

Message Passing

Johannes Åman Pohjola
CSE, UNSW
Term 2 2022

Where we are at

Last week, we saw *semaphores* and *monitors*, concluding our examination of shared variable concurrency.

For the rest of this course, our focus will be on *message passing*.

Distributed Programming

- distributed program:** processes can be distributed across machines → no shared memory (usually)
processes share *communication channels* for message passing
- languages:** Promela (synchronous and asynchronous MP), Java (RPC, RMI, ...)
- libraries:** sockets, message passing interface (MPI), parallel virtual machine (PVM) etc.

Message Passing

A *channel* is a typed FIFO queue between processes.

	Ben-Ari	Promela
send a message	$ch \leftarrow x$	$ch ! x$
recieve a message	$ch \Rightarrow y$	$ch ? y$

Message Passing

A *channel* is a typed FIFO queue between processes.

	Ben-Ari	Promela
send a message	$ch \leftarrow x$	$ch ! x$
recieve a message	$ch \Rightarrow y$	$ch ? y$

Synchronous channels

A *synchronous* channel has queue capacity 0. Both the send and the receive operation block until both are ready. When they are, they execute at the same time, and assign the value of x to y .

Message Passing

A *channel* is a typed FIFO queue between processes.

	Ben-Ari	Promela
send a message	$ch \leftarrow x$	$ch ! x$
recieve a message	$ch \Rightarrow y$	$ch ? y$

Synchronous channels

A *synchronous* channel has queue capacity 0. Both the send and the receive operation block until both are ready. When they are, they execute at the same time, and assign the value of x to y .

Asynchronous channels

For *asynchronous* channels, send doesn't block. It appends the value of x to the queue of ch . Receive blocks, until ch contains a message. When it does, the oldest message is removed, and its content is stored in y .

Taxonomy of Asynchronous Message Passing

Asynchronous channels may be...

Reliable: all messages sent will eventually arrive.

Lossy: messages may be lost in transit.

FIFO: messages will arrive in order.

Unordered: messages can arrive out-of-order.

Error-detecting: received messages aren't garbled in transit (or if they are, we can tell).

Taxonomy of Asynchronous Message Passing

Asynchronous channels may be...

Reliable: all messages sent will eventually arrive.

Lossy: messages may be lost in transit.

FIFO: messages will arrive in order.

Unordered: messages can arrive out-of-order.

Error-detecting: received messages aren't garbled in transit (or if they are, we can tell).

Example

TCP is reliable and FIFO. UDP is lossy and unordered, but error-detecting.

Algorithm 2.1: Producer-consumer (channels)

channel of integer ch

producer**consumer**

integer x

loop forever

p1: x \leftarrow producep2: ch \leftarrow x

integer y

loop forever

q1: ch \Rightarrow y

q2: consume(y)

Conway's Problem

Example

Input on channel inC: a sequence of characters

Output on channel outC:

- The sequence of characters from inC, with runs of $2 \leq n \leq 9$ occurrences of the same character c replaced by the n and c
- a newline character after every K th character in the output.

Conway's Problem

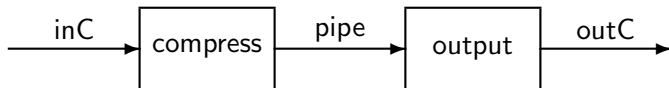
Example

Input on channel inC: a sequence of characters

Output on channel outC:

- The sequence of characters from inC, with runs of $2 \leq n \leq 9$ occurrences of the same character c replaced by the n and c
- a newline character after every K th character in the output.

