

Parallel Programming

Prof. Dr. Uwe Kastens

Winter 2014 / 2015

Objectives

The participants are taught to understand and to apply

- **fundamental concepts** and **high-level paradigms** of parallel programs,
- **systematic methods** for developing parallel programs,
- **techniques** typical for parallel programming in Java;
- English language in a lecture.

Exercises:

- The exercises will be tightly integrated with the lectures.
- Small teams will solve given assignments practically supported by a lecturer.
- Homework assignments will be solved by those teams.

Contents

Week	Topic
1	1. Introduction
2	2. Properties of Parallel Programs
4	3. Monitors in General and in Java
5	4. Systematic Development of Monitors
6	5. Data Parallelism: Barriers
7	6. Data Parallelism: Loop Parallelization
11	7. Asynchronous Message Passing
12	8. Messages in Distributed Systems
14	9. Synchronous message Passing
	10. Conclusion

Prerequisites

Topic	Course that teaches it
practical experience in programming Java	Grundlagen der Programmierung 1, 2
foundations in parallel programming	Grundlagen der Programmierung 2, Konzepte und Methoden der Systemsoftware (KMS)
process, concurrency, parallelism, interleaved execution address spaces, threads, process states monitor	KMS KMS KMS KMS
process, concurrency, parallelism, threads, synchronization, monitors in Java	GP, KMS GP, KMS GP, KMS
verification of properties of programs	Modellierung

Organization of the course

Lecturer

Prof. Dr. Uwe Kastens:

Office hours: on appointment by email

Teaching Assistant:

- Peter Pfahler

Lecture

- V2 Mon 11:15 - 12:45, F1.110

Start date: Oct 13, 2014

Tutorials

- Grp 1 Mon 09.30 - 11.00 Even Weeks, F2.211 / F1 pool, Start Oct. 27
- Grp 2 Fri 11.00 - 12.30 Odd Weeks, F2.211 / F1 pool, Start Oct. 24

Schedule

Tutorial	Group 1, Mon 09:30	Group 2, Fri 11:00
1	Oct 27	Oct 24
2	Nov 10	Nov 07
3	Nov 24	Nov 21
4	Dec 08	Dec 05
5	Jan 05	Dec 19
6	Jan 19	Jan 16
7	Feb 02	Jan 30

Examination

Oral examinations of 20 to 30 min duration. For students of the Computer Science Masters Program the examination is part of a module examination, see [Registering for Examinations](#). In general the examination is held in English. As an alternative, the candidates may choose to give a short presentation in English at the begin of the exam; then the remainder of the exam is held in German. In this case the candidate has to ask via email for a topic of that presentation latest a week before the exam.

Literature

Course material „**Parallel Programming**“

<http://ag-kastens.upb.de/lehre/material/ppje>

Course material „Grundlagen der Programmierung“ (in German)

Course material „**Software-Entwicklung I + II**“ WS, SS 1998/1999:(in German)

<http://ag-kastens.upb.de/lehre/material/swei>

Course material „**Konzepte und Methoden der Systemsoftware**“ (in German)

Course material „**Modellierung**“ (in German)

<http://ag-kastens.upb.de/lehre/material/model>

Gregory R. Andrews: **Concurrent Programming**, Addison-Wesley, 1991

Gregory R. Andrews: **Foundations of multithreaded, parallel, and distributed programming**, Addison-Wesley, 2000

David Gries: **The Science of Programming**, Springer-Verlag, 1981

Scott Oaks, Henry Wong: **Java Threads**, 2nd ed., O'Reilly, 1999

Jim Farley: **Java Distributed Computing**, O'Reilly, 1998

Doug Lea: **Concurrent Programming in Java**, Addison-Wesley, 2nd Ed., 2000


Fundamental notions (repeated): Parallel processes

process:

Execution of a sequential part of a program in its storage (address space).
Variable state: contents of the storage and the position of execution


parallel processes:

several processes, which are executed simultaneously on several processors



interleaved processes:

several processes, which are executed piecewise alternately on a single processor
processes are switched by a common process manager or by the processes themselves.

