

Reductions and Causality (I)



jean-jacques.levy@inria.fr

Tsinghua University,

October 31, 2011

<http://pauillac.inria.fr/~levy/courses/tsinghua/reductions>

Plan

- independent statements
- control flow
- relaxed control flow
- data dependencies
- parallel processes
- functional languages and multi-threading

Weak memory models

CENTRE DE RECHERCHE
COMMUN



INRIA
MICROSOFT RESEARCH

Intel whitepaper (1/3)

2.3 Loads may be reordered with older stores to different locations

Intel 64 memory ordering allows load instructions to be reordered with prior stores to a different location. However, loads are not reordered with prior stores to the same location.

The first example in this section illustrates the case in which a load may be reordered with an older store – i.e. if the store and load are to different non-overlapping locations.

Processor 0	Processor 1
<code>mov [_x], 1 // M1</code>	<code>mov [_y], 1 // M3</code>
<code>mov r1, [_y] // M2</code>	<code>mov r2, [_x] // M4</code>
Initially <code>x == y == 0</code>	

Table 2.3.a: Loads may be reordered with older stores

Demo1

Intel whitepaper (2/3)

2.1 Loads are not reordered with other loads and stores are not reordered with other stores

Intel 64 memory ordering ensures that loads are seen in program order, and that stores are seen in program order.

Processor 0	Processor 1
<code>mov [_x], 1 // M1</code>	<code>mov r1,[_y] // M3</code>
<code>mov [_y], 1 // M2</code>	<code>mov r2, [_x] // M4</code>
Initially <code>x == y == 0</code>	

Table 2.1: Stores are not reordered with other stores

Intel whitepaper (3/3)

2.5 Stores are transitively visible

Intel 64 memory ordering ensures transitive visibility of stores – i.e. stores that are causally related appear to execute in an order consistent with the causal relation.

Processor 0	Processor 1	Processor 2
<code>mov [_x], 1 // M1</code>	<code>mov r1, [_x] // M2</code> <code>mov [_y], 1 // M3</code>	<code>mov r2, [_y] // M4</code> <code>mov r3, [_x] // M5</code>
Initially <code>x == y == 0</code> <code>r1 == 1, r2 == 1, r3 == 0</code> is not allowed		

Table 2.5: Stores are transitively visible

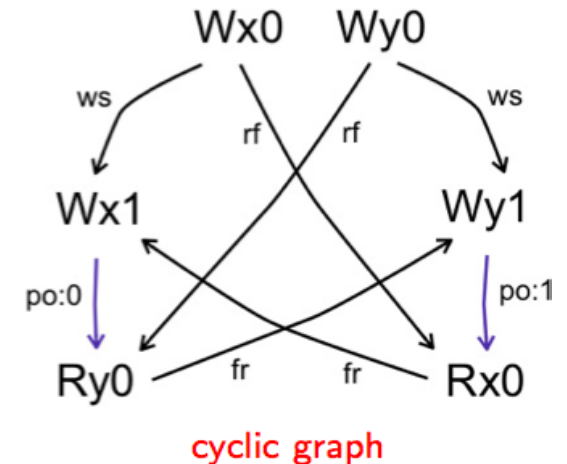
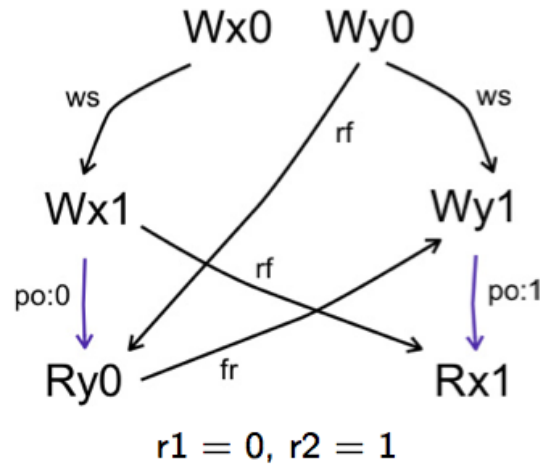
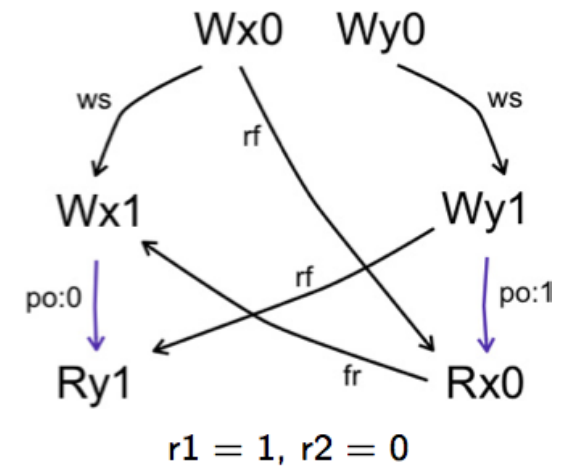
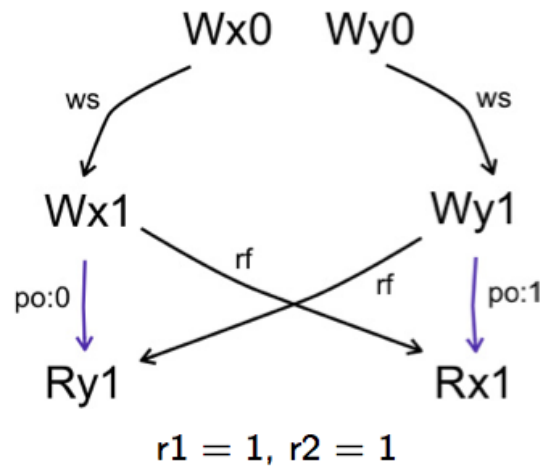
Causality

- Computations steps may be independent
- Others are causally related
- Refined to memory accesses
- Theory of causality ?
 - reductions steps in the λ -calculus (this course)
 - event structures (processes)
 - true concurrency
 - slicing (program analysis)
 - etc.