Let's use message-passing for separation of concerns:



Algorithm 2.2: Conway's problemconstant integer $MAX \leftarrow 9$ constant integer $K \leftarrow 4$

channel of integer inC, pipe, outC

compresschar c, previous $\leftarrow 0$ integer n $\leftarrow 0$ inC \Rightarrow previous**loop** foreverp1: inC \Rightarrow cp2: **if** (c = previous) and
(n < MAX - 1)p3: n \leftarrow n + 1**else**p4: **if** n > 0p5: pipe \leftarrow i2c(n+1)p6: n \leftarrow 0p7: pipe \leftarrow previousp8: previous \leftarrow c**output**

char c

integer m $\leftarrow 0$ **loop** foreverq1: pipe \Rightarrow cq2: outC \leftarrow cq3: m \leftarrow m + 1q4: **if** m \geq Kq5: outC \leftarrow newlineq6: m \leftarrow 0

q7:

q8:

Reminder: Matrix Multiplication

Example

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \times \begin{pmatrix} 1 & 0 & 2 \\ 0 & 1 & 2 \\ 1 & 0 & 0 \end{pmatrix} = \begin{pmatrix} 4 & 2 & 6 \\ 10 & 5 & 18 \\ 16 & 8 & 30 \end{pmatrix}$$

Reminder: Matrix Multiplication

Example

$$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \times \begin{pmatrix} 1 & 0 & 2 \\ 0 & 1 & 2 \\ 1 & 0 & 0 \end{pmatrix} = \begin{pmatrix} 4 & 2 & 6 \\ 10 & 5 & 18 \\ 16 & 8 & 30 \end{pmatrix}$$

Let $p, q, r \in \mathbb{N}$. Let $A = (a_{i,j})_{\substack{1 \leq i \leq p \\ 1 \leq j \leq q}} \in \mathbb{T}^{p \times q}$ and

$B = (b_{j,k})_{\substack{1 \leq j \leq q \\ 1 \leq k \leq r}} \in \mathbb{T}^{q \times r}$ be two (compatible) matrices. Recall that

the matrix $C = (c_{i,k})_{\substack{1 \leq i \leq p \\ 1 \leq k \leq r}} \in \mathbb{T}^{p \times r}$ is their *product*, $A \times B$, iff, for

all $1 \leq i \leq p$ and $1 \leq k \leq r$:

$$c_{i,k} = \sum_{j=1}^q a_{i,j} b_{j,k}$$

Algorithms for Matrix Multiplication

The standard algorithm for matrix multiplication is:

for all rows i of A do:

 for all columns k of B do:

 set $c_{i,k}$ to 0

 for all columns j of A do:

 add $a_{i,j}b_{j,k}$ to $c_{i,k}$

Because of the three nested loops, its complexity is $\mathcal{O}(p \cdot q \cdot r)$.

Algorithms for Matrix Multiplication

The standard algorithm for matrix multiplication is:

for all rows i of A do:

 for all columns k of B do:

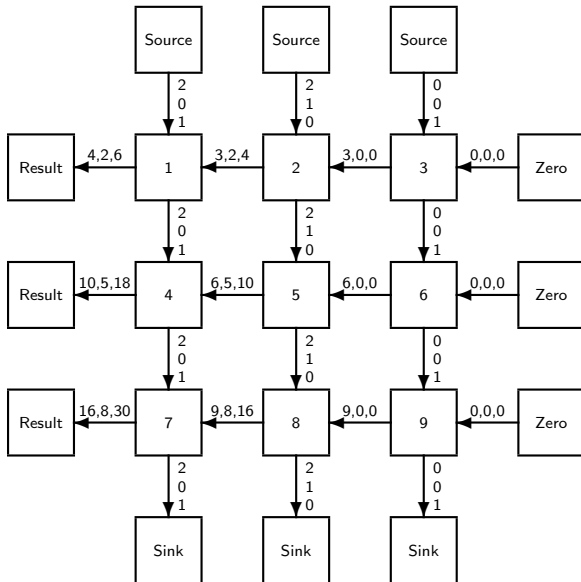
 set $c_{i,k}$ to 0

 for all columns j of A do:

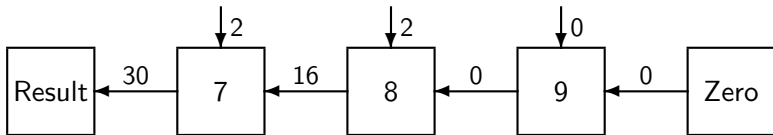
 add $a_{i,j}b_{j,k}$ to $c_{i,k}$

Because of the three nested loops, its complexity is $\mathcal{O}(p \cdot q \cdot r)$. In case both matrices are quadratic, i.e., $p = q = r$, that's $\mathcal{O}(p^3)$.