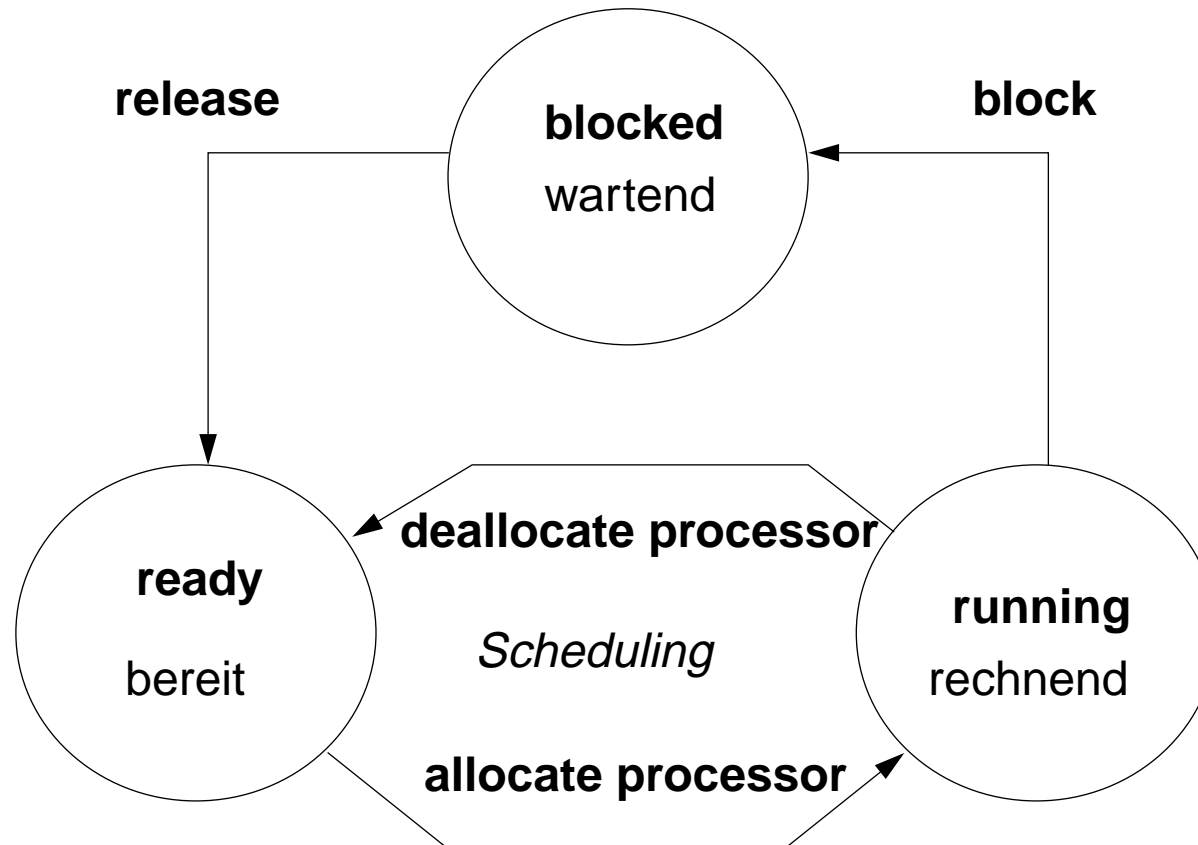


interleaved execution can simulate parallel execution;
frequent process switching gives the illusion that all process execute steadily.

concurrent processes:

processes, that can be executed in parallel or interleaved

Fundamental notions (repeated): States and transitions of processes



see KMS 2-17, 2-18

Threads (lightweight processes, Leichtgewichtsprozesse):

Processes, that are executed in parallel or interleaved in one common address space; process switching is easy and fast.