Intel 64 revisited (1/4)

- In SC (sequentially consistent), program order is strictly respected

P0	P1
mov [x], 1	mov [y], 1
mov r1, [y]	mov r2, [x]



Intel 64 revisited (2/4)

- Dependency relations

$A \xrightarrow{po} B$ is program order

$W \xrightarrow{ws} W$ is write serialization (acyclic) relation

$W \xrightarrow{rf} R$ is read from relation

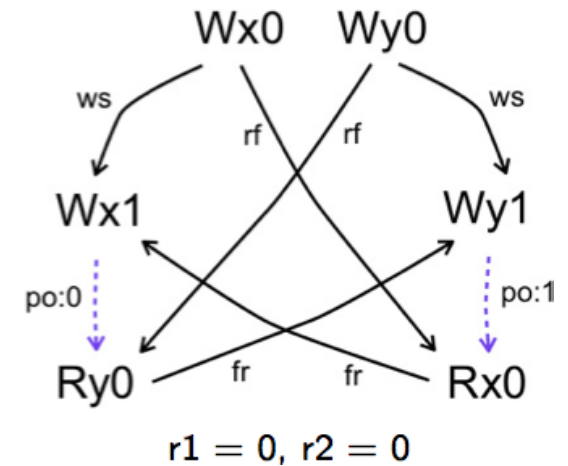
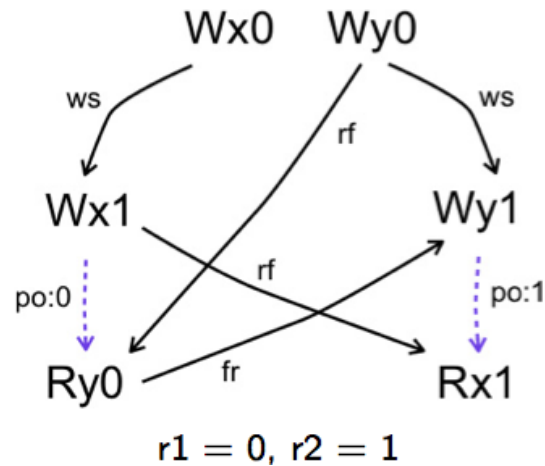
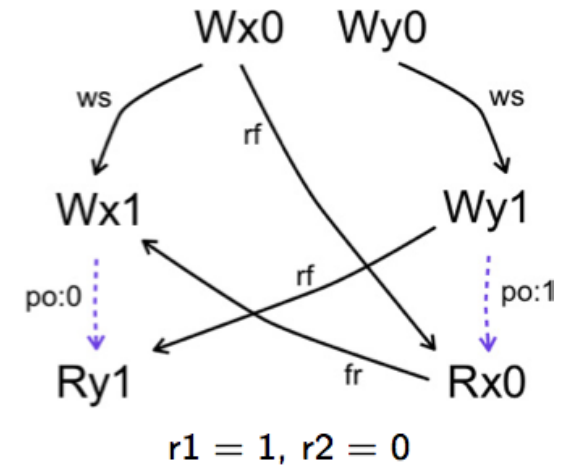
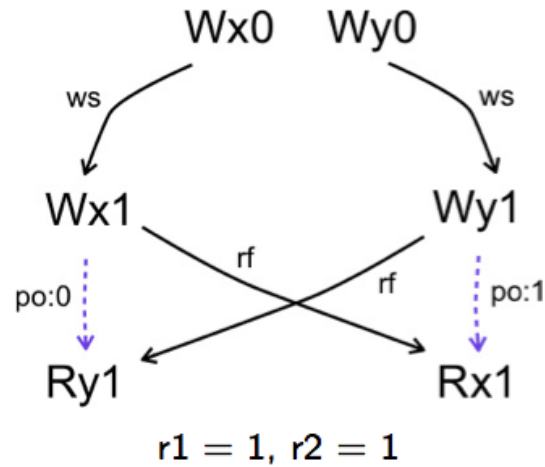
$R \xrightarrow{fr} W$ is from read relation when there is W' such that

$W' \xrightarrow{rf} R$ and $W' \xrightarrow{ws} W$

Intel 64 revisited (3/4)

- In TSO, W followed by R can be relaxed within program order

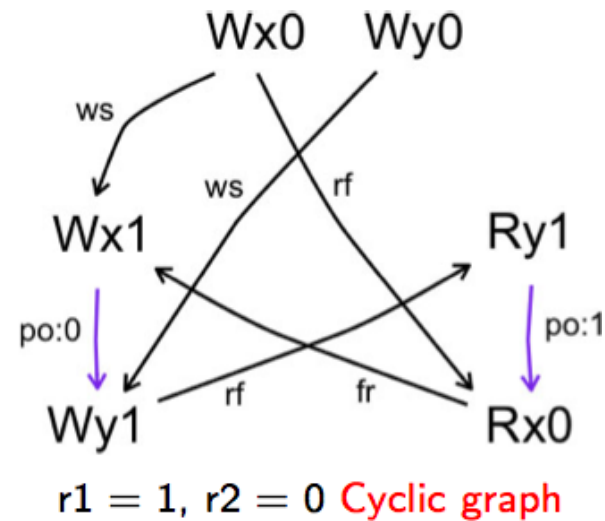
P0	P1
mov [x], 1	mov [y], 1
mov r1, [y]	mov r2, [x]



Intel 64 revisited (4/4)

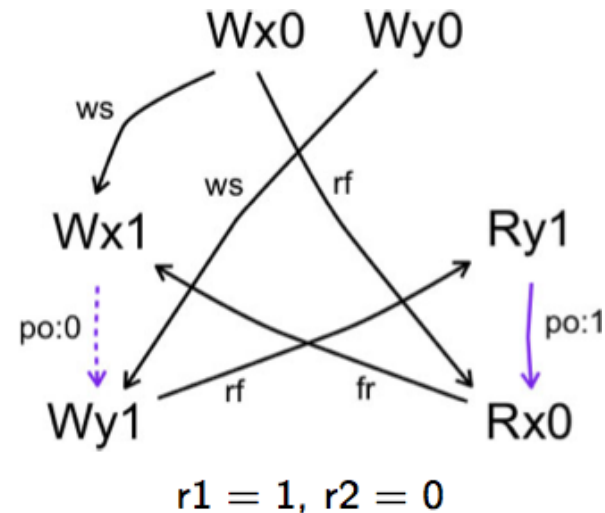
- In TSO, W followed by R is relaxed

P0	P1
mov [x], 1	mov r1, [y]
mov [y], 1	mov r2, [x]