Process Array for Matrix Multiplication



Computation of One Element



Algorithm 2.3: Multiplier process with channels

integer FirstElement
channel of integer North, East, South, West
integer Sum, integer SecondElement

loop forever

p1: North \Rightarrow SecondElement

p2: East \Rightarrow Sum

p3: Sum \leftarrow Sum + FirstElement · SecondElement

p4: South \leftarrow SecondElement

p5: West \leftarrow Sum

Algorithm 2.4: Multiplier with channels and selective input

integer FirstElement
channel of integer North, East, South, West
integer Sum, integer SecondElement

```
loop forever
  either
p1:   North  $\Rightarrow$  SecondElement
p2:   East  $\Rightarrow$  Sum
    or
p3:   East  $\Rightarrow$  Sum
p4:   North  $\Rightarrow$  SecondElement
p5:   South  $\Leftarrow$  SecondElement
p6:   Sum  $\Leftarrow$  Sum + FirstElement · SecondElement
p7:   West  $\Leftarrow$  Sum
```

Multiplier Process in Promela

```
1  proctype Multiplier(byte Coeff;
2      chan North;
3      chan East;
4      chan South;
5      chan West)
6  {
7      byte Sum, X;
8      for (i : 0..(SIZE-1)) {
9          if :: North ? X -> East ? Sum;
10             :: East ? Sum -> North ? X;
11             fi;
12             South ! X;
13             Sum = Sum + X*Coeff;
14             West ! Sum;
15     }
16 }
```

Algorithm 2.5: Dining philosophers with channels

channel of boolean forks[5]

philosopher i	fork i
boolean dummy	boolean dummy
loop forever	loop forever
p1: think	q1: forks[i] \Leftarrow true
p2: forks[i] \Rightarrow dummy	q2: forks[i] \Rightarrow dummy
p3: forks[i+1] \Rightarrow dummy	q3:
p4: eat	q4:
p5: forks[i] \Leftarrow true	q5:
p6: forks[i+1] \Leftarrow true	q6:

NB

The many shared channels make it possible to give forks directly to other philosophers, rather than putting them back on the table.

Synchronous Message Passing

Recall that, when message passing is synchronous, the exchange of a message requires **coordination** between sender and receiver (sometimes called a *handshaking* mechanism).

In other words, the sender is **blocked** until the receiver is ready to cooperate.

Synchronous Transition Diagrams

Definition

A *synchronous transition diagram* is a parallel composition $P_1 \parallel \dots \parallel P_n$ of n (sequential) transition diagrams P_1, \dots, P_n called *processes*.

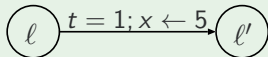
The processes P_i

- do **not** share variables
- communicate along channels C, D, \dots connecting processes by way of
 - *output* statements $C \Leftarrow e$
for sending the value of expression e along channel C
 - *input* statements $C \Rightarrow x$
for receiving a value along channel C into variable x

Edges in (Sequential) Transition Diagrams

For shared variable concurrency, labels $b; f$, where b is a Boolean condition and f is a state transformation sufficed.

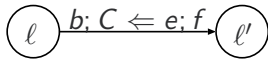
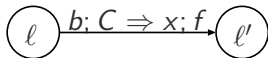
Example



Now, we call such transitions *internal*.

I/O Transitions

We extend this notation to message passing by allowing the guard to be combined with an input or an output statement:



Example 1

Let $P = P_1 \parallel P_2$ be given as:



Obviously, $\{\top\} P \{x = 1\}$, but how to prove it?

Some notation

For an n -tuple $x = \langle x_1, \dots, x_i, \dots, x_n \rangle$, we define

$$x[i \leftarrow e] = \langle x_1, \dots, e, \dots, x_n \rangle$$

$x[i \leftarrow e]$ is like x , except the i :th element is replaced with e .