Applications of parallel processes

- **Event-based user interfaces:**
Events are propagated by a specific process of the system.
Time consuming computations should be implemented by concurrent processes,
to avoid blocking of the user interface.
- **Simulation** of real processes:
e. g. production in a factory
- **Animation:**
visualization of processes, algorithms; games
- **Control** of machines in **Real-Time:**
processes in the computer control external facilities,
e. g. factory robots, airplane control
- **Speed-up of execution** by parallel computation:
several processes cooperate on a common task,
e. g. parallel sorting of huge sets of data

The application classes follow **different objectives**.

Create threads in Java - technique: implement Runnable

Processes, threads in Java:

concurrently executed in the **common address space** of the program (or applet),
objects of class `Thread` with certain properties

Technique 1: A user's class implements the interface Runnable:

```
class MyTask implements Runnable
{
    ...
    public void run ()           The interface requires to implement the method run
    {...}                       - the program part to be executed as a process.
    public MyTask(...) {...}    The constructor method.
}
```

The process is created as an **object of the predefined class Thread**:

```
Thread aTask = new Thread (new MyTask (...));
```

The following call starts the process:

```
aTask.start();
```

The new process starts executing in parallel with the initiating one.

This technique (implement the interface `Runnable`) should be used if

- the **new process need not be influenced** any further;
i. e. it performs its task (method `run`) and then terminates, or
- the **user's class is to be defined as a subclass** of a class different from `Thread`

Create threads in Java - technique: subclass of Thread

Technique 2:

The user's class is defined as a **subclass of the predefined class Thread**:

```
class DigiClock extends Thread
{ ...
  public void run ()                Overrides the Thread method run.
  {...}                             The program part to be executed as a process.
  DigiClock (...) {...}             The constructor method.
}
```

The process is created as an **object of the user's class** (it is a **Thread** object as well):

```
Thread clock = new DigiClock (...);
```

The following call starts the process:

```
clock.start();    The new process starts executing in parallel with the initiating one.
```

This technique (subclass of **Thread**) should be used if the new process **needs to be further influenced**; hence, **further methods** of the user's class are to be defined and called from outside the class, e. g. to interrupt the process or to terminate it. The class can not have another superclass!

Important methods of the class Thread

```
public void run ();
```

is to be overridden with a method that contains the code to be executed as a process

```
public void start ();
```

starts the execution of the process

```
public void suspend ();
```

(deprecated, deadlock-prone),

suspends the indicated process temporarily: e. g. `clock.suspend()`;

```
public void resume ();
```

(deprecated), resumes the indicated process: `clock.resume()`;

```
public void join () throws InterruptedException;
```

the calling process waits until the indicated process has terminated

```
try { auftrag.join(); } catch (Exception e){}
```

```
public static void sleep (long millisec) throws InterruptedException;
```

the calling process waits at least for the given time span (in milliseconds), e. g.


```
try { Thread.sleep (1000); } catch (Exception e){}
```

```
public final void stop () throws SecurityException;
```

not to be used! May terminate the process in an inconsistent state.

Example: Digital clock as a process in an applet (1)

The process displays the **current date and time** every second as a formatted text.



```
class DigiClock extends Thread
{
  public void run ()
  {
    while (running)
    {
      line.setText(new Date().toString());
      try { sleep (1000); } catch (Exception ex) {}
    }
  }
}
```

iterate until it is terminated from the outside
write the date
pause

Method, that terminates the process from the outside:

```
public void stopIt () { running = false; }
private volatile boolean running = true;

public DigiClock (Label t) {line = t;}
private Label line;
}
```

state variable
label to be used for the text

Technique **process as a subclass of Thread**, because it

- **is to be terminated** by a call of `stopIt`,
- **is to be interrupted** by calls of further `Thread` methods,
- other **super classes are not needed**.