- In PSO, W followed by W to distinct location is relaxed

P0	P1
mov [x], 1	mov r1, [y]
mov [y], 1	mov r2, [x]

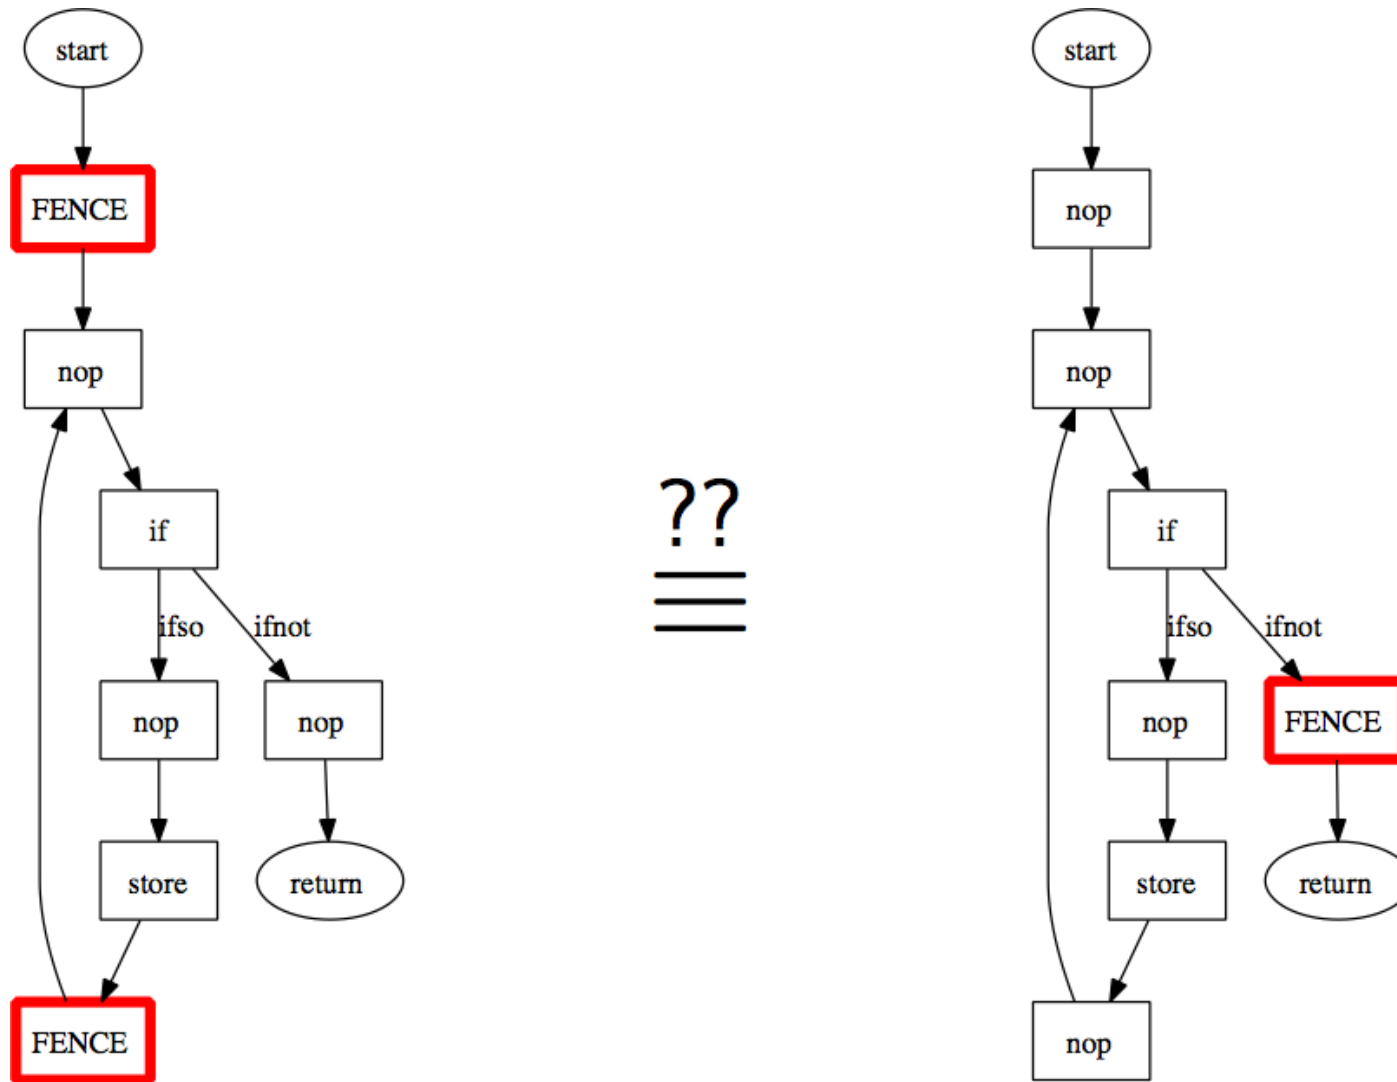


WMM

- axiomatic + operational models for Intel [[~Cambridge](#)] / Power [[~INRIA](#)]
- formalisation in HOL/Coq
- tests on real processor behaviour
<http://www.cl.cam.ac.uk/~pes20/ppc-supplemental/ppc003.html>
- formal proof of simple concurrent code (eg. Linux spinlocks)
- operational reasoning: data-race freedom, separation logic
- certified compiler for concurrent languages
<http://www.cl.cam.ac.uk/~pes20/CompCertTSO>