Example

$$\langle 1, 5, 7 \rangle [2 \leftarrow 3] = \langle 1, 3, 7 \rangle$$

Semantics: Closed Product

Definition

The *closed product* of $P_i = (L_i, T_i, s_i, t_i)$, for $1 \leq i \leq n$ (with disjoint local variable sets), is defined as $P = (L, T, s, t)$, where:

$$L = L_1 \times \dots \times L_n \quad s = \langle s_1, \dots, s_n \rangle \quad t = \langle t_1, \dots, t_n \rangle$$

and $l \xrightarrow{a} l' \in T$ holds iff

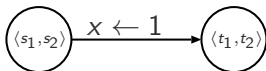
- 1 $l' = l[i \rightarrow l'_i]$ and $l_i \xrightarrow{a} l'_i \in T_i$ is an internal transition, or
- 2 $l' = l[i \rightarrow l'_i][j \rightarrow l'_j]$ and $i \neq j$,

with $l_i \xrightarrow{b; C \leftarrow e; f} l'_i \in T_i$ and $l_j \xrightarrow{b'; C \Rightarrow x; g} l'_j \in T_j$, and
 $a = b \wedge b'; f \circ g \circ [x \leftarrow e]$

Example 1 cont'd



Observe that the closed product is just



so validity of $\{\top\} P \{x = 1\}$ follows from

$$\models \top \implies (x = 1) \circ \llbracket x \leftarrow 1 \rrbracket$$

which is immediate.

(See the glossary of notation for the meaning of all these strange symbols.)

Verification

To show that the Hoare triple

$$\{\phi\} P_1 \parallel \dots \parallel P_n \{\psi\}$$

is valid, it suffices to prove

$$\{\phi\} P \{\psi\}$$

where P is the closed product of the P_i .

Verification

To show that the Hoare triple

$$\{\phi\} P_1 \parallel \dots \parallel P_n \{\psi\}$$

is valid, it suffices to prove

$$\{\phi\} P \{\psi\}$$

where P is the closed product of the P_i .

There are no I/O transitions in P , so Floyd's method works.

Verification

To show that the Hoare triple

$$\{\phi\} P_1 \parallel \dots \parallel P_n \{\psi\}$$

is valid, it suffices to prove

$$\{\phi\} P \{\psi\}$$

where P is the closed product of the P_i .

There are no I/O transitions in P , so Floyd's method works.

Disadvantage

As with the product construction for shared-variable concurrency, the closed product is **exponential** in the number of processes.

Verification

To show that the Hoare triple

$$\{\phi\} P_1 \parallel \dots \parallel P_n \{\psi\}$$

is valid, it suffices to prove

$$\{\phi\} P \{\psi\}$$

where P is the closed product of the P_i .

There are no I/O transitions in P , so Floyd's method works.

Disadvantage

As with the product construction for shared-variable concurrency, the closed product is **exponential** in the number of processes.

Is there an **Owicki-Gries** equivalent for synchronous message passing?

A Simplistic Method

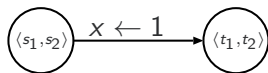
For each location ℓ in some L_i , find a local predicate Q_ℓ , only depending on P_i 's local variables.

- 1 Prove that, for all i , the local verification conditions hold, i.e.,
 $\models Q_\ell \wedge b \rightarrow Q_{\ell'} \circ f$ for each $\ell \xrightarrow{b;f} \ell' \in T_i$.
- 2 For all $i \neq j$ and *matching* pairs of I/O transitions
 $\ell_i \xrightarrow{b;C \leftarrow e;f} \ell'_i \in T_i$ and $\ell_j \xrightarrow{b';C \Rightarrow x;g} \ell'_j \in T_j$ show that

$$\models Q_{\ell_i} \wedge Q_{\ell_j} \wedge b \wedge b' \implies (Q_{\ell'_i} \wedge Q_{\ell'_j}) \circ f \circ g \circ \llbracket x \leftarrow e \rrbracket .$$

- 3 Prove $\models \phi \implies Q_{s_1} \wedge \dots \wedge Q_{s_n}$ and $\models Q_{t_1} \wedge \dots \wedge Q_{t_n} \implies \psi$.