Example: Digital clock as a process in an applet (2)

The process is created in the `init` method of the subclass of `Applet`:

```
public class DigiApp extends Applet
{
    public void init ()
    {
        Label clockText = new Label ("-----");
        add (clockText);

        clock = new DigiClock (clockText);
        clock.start();
    }

    public void start ()    { /* see below */ }
    public void stop ()    { /* see below */ }
    public void destroy () { clock.stopIt(); }

    private DigiClock clock;
}
```

create process
start process

resume process
suspend process
terminate process

Processes, which are started in an applet

- may be suspended, while the applet is invisible (`stop`, `start`);
better use synchronization or control variables instead of `suspend`, `resume`
- are to be terminated (`stopIt`), when the applet is deallocated (`destroy`).

Otherwise they bind resources, although they are not visible.

2. Properties of Parallel Programs

Goals:

- **formal reasoning** about parallel programs
- **proof properties** of parallel programs
- **develop** parallel programs such that certain **properties can be proven**

Example A:

```

x := 0; y := 0
co  x := x + 1 //
   y := y + 1
oc
z := x + y

```

Branches of `co-oc` are executed in parallel.

Proof that $z = 2$ holds at the end.

Methods:

Hoare Logic, Weakest Precondition, techniques for parallel programs

Example B:

```

x := 0; y := 0
co  x := y + 1 //
   y := x + 1
oc
z := x + y

```

Show that $z = 2$ can not be proven.

Proofs of parallel programs

Example A:

```

x := 0; y := 0 {x=0 ∧ y=0}
co
  {x+1=1} x := x + 1 {x=1} //
  {y+1=1} y := y + 1 {y=1}
oc
{x=1 ∧ y=1} → {x+y=2}
z := x + y {z=2}

```

Example B₁:

```

x := 0; y := 0 {x=0 ∧ y=0}
co
  {y+1=1} x := y + 1 {x=1} //
  {x+1=1} y := x + 1 {y=1}
oc
{x=1 ∧ y=1} → {x+y=2}
z := x + y {z=2}

```

Check each proof for correctness!

Explain!

Example B₂:

```

x := 0; y := 0 {x≥0 ∧ y≥0}
co
  {y+1>0} x := y + 1 {x>0} //
  {x+1>0} y := x + 1 {y>0}
oc
{x>0 ∧ y>0} → {x+y≥2}
z := x + y {z≥2}

```

Does an assignment of process **p** interfere with an **assertion of process q**?

Hoare Logic: a brief reminder

Formal calculus for **proving properties of algorithms or programs** [C. A. R. Hoare, 1969]

Predicates (assertions) are stated for program positions:

$$\{P\} \ S \ \{Q\}$$

A predicate, like Q , characterizes the **set of states** that any execution of the program can achieve at that position. The predicates are expressions over variables of the program.

Each triple $\{P\} \ S \ \{Q\}$ describes an effect of the execution of S . P is called a precondition, Q a postcondition of S .

The triple $\{P\} \ S \ \{Q\}$ is correct, if the following holds:

If the execution of S is begun in a state of P and **if it terminates**, the the final state is in Q (partial correctness).

Two special assertions are:

$\{\mathbf{true}\}$ characterizing all states, and $\{\mathbf{false}\}$ characterizing no state.

Proofs of program properties are constructed using **axioms** and **inference rules** which describe the effects of each kind of statement, and define how proof steps can be correctly combined.

Axioms and inference rules for sequential constructs

statement sequence

$$\frac{\begin{array}{l} \{P\} \quad S_1 \quad \{Q\} \\ \{Q\} \quad S_2 \quad \{R\} \end{array}}{\{P\} \quad S_1; S_2 \quad \{R\}}$$

1

stronger precondition

$$\frac{\begin{array}{l} \{P\} \rightarrow \{R\} \\ \{R\} \quad S \quad \{Q\} \end{array}}{\{P\} \quad S \quad \{Q\}}$$

3

weaker postcondition

$$\frac{\begin{array}{l} \{P\} \quad S \quad \{R\} \\ \{R\} \rightarrow \{Q\} \end{array}}{\{P\} \quad S \quad \{Q\}}$$