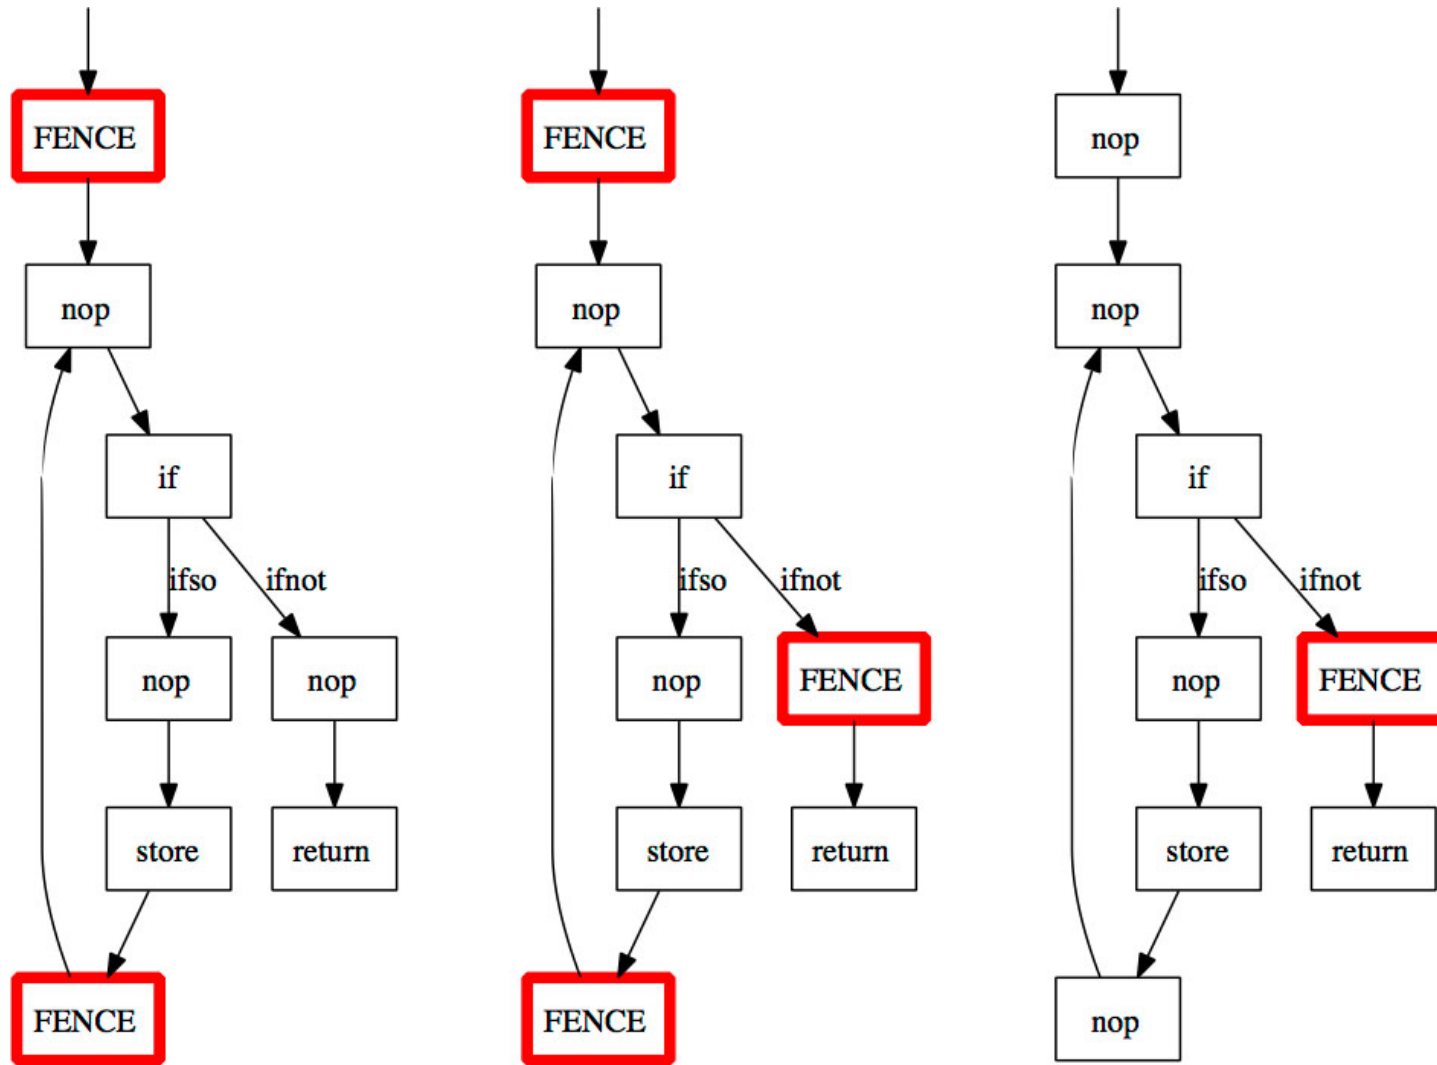
[Zappa Nardelli, Maranget, Alglave, Braibant, Sewell et al]
[POPL 09, CACM 10; DAMP 09, CAV 10, PLDI 11; TACAS 11; POPL 11]

WMM and optimization (1/2)



Fences elimination with TSO \simeq 3 kloCoq

WMM and optimization (2/2)



Parallel computations

CENTRE DE RECHERCHE
COMMUN



INRIA
MICROSOFT RESEARCH

Pure functional languages (1/2)

- Evaluations of subexpressions are independent

$$M M_1 M_2 \cdots M_n$$

- Evaluations of M and M_i can be done in parallel
- No longer true if effects within some of M_i
 - In Haskell, monads are the only effect-ful subterms
 - monads may call pure functional terms
 - pure functional terms cannot call monads
 - effects are visible in types
- In other languages, static analysis necessary to detect effects

Pure functional languages (2/2)