Proof of Example 1



There are no internal transitions. There's one matching pair.

$$\begin{aligned} \top &\Longrightarrow (x = 1) \circ \llbracket x \leftarrow 1 \rrbracket \equiv 1 = 1 \\ &\equiv \top \end{aligned}$$

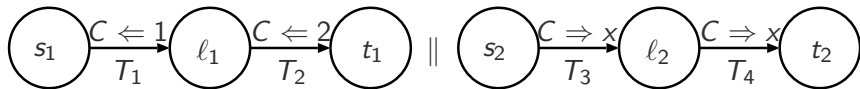
Soundness & Incompleteness

The simplistic method is sound but not complete.

It generates proof obligations for all *syntactically matching* I/O transition pairs, regardless of whether these pairs can actually be matched *semantically* (in an execution).

Example 2

Let $P = P_1 \parallel P_2$ be given as:



We cannot prove $\{\top\} P \{x = 2\}$ using the simplistic method. Proof obligations for the transition pairs (T_1, T_4) and (T_2, T_3) must be discharged. This leads to a contradiction: no assertion network can make the simplistic method work for this example.

Remedy 1: Adding Shared Auxiliary Variables

Use *shared*, *write-only* auxiliary variables to relate locations in different processes. Only output transitions need to be augmented with assignments to these shared auxiliary variables.

Pro easy

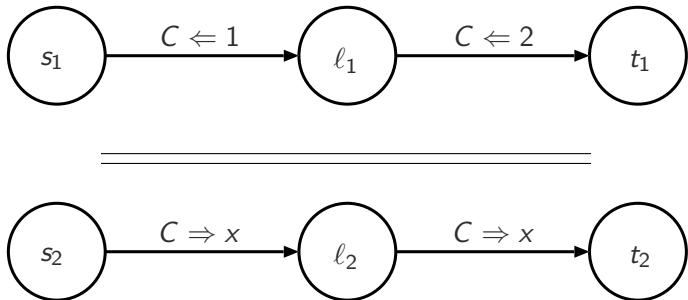
Con incomplete when channels are shared between more than two process.

Con re-introduces *interference freedom tests* for matching pairs $l_i \xrightarrow{b_i; C \leftarrow e; f_i} l'_i \in T_i$ and $l_j \xrightarrow{b_j; C \Rightarrow x; f_j} l'_j \in T_j$, and location l_m of process P_m , $m \neq i, j$:

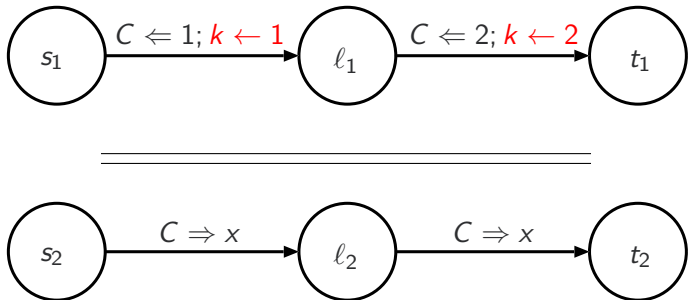
$$\models Q_{l_i} \wedge Q_{l_j} \wedge Q_{l_m} \wedge b_i \wedge b_j \implies Q_{l_m} \circ f_i \circ f_j \circ [x \leftarrow e]$$

[This method is due to Levin & Gries.]

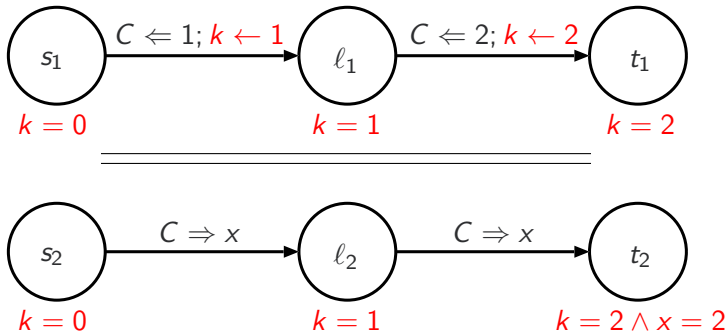
Example 2 cont'd



Example 2 cont'd



Example 2 cont'd



Levin & Gries-style Proof for Example 2

There are no internal transitions. Four matching I/O transition pairs exist, the same as in the simplistic method. Proof obligations:

