

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.

Objectives:

Understand the objectives

In the lecture:

Explanation of the objectives

Questions:

Do these objectives coincide with yours?

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

Lecture Parallel Programming WS 2014/2015 / Slide 03

Objectives:

Overview over the topics of the course

In the lecture:

Brief explanations of the topics

Questions:

- Which topics are you mostly interested in?
- Which are of less interest?
- Do you miss any topic?

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	KMS
address spaces, threads, process states	KMS
monitor	KMS
process, concurrency, parallelism, threads, synchronization, monitors in Java	GP, KMS GP, KMS GP, KMS
verification of properties of programs	Modellierung

Lecture Parallel Programming WS 2014/2015 / Slide 04

Objectives:

Sources for prerequisites

In the lecture:

- Explanations.
- The notions will be briefly repeated in the first chapter of this lecture - as introduced in GP and KMS

Suggested reading:

Relevant sections of lecture material of

- Grundlagen der Programmierung 1, 2,
- Konzepte und Methoden der Systemsoftware

Questions:

- Did you attend those lectures?
- Are you going to learn or to repeat those topics?

Organization of the course

PPJ-5

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.

Lecture Parallel Programming WS 2014/2015 / Slide 05

Objectives:

Introduce the form of the material.

In the lecture:

- Explain the organization of the material.

Questions:

- Did you already explore the material?
- Did you place bookmarks into it?

Literature

PPJ-6

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

Lecture Parallel Programming WS 2014/2015 / Slide 06

Objectives:

Reference to books

In the lecture:

Explain

- Andrews' book treats the concepts very thoroughly - deeper than we can do it in this lecture.
- The 3 books on Java present programming techniques very extensively with many elaborated examples; in some parts orientation is a bit missing.

Questions:

Are you going to dive into the matter along those books?


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

Lecture Parallel Programming WS 2014/2015 / Slide 07

Objectives:

Repeat fundamental notions of processes

In the lecture:

- The notions are explained.
- Interleaved execution is also used as a model for describing properties of a system of processes.

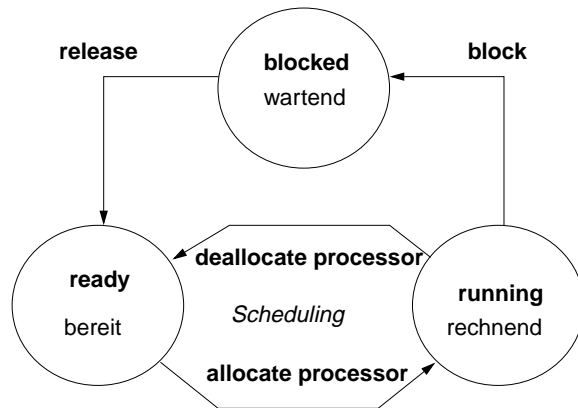
Suggested reading:

SWE-131

Questions:

- Which are the situations when a process switch is performed?

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.

Lecture Parallel Programming WS 2014/2015 / Slide 08

Objectives:

Understand process switching

In the lecture:

- Explain states and transitions.
- Role of the scheduler.

Questions:

- Give reasons and examples for state transitions.

Applications of parallel processes

PPJ-9

- **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**.

© 2015 bei Prof. Dr. Uwe Kastens

Lecture Parallel Programming WS 2014/2015 / Slide 09

Objectives:

recognize different goals of parallelism

In the lecture:

Example are used to explain the classes of applications

Suggested reading:

SWE-132

Questions:

- Give further examples for the use of parallel processes, and point out their category.

Create threads in Java - technique: implement Runnable

PPJ-10

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`

© 2005 bei Prof. Dr. Uwe Kastens

Lecture Parallel Programming WS 2014/2015 / Slide 10

Objectives:

Understand declaration of process classes

In the lecture:

3 development steps:

- declare the class with its run method
- create a process object
- start the execution of the process object

If the user's class would need further object methods, they would be difficult to access. In that case one should better apply the second technique for process class declaration.

Suggested reading:

SWE-133

Questions:

- The class `Thread` has class methods and object methods. Which can be called from the run method in which way?

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!

Lecture Parallel Programming WS 2014/2015 / Slide 11

Objectives:

Understand declaration of process classes

In the lecture:

3 development steps:

- declare the class with its run method
- create a process object
- start the execution of the process object

Compare to the variant with the interface Runnable.