4

assignment

$$\{P_{[x/e]}\} \quad x := e \quad \{P\}$$

2

$P_{[x/e]}$ means: P with all free occurrences of x substituted by e

multiple alternative (guarded command)

$$\frac{\begin{array}{l} P \wedge \neg(B_1 \vee \dots \vee B_n) \Rightarrow Q \\ \{P \wedge B_i\} \quad S_i \quad \{Q\}, \quad 1 \leq i \leq n \end{array}}{\{P\} \quad \mathbf{if} \quad B_1 \rightarrow S_1 \quad [] \quad \dots \quad [] \quad B_n \rightarrow S_n \quad \mathbf{fi} \quad \{Q\}}$$

5

selecting iteration

$$\frac{\begin{array}{l} \{INV \wedge B_i\} \quad S_i \quad \{INV\}, \quad 1 \leq i \leq n \end{array}}{\{INV\} \quad \mathbf{do} \quad B_1 \rightarrow S_1 \quad [] \quad \dots \quad [] \quad B_n \rightarrow S_n \quad \mathbf{od} \quad \{INV \wedge \neg(B_1 \vee \dots \vee B_n)\}}$$

6

no operation

$$\{P\} \quad \mathbf{skip} \quad \{P\}$$

7

Verification: algorithm computes gcd

precondition: $x, y \in \mathbb{N}$, i. e. $x > 0, y > 0$; let G be greatest common divisor of x and y

postcondition: $a = G$

algorithm with { assertions over variables }:

{ G is gcd of x and $y \wedge x > 0 \wedge y > 0$ }

$a := x; b := y;$

{ INV: G is gcd of a and $b \wedge a > 0 \wedge b > 0$ }

do $a \neq b \rightarrow$

{ INV $\wedge a \neq b$ }

if $a > b \rightarrow$

{ G is gcd of a and $b \wedge a > 0 \wedge b > 0 \wedge a > b$ } \rightarrow

{ G is gcd of $a-b$ and $b \wedge a-b > 0 \wedge b > 0$ }

$a := a - b$

{ INV }

[] $a \leq b \rightarrow$

{ G is gcd of a and $b \wedge a > 0 \wedge b > 0 \wedge b > a$ } \rightarrow

{ G is gcd of a and $b-a \wedge a > 0 \wedge b-a > 0$ }

$b := b - a$

{ INV }

fi { INV $\wedge a \neq b \wedge \neg(a > b \vee a \leq b) \rightarrow$ INV } „there is no 3rd case for the if \rightarrow INV“

{ INV }

od

{ INV $\wedge a = b$ } \rightarrow

{ $a = G$ }

the loop terminates:

- $a+b$ decreases monotonic
- $a+b > 0$ is invariant

Weakest precondition

A similar calculus as Hoare Logic is based on the notion of weakest preconditions [Dijkstra, 1976; Gries 1981]:

Program positions are also annotated by assertions that characterize program states.

The **weakest precondition** $wp (S, Q) = P$ of a statement S maps a predicate Q on a predicate P (wp is a **predicate transformer**).

$wp (S, Q) = P$ characterizes **the largest set of states** such that if the execution of S is begun in any state of P , then the execution is **guaranteed to terminate** in a state of Q (**total correctness**).

If $P \Rightarrow wp (S, Q)$ then $\{P\} S \{Q\}$ holds in Hoare Logic.

This concept is a more goal oriented proof method compared to Hoare Logic. We need weakest precondition only in the definition of „non-interference“ in proof for parallel programs.

Examples for weakest preconditions

1. $P = \text{wp}(\text{statement}, Q)$
2. $i \leq 0 = \text{wp}(i := i + 1, i \leq 1)$
3. $\text{true} = \text{wp}(\text{if } x \geq y \text{ then } z := x \text{ else } z := y, z = \max(x, y))$
4. $(y \geq x) = \text{wp}(\text{if } x \geq y \text{ then } z := x \text{ else } z := y, z = y)$
5. $\text{false} = \text{wp}(\text{if } x \geq y \text{ then } z := x \text{ else } z := y, z = y - 1)$
6. $(x = y + 1) = \text{wp}(\text{if } x \geq y \text{ then } z := x \text{ else } z := y, z = y + 1)$
7. $\text{wp}(S, \text{true}) =$ the set of all states such that the execution of S begun in one of them is guaranteed to terminate

Interleaving - used as an abstract execution model

Processes that are not blocked may be switched **at arbitrary points** in time.
A **scheduling strategy** reduces that freedom of the scheduler.