- Inside functional languages, there exists dependencies

$$M N \xrightarrow{\star} (\lambda x . P) Q \longrightarrow P\{x := Q\}$$

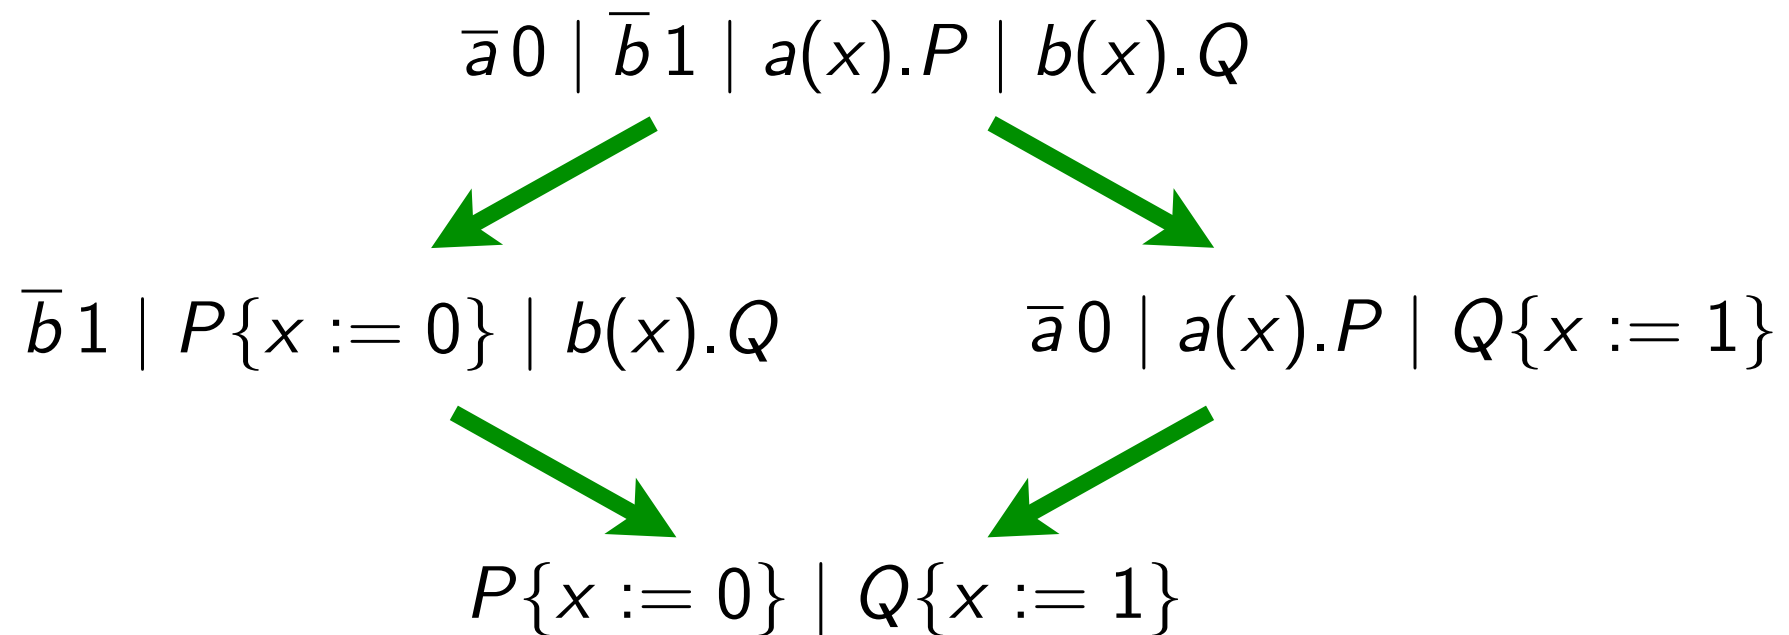
- Evaluations in M has to be done before toplevel redex

$$I I(la) \longrightarrow I(la) \longrightarrow la$$

where $I = \lambda x . x$

CCS or π -calculus (1/3)

- Communications may be independent

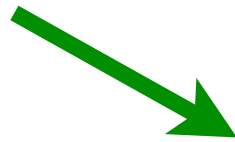
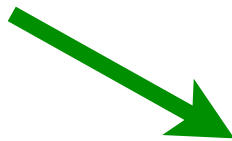


- but other transitions may be causally related

[Jean Krivine] (reversible CCS)

CCS or π -calculus (2/3)

- example of causally related transitions

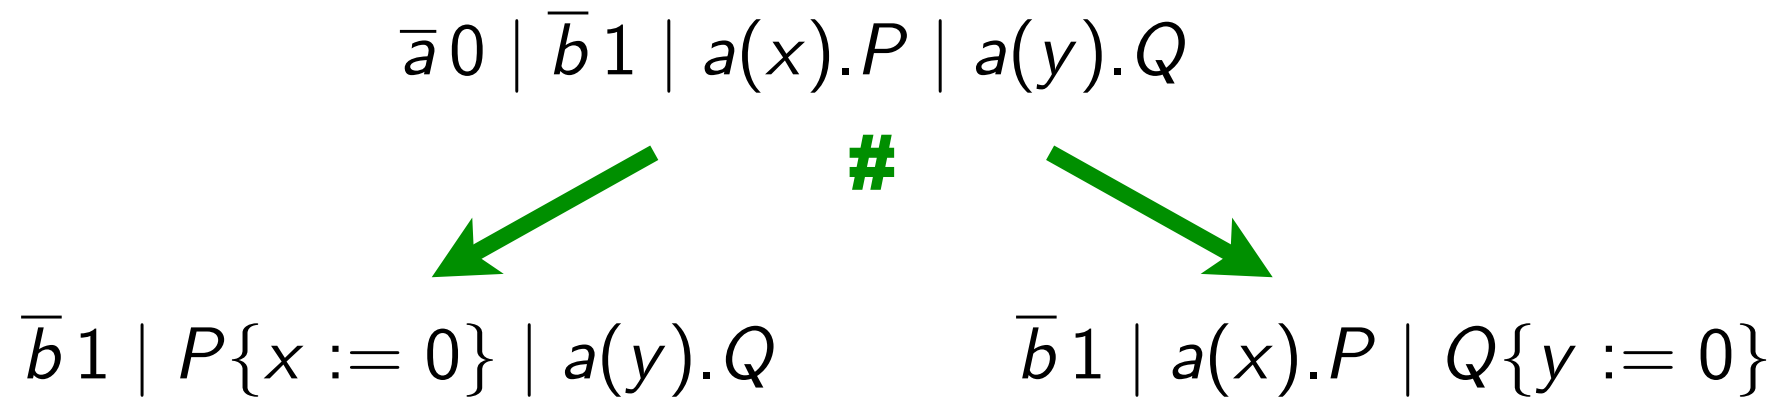
$$\bar{a}0 \mid a(x).b(y).P \mid \bar{b}1$$

$$b(y).P\{x := 0\} \mid \bar{b}1$$

$$P\{x := 0, y := 1\}$$

- causality in process algebras == event structures or Petri nets
- « true concurrency »

