.(\bar{b}|\bar{a})|\bar{a}$ and $Q = \text{Ref}\langle r, 3 \rangle|\bar{r} \oplus \text{read}\langle c \rangle$.

Then $P|Q$ admits an infinite unfair transition sequence.

Fairness induces a *fair pre-order* \lesssim_{fair} , where:

$$P \lesssim_{\text{fair}} Q \iff P \approx Q \wedge \text{WFT}(P) \supseteq \text{WFT}(Q)$$

Partial Failure

We first augment the dynamics with *failing reductions*:

$$\bar{u} \oplus^M \mathbf{1}_j \langle \vec{e} \rangle \longrightarrow \mathbf{0} \quad (M \neq L) \quad \text{if } v \text{ then } P \text{ else } Q \longrightarrow \mathbf{0}$$

Then we augment the τ -transition accordingly.

Only affine outputs may fail!

NOTE :

⇓

Linearity \Rightarrow *Atomicity*

Global Progress

Definition 3.11 (Resilience) Let $\Gamma \vdash P$ such that for all $\text{wft}(\Phi) \in \text{WFT}(P)$, the set $\text{blocked}(\Phi)$ is finite. Then we say that $\Gamma \vdash P$ is *resilient*.

Definition 3.12 (NB/WF) A closed process $\Gamma \vdash P$ is:

1. *non-blocking* (NB) when it is resilient and, for any $\Phi \in \text{FT}(P)$ s.t. $\Delta \vdash Q$ is in Φ and $\text{allowed}(\Delta) \setminus \text{blocked}(\Phi) \neq \emptyset$, some affine output occurs in Φ .
2. *wait-free* (WF) when it is resilient and, for any $\Phi \in \text{FT}(P)$ s.t. $\Delta \vdash Q$ is in Φ and $c \in \text{allowed}(\Delta) \setminus \text{blocked}(\Phi)$, an output at c occurs in Φ .

Weak Global Progress

Let **WNB** be as **NB** but without failures.

Let **WWF** be as **WF** but without failures.

Then **WF** \subsetneq **NB** \cap **WWF** and **NB** \cup **WWF** \subsetneq **WNB**.

Abstract Queue Specification

Abstract Queue:

$$AQ(r, \langle Rs; Vs; As \rangle)$$

Abstract Queue Specification

Abstract Queue:

$$AQ(r, \langle Rs; Vs; As \rangle)$$

Example:

$$AQ(r, \langle \{ \text{enq}(6, g_1), \text{deq}(g_2) \}, 2 \cdot 3 \cdot 1, \{ \overline{g_3} \langle 5 \rangle \} \rangle)$$

State Space Abstraction

Abstract Queue:

$$\text{AQ}(r, \langle \text{Rs}; \text{Vs}; \text{As} \rangle)$$

State Transitions

$$\text{AQ}(r, \langle \text{Rs}, \text{Vs}, \text{As} \rangle) \xrightarrow{r \& \text{enq}(v, g)} \text{AQ}(r, \langle \text{Rs} \uplus \text{enq}(v, g), \text{Vs}, \text{As} \rangle)$$

$$\text{AQ}(r, \langle \text{Rs}, \text{Vs}, \text{As} \rangle) \xrightarrow{r \& \text{deq}(g)} \text{AQ}(r, \langle \text{Rs} \uplus \text{deq}(g), \text{Vs}, \text{As} \rangle)$$

$$\text{AQ}(r, \langle \text{Rs} \uplus \text{enq}(v, g), \text{Vs}, \text{As} \rangle) \xrightarrow{\tau} \text{AQ}(r, \langle \text{Rs}, \text{Vs} \cdot v, \text{As} \uplus \bar{g} \rangle)$$

$$\text{AQ}(r, \langle \text{Rs} \uplus \text{deq}(g), v \cdot \text{Vs}, \text{As} \rangle) \xrightarrow{\tau} \text{AQ}(r, \langle \text{Rs}, \text{Vs}, \text{As} \uplus \{\bar{g}\langle v \rangle\} \rangle)$$

$$\text{AQ}(r, \langle \text{Rs} \uplus \text{deq}(g), , \text{As} \rangle) \xrightarrow{\tau} \text{AQ}(r, \langle \text{Rs}, , \text{As} \uplus \{\bar{g}\langle KO \rangle\} \rangle)$$

$$\text{AQ}(r, \langle \text{Rs}, \text{Vs}, \text{As} \uplus \{\bar{g}\langle v \rangle\} \rangle) \xrightarrow{\bar{g}\langle v \rangle} \text{AQ}(r, \langle \text{Rs}, \text{Vs}, \text{As} \rangle)$$

cas-based Queue in π

Auxiliary data structures :

$$\begin{aligned} \text{Node}(r, v, ptr) &\stackrel{\text{def}}{=} \text{Ref}\langle r, \langle v, ptr \rangle \rangle \\ \text{Ptr}(r, nd, ctr) &\stackrel{\text{def}}{=} \text{Ref}^{\text{cas}}\langle r, \langle nd, ctr \rangle \rangle \end{aligned}$$

Shortened forms :

$$x \triangleleft \text{read}(\vec{y}).P \stackrel{\text{def}}{=} (\nu c)(\bar{x} \oplus \text{read}\langle c \rangle | c(\vec{y}).P)$$

We use ad-hoc names when we only want to keep a projection of the contents (*getPtr*, *getVal*, ...).

cas-based Queue in π

$$\text{CQemp}(r) \stackrel{\text{def}}{=} (\nu h)(\nu t)(\nu s)(\text{CQ}(r, h, t) \mid \text{Ptr}(h, s, 0) \mid \text{Ptr}(t, s, 0) \mid \text{Node}(s, 0, \text{null}))$$

$$\text{CQ}(r, h, t) \stackrel{\text{def}}{=} r \& \left\{ \begin{array}{l} \text{enqueue}(x, u) : \\ (\text{CQ}(r, h, t) \mid P_{\text{enq}}(x, t)) \\ \text{dequeue}(u) : \\ (\text{CQ}(r, h, t) \mid P_{\text{deq}}(h, t)) \end{array} \right\}$$

cas-based Queue in π

$P_{enq}(x, tail) =$

NullPtr($nlPtr$) | Node($node, v, nlPtr$) |

$(\mu X_{tag}(u'))$.