An example shows how different results are exhibited by switching processes differently.
Two processes operate on a common variable **account**:

<code>account = 50;</code>			
<u> <i>a</i> </u>	<u> <i>b</i> </u>	<u> <i>c</i> </u>	
Process1: <code>t1 = account;</code>	<code>t1 = t1 + 10;</code>	<code>account = t1;</code>	
Process2: <code>t2 = account;</code>	<code>t2 = t2 - 5;</code>	<code>account = t2;</code>	
<u> <i>d</i> </u>	<u> <i>e</i> </u>	<u> <i>f</i> </u>	

Assume that the assignments *a* - *f* are atomic. Try any interleaved execution order of the two processes on a single processor. Check what the value of **account** is in each case.

Assume the sequences of statements $\langle a, b \rangle$ and $\langle d, e \rangle$ (or $\langle b, c \rangle$ and $\langle e, f \rangle$) are atomic and check the results of any interleaved execution order.

We get the **same variety of results**, because there are **no global variables** in *b* or *e*
The coarser execution model is sufficient.

Atomic actions

Atomic action: A sequence of (one or more) operations, the internal states of which can not be observed because it has one of the following properties:

- it is a **non-interruptable machine instruction**,
- it has the **AMO** property, or
- **Synchronization** prohibits, that the action is interleaved with those of other processes, i. e. explicitly atomic.

At-most-once property (AMO):

The construct has **at most one** point where an other process can interact:

- **Expression E:**
E has at most one variable v , that is written by a different process, and v occurs only once in E.
- **Assignment $x := E$:**
E is AMO and x is not read by a different process, or x may be read by a different process, but E does not contain any global variable.
- **Statement sequence S:**
one statement in S is AMO and all other statements in S do not have any global variable.

Atomic by AMO

Interleaving analysis is **simpler**, if **atomic decomposition is coarser**.

Check AMO property for nested constructs. Consider the most enclosing one to be atomic.

Examples: assume $x = 0; y = 0; z = 0;$ to be global
atomic AMO constructs $\langle \dots \rangle$:

$\langle t = \langle \langle x \rangle + \langle 1 \rangle \rangle; \rangle \langle x = \langle 1 \rangle; \rangle$

interleaving actions of two processes:

(1)

p1: $\langle t = 0; t = t + 1; \rangle$ ^a
p2: $\langle s = 0; s = s + 1; \rangle$ _b

p1:

$\langle x = 2; \rangle$ ^a

(2)

p2:

$\langle t = x + 1; \rangle$ _b

(3)

p1: $x = \langle y + 1 \rangle;$ ^{b a}
p2: $y = \langle x + 1 \rangle;$ _{d c}

p1:

$x = \langle y \rangle + \langle z \rangle;$ ^{c a b}

(4)

p2:

$\langle y = 1; \rangle \langle z = 2; \rangle;$ _{d e}

Interference between processes

Critical assertions characterize **observable states** of a process p :


Let $\{P\} S \{Q\}$ be the statement sequence of process p with its pre- and postcondition.

Then Q is critical.

Let T be a statement in S that is not part of an atomic statement and R its postcondition; then $C = wp(T, R)$ is critical.

For every critical assertion of the proof of p , it has to be proven that

non-interference NI (A, C) holds for each **assignment A** of every other process q :



non-interference NI (A, C) holds between

assignment A: $\{D\} x = e$ in q having precondition D in a proof of q and

assertion C on p , if the following can be proven in programming logic:

$$\{C \wedge D\} A \{C\}$$

i. e. **the execution of A does not interfere with C (can not change C)**, provided that the precondition D allows to execute A in a state where C holds.