[Winskel, Boudol-Castellani]

CCS or π -calculus (3/3)

- possible conflicts



Security

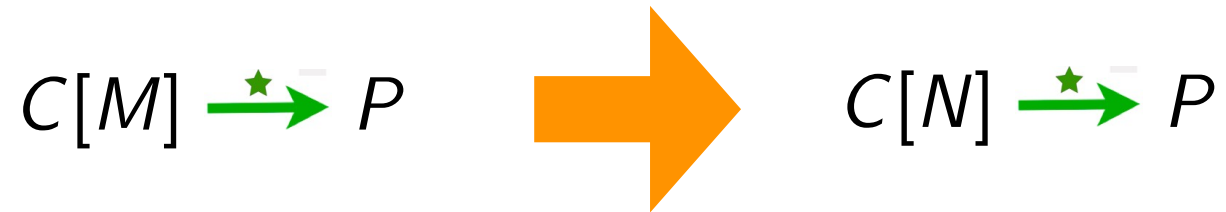
CENTRE DE RECHERCHE
COMMUN



INRIA
MICROSOFT RESEARCH

Non interference

- Private and Public expressions



when P is public and M, N are private

- Computations in M do not interfere on result P

[Volpano-Smith, Pottier-Simonet, Boudol]

- ... information flow

λ -calculus

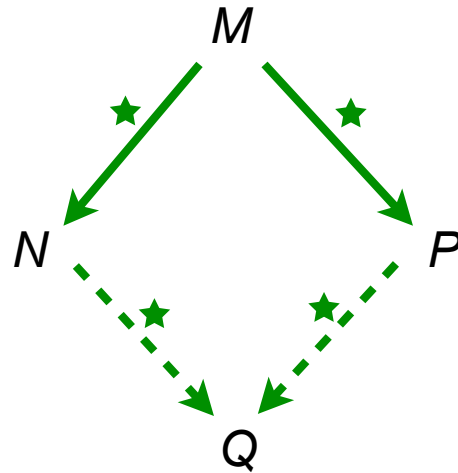
CENTRE DE RECHERCHE
COMMUN



INRIA
MICROSOFT RESEARCH

Independent reductions (1/3)

- In the λ -calculus, there are no conflicts (Church-Rosser thm)



- Need closer look at reduction steps
and notice when they can be permuted

Independent reductions (2/3)

- permutation of reduction steps (disjoint redexes)

$$\begin{array}{ccc} (\lambda x.x(Ix))(Ia) & \longrightarrow & (\lambda x.xx)(Ia) \\ \swarrow & & \searrow \\ & (\lambda x.x(Ix))a & \longrightarrow & (\lambda x.xx)a \end{array}$$

- permutation of reduction steps (nested redexes)

$$\begin{array}{ccc} (\lambda x.x(Ix))(Ia) & \longrightarrow & Ia(I(Ia)) \\ \swarrow & & \searrow \\ & (\lambda x.xx)(Ia) & \longrightarrow & Ia(Ia) \end{array}$$

Independent reductions (3/3)

- problem: copies of redexes

$$\begin{array}{l} (\lambda x.xx)(la) \longrightarrow la(la) \longrightarrow a(la) \\ \searrow \\ (\lambda x.xx)a \longrightarrow aa \end{array}$$

- how to define easily equivalence by permutations ?