$tail \triangleleft read(last, ctrT)$. $last \triangleleft getPtr(tPtr)$. $tPtr \triangleleft read(next, ctr)$.

if ($next = null$) then

if cas($tPtr, \langle next, ctr \rangle, \langle node, ctr + 1 \rangle$) then

if cas($tail, \langle last, ctrT \rangle, \langle node, ctrT + 1 \rangle$) then \bar{u}' else \bar{u}'

else $X_{tag}(u')$)

else

if cas($tail, \langle last, ctrT \rangle, \langle next, ctrT + 1 \rangle$) then $X_{tag}(u')$ else $X_{tag}(u')$)

endif) $\langle u \rangle$

cas-based Queue in π

$$P_{deq}(head, tail) = (\mu X_{tag}(u').$$

$$head \triangleleft read(hNdRef, ctrH). \quad tail \triangleleft read(tNdRef, ctrT).$$

$$hNdRef \triangleleft getPtr(hNdPtr). \quad hNdPtr \triangleleft getNext(next).$$

$$\mathbf{if} (hNdRef = tNdRef) \mathbf{then}$$

$$\quad \mathbf{if} (next = null) \mathbf{then} \overline{u'}\langle null \rangle \mathbf{else}$$

$$\quad \quad \mathbf{if} (\mathbf{cas}(tail, \langle tNdRef, ctrT \rangle, \langle next, ctrT + 1 \rangle)) \mathbf{then} X_{tag}\langle u' \rangle$$

$$\quad \quad \mathbf{else} X_{tag}\langle u' \rangle)$$

$$\mathbf{else} next \triangleleft getVal(x).$$

$$\quad \mathbf{if} (\mathbf{cas}(head, \langle hNdRef, ctrH \rangle, \langle next, ctr + 1 \rangle)) \mathbf{then} \overline{u'}\langle x \rangle$$

$$\quad \mathbf{else} X_{tag}\langle u' \rangle)\langle u \rangle$$

Lock-Based Queue in π

$$\begin{aligned}
 \text{LQemp}(r) &\stackrel{\text{def}}{=} (\nu u)(\nu h)(\nu t)(\nu s)(\text{LQ}(r, h, t) | \text{Mtx}\langle u \rangle | \\
 &\quad \text{LPtr}(h, s) | \text{LPtr}(t, s) | \text{LENode}(s, 0)) \\
 \\
 \text{LQ}(r, h, t, l) &\stackrel{\text{def}}{=} r \& \left\{ \begin{array}{l} \text{enqueue}(v, u) : \text{LQ}(r, h, t, l) | \\ (\bar{l}(g)g(y).P_{enq}^{lck}(v, t, c') | c'.(\bar{y}|\bar{u})), \\ \text{dequeue}(u) : \text{LQ}(r, h, t, l) | \\ (\bar{l}(g)g(y).P_{deq}^{lck}(h, t, c') | c'.(\bar{y}|\bar{u})) \end{array} \right\}
 \end{aligned}$$

Correctness

- **state space abstraction** (*correctness by bisimilarity*)
- **molecular transitions** (\rightsquigarrow *atomic operation*)
- **commit event** (*i.e. cas on successor pointer*)
- **normal form** (*each thread is either:*
 1. *ready to commit,*
 2. *or ready to output*))
- **normalisation through linearisation** (*local permutations*)

Global Progress

Proposition 4.6 Let $\Phi : \text{CQemp}(r) \longrightarrow^* P$ be a queue process. Then an output is blocked in Φ iff its thread fails in Φ .

Proposition 4.8 $\Gamma_Q \vdash \text{CQemp}(r)$ is non-blocking.

Behavioural Classification

$\Gamma_Q \vdash \text{LQemp}(r)$ is blocking (it is weakly wait-free).

$\Gamma_Q \vdash \text{LQemp}(r)$ has more fair sequences

Theorem 4.11 $\text{AQ}(r, \varepsilon) \approx \text{LQemp}(r) \not\approx_{\text{fair}} \text{CQemp}(r)$

Results and Future Works

- Fairness
 - Fine-grained analysis
 - Generality
 - Extensionality
 - State space abstraction
- ⇒ Automated verification tools
- ⇒ Encoding from imperative languages

References

- [1] D. Cacciagrano, F. Corradini and C. Palamidessi. Fair π .
Proc. *EXPRESS'06. ENTCS*, 175:3–26. Elsevier Science, 2007.
- [2] M. M. Michael and M. L. Scott. Simple, fast, and practical
non-blocking and blocking concurrent queue algorithms.
PODC-15. 267–275. ACM, 1996.
- [3] R. Milner and J. G. Parrow and D. J. Walker, A calculus of
Mobile Processes, *Information and Computation* 100(1), 1–77,
1992.
- [4] G. Taubenfeld. *Synchronization Algorithms and Concurrent
Programming*. Pearson–Prentice Hall, 2006.

Thank you for your attention!