$$k = 0 \implies (k = 1) \circ \llbracket k \leftarrow 1 \rrbracket \circ \llbracket x \leftarrow 1 \rrbracket \quad (1)$$

$$k = 0 \wedge k = 1 \implies (k = 1 \wedge k = 2 \wedge x = 2) \circ \llbracket k \leftarrow 1 \rrbracket \circ \llbracket x \leftarrow 1 \rrbracket \quad (2)$$

$$k = 1 \wedge k = 0 \implies (k = 2 \wedge k = 1) \circ \llbracket k \leftarrow 2 \rrbracket \circ \llbracket x \leftarrow 2 \rrbracket \quad (3)$$

$$k = 1 \implies (k = 2 \wedge x = 2) \circ \llbracket k \leftarrow 2 \rrbracket \circ \llbracket x \leftarrow 2 \rrbracket \quad (4)$$

No interference freedom proof obligations are generated in this example since there is no third process.

Levin & Gries-style Proof for Example 2 cont'd

Thanks to contradictory assumptions about k , (2) and (3) are vacuously true.

The right-hand-sides of the implications (1) and (4) simplify to \top , which discharges those proof obligations, e.g., for the RHS of (1):

$$\begin{aligned} (k = 1) \circ \llbracket k \leftarrow 1 \rrbracket \circ \llbracket x \leftarrow 1 \rrbracket &\equiv 1 = 1 \\ &\equiv \top \end{aligned}$$

Remedy 2: Local Auxiliary Variables + Invariant

Use *local*, *write only* auxiliary variables + a global *communication invariant* I to relate values of local auxiliary variables in the various processes.

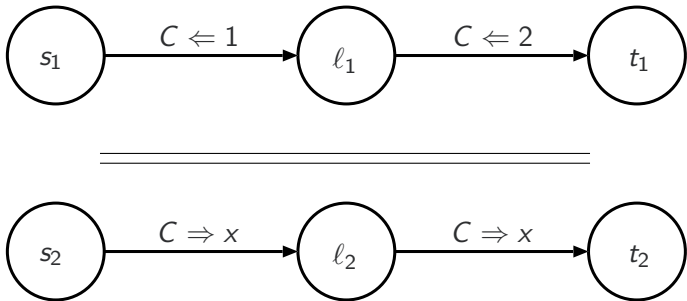
Pro no interference freedom tests

Con more complicated proof obligation for communication steps:

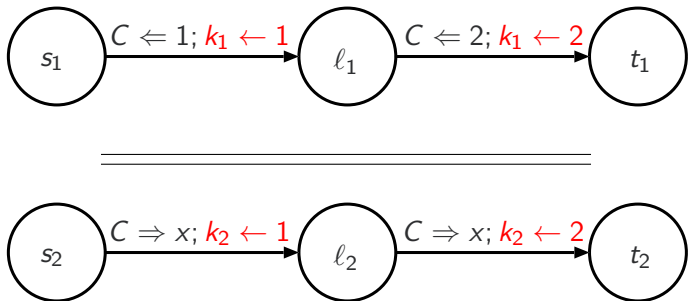
$$\models Q_{\ell_i} \wedge Q_{\ell_j} \wedge b \wedge b' \wedge I \implies (Q_{\ell'_i} \wedge Q_{\ell'_j} \wedge I) \circ f \circ g \circ \llbracket x \leftarrow e \rrbracket$$

[This is the *AFR method*, named for **A**pt, **F**rancez, and de **R**oever.]