Exercices

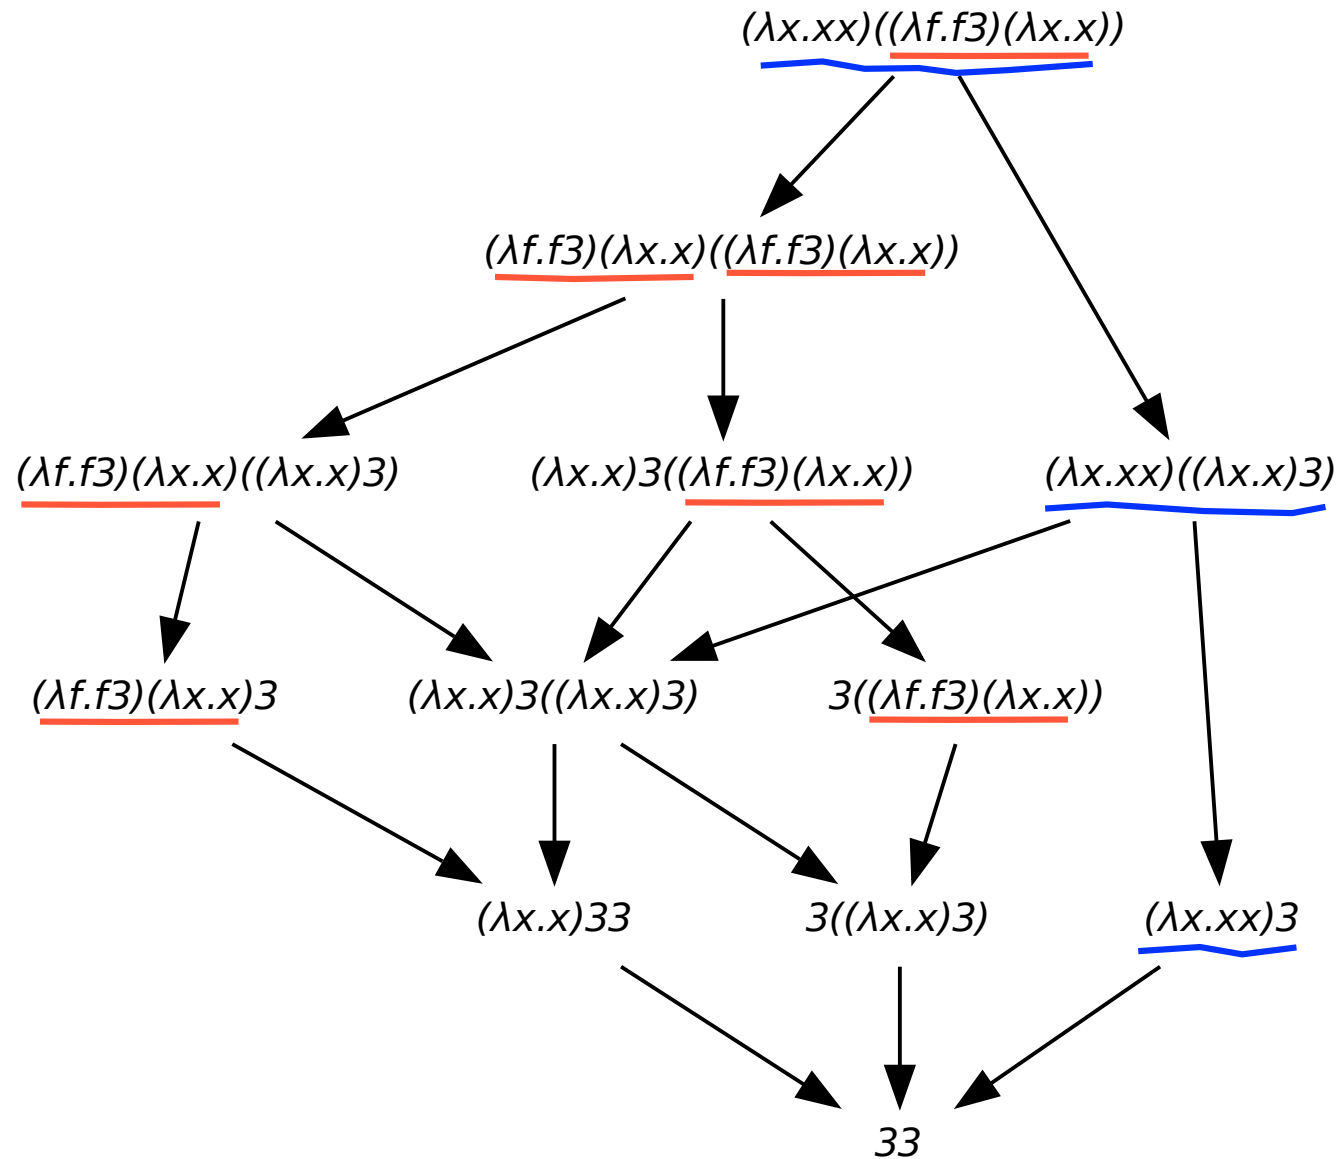
CENTRE DE RECHERCHE
COMMUN



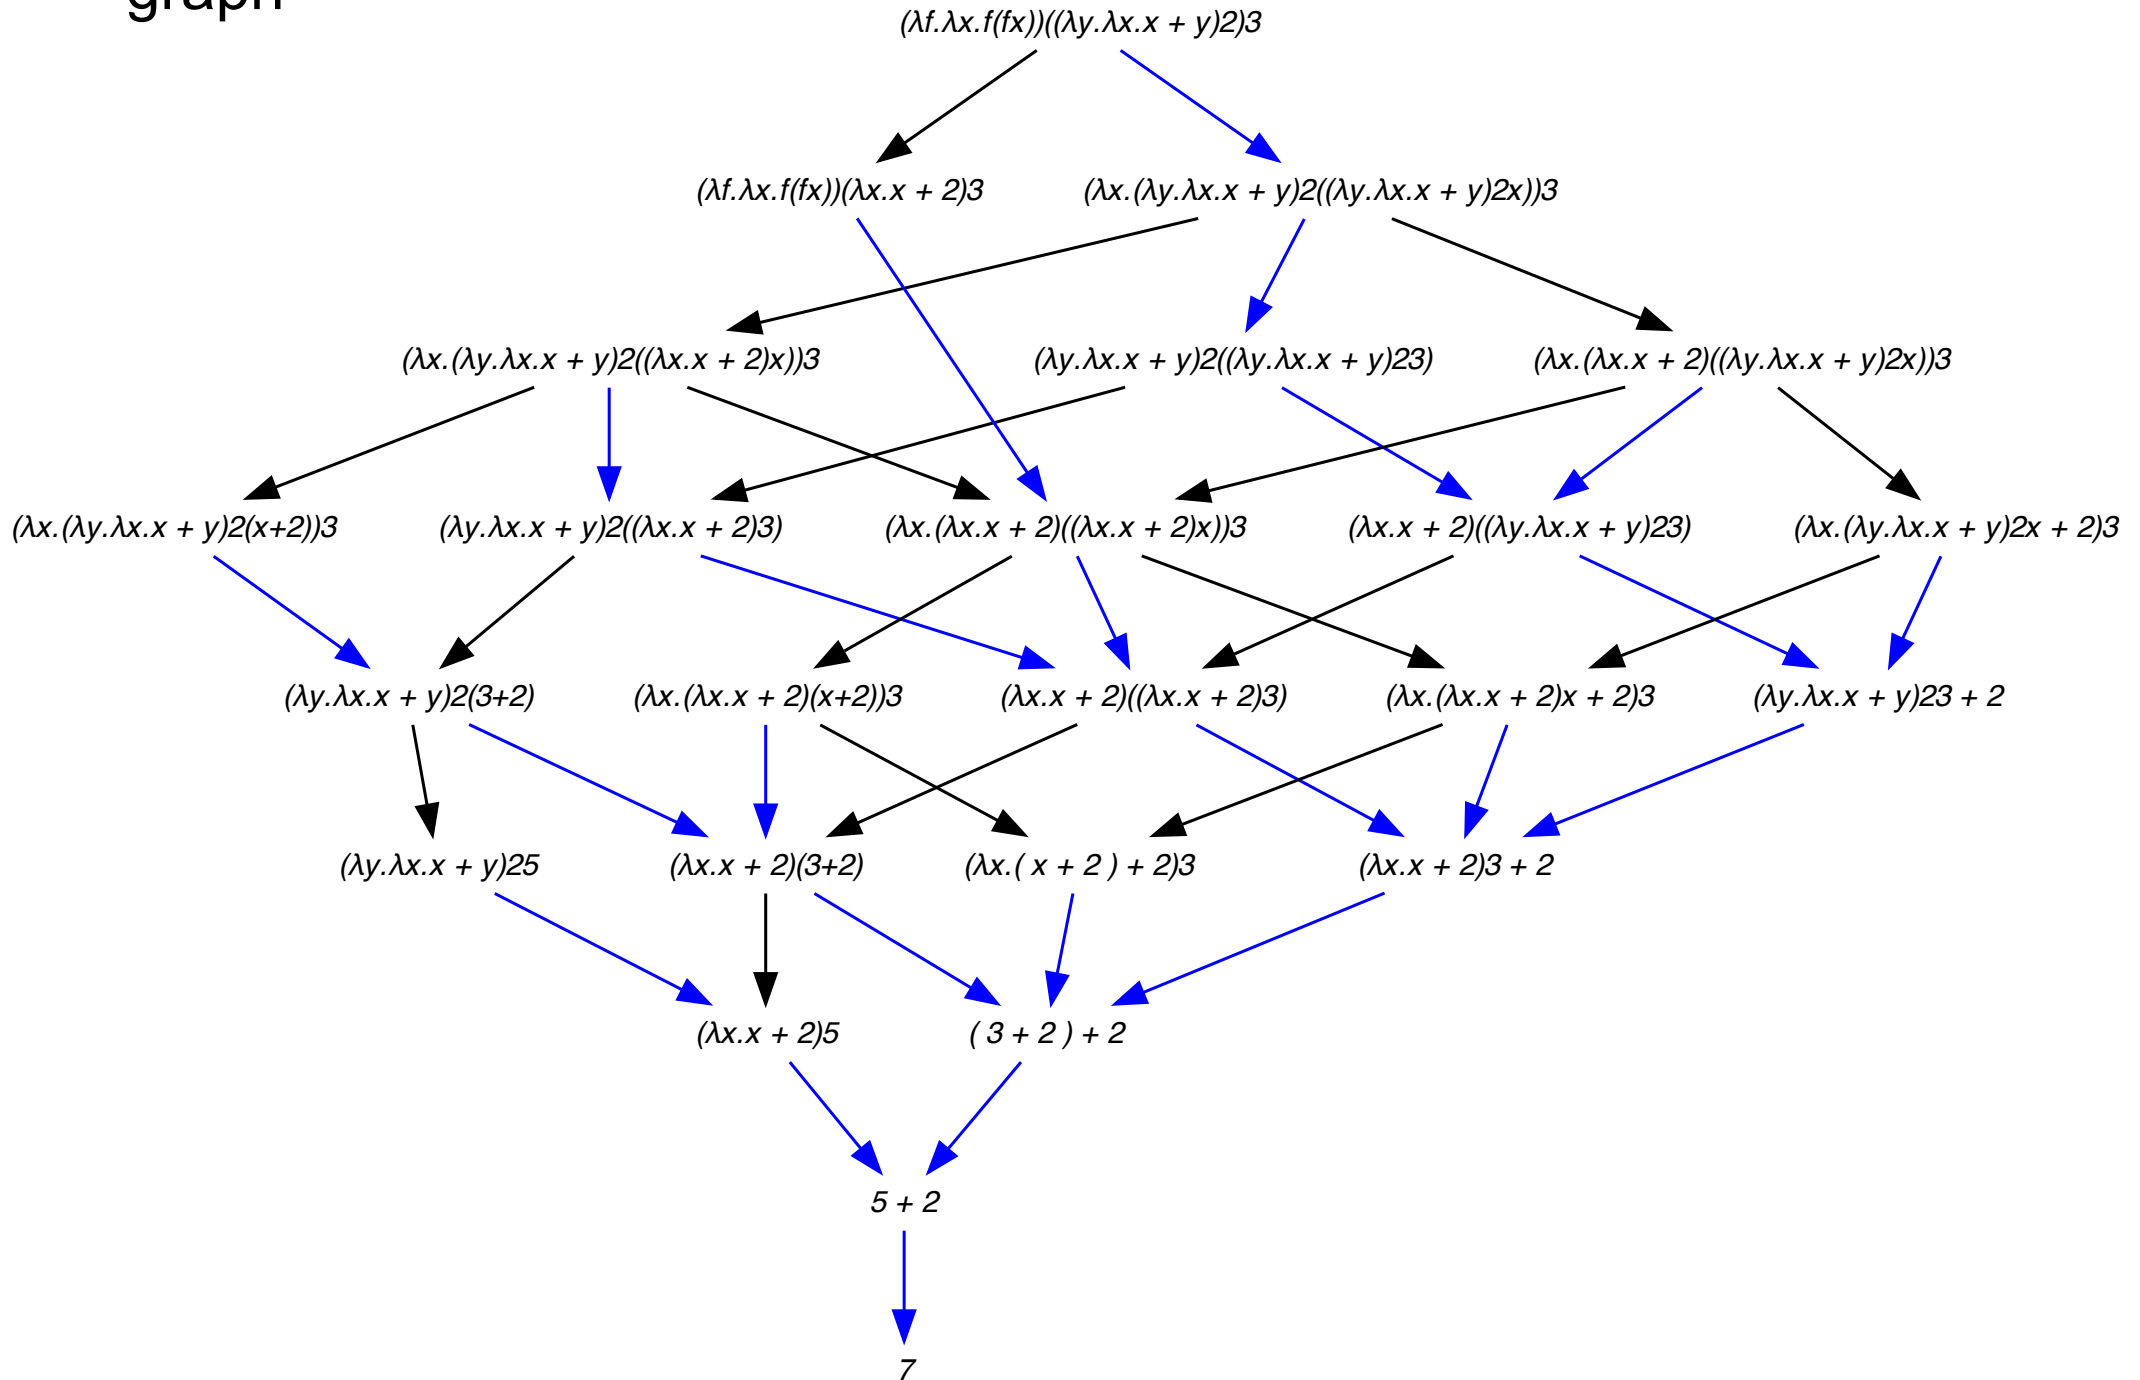
INRIA
MICROSOFT RESEARCH

Exercise

- Show reductions equivalent by permutations in following reduction graph



- Show reductions equivalent by permutations in following reduction graph



- Show reductions equivalent by permutations in following reduction graph