Example: Interference between an assertion and an assignment

Consider processes p and q with **assertions at observable states**.

Consider a single critical **assertion C** in p and a single **assignment A** in q:

p: ...<...> {C} <...>...

q: ...<...> {d+1 > 0} a = d + 1; {Q} <...>...

A

Does A interfere with C? Depends on C:

1. C: a == 1

{a == 1 ∧ d + 1 > 0} a = d + 1 {a == 1} is not provable ⇒ interference

C

2. C: a > 0

{a > 0 ∧ d + 1 > 0} a = d + 1 {a > 0} is provable ⇒ non-interference

3. C: a==1 ∧ d<0

{a==1 ∧ d<0 ∧ d+1>0} a = d + 1 {a==1 ∧ d<0} is provable ⇒ non-interference

_____f_____

Non-interference checks

$x := 0; y := 0;$
 $\{x = 0 \wedge y = 0\}$
co $\{x+1 = 1\} x := x+1 \{x=1\} //$

$\{y+1 = 1\} y := y+1 \{y=1\}$

oc
 $\{x = 1 \wedge y = 1\} \Rightarrow \{x+y = 2\}$
 $z := x+y$
 $\{z = 2\}$

$NI(a, c)$ holds for all 4 cases, e.g.

$\{x+1 = 1 \wedge y+1 = 1\} y := y+1 \{x+1 = 1 \wedge y = 1\} \Rightarrow$
 $\{x+1 = 1\}$

$x := 0; y := 0;$
 $\{x = 0 \wedge y = 0\}$
co $\{y+1 = 1\} x := y+1 \{x=1\} //$

$\{x+1 = 1\} y := x+1 \{y=1\}$

oc
 $\{x = 1 \wedge y = 1\} \Rightarrow \{x+y = 2\}$
 $z := x+y$
 $\{z = 2\}$

$NI(y := x+1, y+1 = 1)$ does not hold:

$\{y+1 = 1 \wedge x+1 = 1\} y := x+1 \{y+1 = 1\}$
 is not correct

is not correct

Two inference rules for concurrent execution

The statement for **condition synchronization**

`<await B -> S>`

causes the executing process to be blocked until the condition **B** is true; then **s** is executed. The whole statement is executed as an atomic action; hence **B** holds at the begin of **s**.

$$\frac{\{P \wedge B\} S \{Q\}}{\{P\} \langle \text{await } B \rightarrow s \rangle \{Q\}}$$

The statement for **concurrent processes**

`co S1 // ... // Sn oc`

executes the statements **s_i** concurrently. It terminates when all **s_i** have terminated.

Non-Interference is to be proven.

$\{P_i\} S_i \{Q_i\}, 1 \leq i \leq n$, are **interference-free theorems**

$\{P_1 \wedge \dots \wedge P_n\} \text{co } S_1 // \dots // S_n \text{ oc } \{Q_1 \wedge \dots \wedge Q_n\}$

Avoiding interference

1. disjoint variables:

Two concurrent processes p and q are interference-free if the set of variables p writes to is disjoint from the set of variables q reads from and vice versa.

2. weakened assertions:

The assertions in the proofs of concurrent processes can in some cases be made interference-free by weakening them.

3. atomic action:

A non-interference-free assertion C can be hidden in an atomic action.

$$p:: \dots x := e \dots$$

$$p:: \dots x := e \dots$$

$$q:: \dots s1 \{C\} s2 \dots$$

$$q:: \dots \langle s1 \{C\} s2 \rangle \dots$$

4. condition synchronization:

A synchronization condition can make an interfering assignment interference-free.

$S2$ can not be executed in this state or C holds after $x:=e$

$$p:: \dots x := e \dots$$

$$p:: \dots \langle \text{await not } C \text{ or } B \rightarrow x:=e \rangle \dots$$

with $B = wp(x:=e, C)$

$$q:: \dots s1 \{C\} s2 \dots$$

$$q:: \dots s1 \{C\} s2 \dots$$