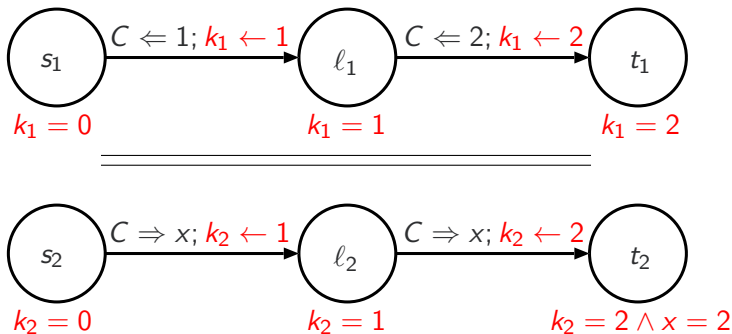
Example 2 cont'd



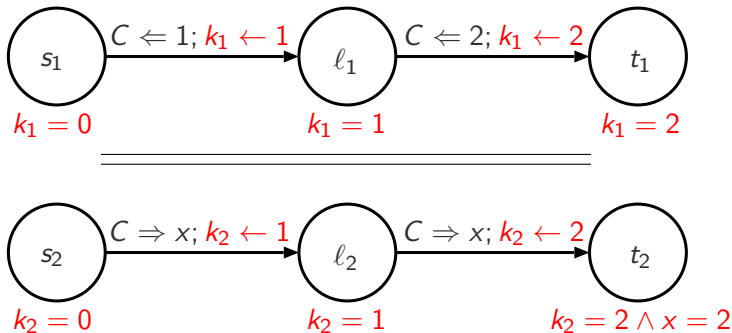
Example 2 cont'd



Example 2 cont'd



Example 2 cont'd



Define $I \equiv (k_1 = k_2)$.

AFR-style Proof for Example 2

There are no internal transitions. Four matching pairs of I/O transitions exist, with these proof obligations:

$$k_1 = 0 \wedge k_2 = 0 \wedge k_1 = k_2 \implies (k_1 = 1 \wedge k_2 = 1 \wedge k_1 = k_2) \circ \llbracket k_1 \leftarrow 1 \rrbracket \circ \llbracket k_2 \leftarrow 1 \rrbracket \circ \llbracket x \leftarrow 1 \rrbracket \quad (5)$$

$$k_1 = 0 \wedge k_2 = 1 \wedge k_1 = k_2 \implies (k_1 = 1 \wedge k_2 = 2 \wedge x = 2 \wedge k_1 = k_2) \circ \llbracket k_1 \leftarrow 1 \rrbracket \circ \llbracket k_2 \leftarrow 2 \rrbracket \circ \llbracket x \leftarrow 1 \rrbracket \quad (6)$$

$$k_1 = 1 \wedge k_2 = 0 \wedge k_1 = k_2 \implies (k_1 = 2 \wedge k_2 = 1 \wedge k_1 = k_2) \circ \llbracket k_1 \leftarrow 2 \rrbracket \circ \llbracket k_2 \leftarrow 1 \rrbracket \circ \llbracket x \leftarrow 2 \rrbracket \quad (7)$$

$$k_1 = 1 \wedge k_2 = 1 \wedge k_1 = k_2 \implies (k_1 = 2 \wedge k_2 = 2 \wedge x = 2 \wedge k_1 = k_2) \circ \llbracket k_1 \leftarrow 2 \rrbracket \circ \llbracket k_2 \leftarrow 2 \rrbracket \circ \llbracket x \leftarrow 2 \rrbracket \quad (8)$$

AFR-style Proof for Example 2 cont'd

Thanks to the invariant $k_1 = k_2$, (6) and (7) are vacuously true.
 The right-hand-sides of the implications (5) and (8) simplify to \top ,
 which discharges those proof obligations, e.g., for the RHS of (8):

$$\begin{aligned}
 & (k_1 = 2 \wedge k_2 = 2 \wedge x = 2 \wedge k_1 = k_2) \circ \llbracket k_1 \leftarrow 2 \rrbracket \circ \llbracket k_2 \leftarrow 2 \rrbracket \circ \llbracket x \leftarrow 2 \rrbracket \\
 \equiv & 2 = 2 \wedge 2 = 2 \wedge 2 = 2 \wedge 2 = 2 \\
 \equiv & \top
 \end{aligned}$$

What Now?

Next lecture, we'll be looking at **proof methods for termination** (convergence and deadlock freedom) in sequential, shared-variable concurrent, and message-passing concurrent settings.

After the break, we'll see a **compositional** proof method for verification, proving properties for **asynchronous communication**, and, if time on Thursday, we'll talk about **process algebra**.

Assignment 1 is out! Read the spec ASAP!.

Levin & Gries in full, part 1

For each $\ell \in L_i$, the annotation Q_ℓ should only depend on P_i 's local variables, and shared write-only auxiliary variables. Shared variables should only be assigned to in output transitions.

- 1 Prove that, for all i , the local verification conditions hold, i.e.,
 $\models Q_\ell \wedge b \rightarrow Q_{\ell'} \circ f$ for each $\ell \xrightarrow{b;f} \ell' \in T_i$.
- 2 For all $i \neq j$ and $\ell_i \xrightarrow{b;C \leftarrow e;f} \ell'_i \in T_i$ and $\ell_j \xrightarrow{b';C \Rightarrow x;g} \ell'_j \in T_j$ show that

$$\models Q_{\ell_i} \wedge Q_{\ell_j} \wedge b \wedge b' \implies (Q_{\ell'_i} \wedge Q_{\ell'_j}) \circ f \circ g \circ [x \leftarrow e] .$$

- 3 For all $i \neq j$ and $\ell_i \xrightarrow{b_i;C \leftarrow e;f_i} \ell'_i \in T_i$ and $\ell_j \xrightarrow{b_j;C \Rightarrow x;f_j} \ell'_j \in T_j$, and location ℓ_m of process P_m , $m \neq i, j$:

$$\models Q_{\ell_i} \wedge Q_{\ell_j} \wedge Q_{\ell_m} \wedge b_i \wedge b_j \implies Q_{\ell_m} \circ f_i \circ f_j \circ [x \leftarrow e]$$

Levin & Gries in full, part 2

Let k_1, \dots, k_m be all auxiliary variables used in the transition diagrams. We assume that ϕ, ψ mentions none of these auxiliaries.

4 Prove

$$\models \phi \implies \exists k_1, \dots, k_m. Q_{s_1} \wedge \dots \wedge Q_{s_n}$$

and

$$\models Q_{t_1} \wedge \dots \wedge Q_{t_n} \implies \psi$$

These four items suffice to prove $\{\phi\} P \{\psi\}$, where P is the closed product of the P_i :s.

NB

The existential quantification over the shared variables allows the final Hoare triple to make no mention of the auxiliary variables.

AFR in full, part 1

For each $\ell \in L_i$, the annotation Q_ℓ should only depend on P_i 's local variables, and local write-only auxiliary variables. Auxiliary variables should only be assigned to in I/O transitions. The communication invariant I should only mention auxiliary variables.

- 1 Prove that, for all i , the local verification conditions hold, i.e.,
 $\models Q_\ell \wedge b \rightarrow Q_{\ell'} \circ f$ for each $\ell \xrightarrow{b;f} \ell' \in T_i$.
- 2 For all $i \neq j$ and $\ell_i \xrightarrow{b_i;C \leftarrow e;f_i} \ell'_i \in T_i$ and $\ell_j \xrightarrow{b_j;C \Rightarrow x;f_j} \ell'_j \in T_j$,
 and location ℓ_m of process P_m , $m \neq i, j$:

$$\models Q_{\ell_i} \wedge Q_{\ell_j} \wedge b \wedge b' \wedge I \implies (Q_{\ell'_i} \wedge Q_{\ell'_j} \wedge I) \circ f \circ g \circ [x \leftarrow e]$$

AFR in full, part 2

Let k_1, \dots, k_m be all auxiliary variables used in the transition diagrams. We assume that ϕ, ψ mentions none of these auxiliaries.

3 Prove

$$\models \phi \implies \exists k_1, \dots, k_m. Q_{s_1} \wedge \dots \wedge Q_{s_n} \wedge I$$

and

$$\models Q_{t_1} \wedge \dots \wedge Q_{t_n} \wedge I \implies \psi$$

These three items suffice to prove $\{\phi\} P \{\psi\}$, where P is the closed product of the P_i :s.

NB

We could allow non-auxiliary variables in I , at the expense of making proof obligation 1 more involved.