Suggested reading:

SWE-134

Questions:

- The class Thread has class methods and object methods. Which can be called from the run method in which way?

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.

Lecture Parallel Programming WS 2014/2015 / Slide 12

Objectives:

Overview over the Thread methods

In the lecture:

- Explain the methods.
- Demonstrate the execution of the involved processes graphically.
- Point to examples.

Suggested reading:

SWE-137

Assignments:

Demonstrate the execution of the methods graphically.


Questions:

- Which method calls involve two processes, which only one?

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

PPJ-13

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



```
class DigiClock extends Thread
{ public void run ()
  { while (running)           iterate until it is terminated from the outside
    { line.setText(new Date().toString());           write the date
      try { sleep (1000); } catch (Exception ex) {}           pause
    }
  }

  Method, that terminates the process from the outside:
  public void stopIt () { running = false; }
  private volatile boolean running = true;           state variable

  public DigiClock (Label t) {line = t;}           label to be used for the text
  private Label line;
}
```

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**.

© 2011 bei Prof. Dr. Uwe Kastens

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

PPJ-14

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);           create process
    clock.start();                               start process
  }

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

  private DigiClock clock;
}
```

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.

© 2008 bei Prof. Dr. Uwe Kastens

Lecture Parallel Programming WS 2014/2015 / Slide 13

Objectives:

A first complete example

In the lecture:

Explanation of

- the execution until termination from the outside.
- the `stopIt` method,
- the reason for the variant "subclass of `Thread`".

Demonstrate the applet [Digital Clock Process](#)

Suggested reading:

SWE-135

Assignments:

Install the example and modify it.

Lecture Parallel Programming WS 2014/2015 / Slide 14

Objectives:

Start a process from an applet

In the lecture:

Explain how to start, suspend, resume, and terminate a process from an applet.

Suggested reading:

SWE-136

Assignments:

Modify the classes of this example such that `DigiClock` implements `Runnable` instead of being a subclass of `Thread`.

Questions:

- Explain why `DigiClock` extends `Thread` in the presented version.

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.

Lecture Parallel Programming WS 2014/2015 / Slide 15a

Objectives:

Recognize the necessity for formal reasoning

In the lecture:

Explain the goals and the examples.

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**?

Lecture Parallel Programming WS 2014/2015 / Slide 15ab

Objectives:

How to proof parallel programs

In the lecture:

Check the correctness of proofs

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:

$$\{P\} S_1 \{Q\} S_2 \{R\}$$

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:

$\{true\}$ characterizing all states, and $\{false\}$ 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.

Lecture Parallel Programming WS 2014/2015 / Slide 15b

Objectives:

Recall the fundamental notions of Hoare logic

In the lecture:

The notions are explained. (see lecture material "Modellierung", slides Mod-4.51 to Mod-4.68)

Axioms and inference rules for sequential constructs

statement sequence

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

1

stronger precondition

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

3

weaker postcondition

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

4

assignment

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

2

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

multiple alternative (guarded command)

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

5

selecting iteration

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

6

no operation

$$\{P\} \text{ skip } \{P\}$$

7

Lecture Parallel Programming WS 2014/2015 / Slide 15c

Objectives:

Understand the inference rules

In the lecture:

The rules are explained:

- 1, 2, 3 are explained in "Modellierung" Mod-4.57 to MOD-4.60.
- guarded commands and iteration are generalized form of those explained in Mod-4.61 to Mod-4.66b.
- skip is clear.

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 ->
  { INV  $\wedge$   $a \neq b$  }
  if a > b ->
    { 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 <= b ->
    { 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 -> INV“
{ INV }
od
{ INV  $\wedge$   $a = b$  }  $\rightarrow$ 
{ a = G }
  
```

the loop terminates:

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

Lecture Parallel Programming WS 2014/2015 / Slide 15d

Objectives:

Example for application of inference rules

In the lecture:

The verification steps are explained in "Modellierung" slide and PDF-file Mod-4.68

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.

Lecture Parallel Programming WS 2014/2015 / Slide 15e

Objectives:

Understand the notion of weakest precondition

In the lecture:

The notion is explained using some examples.

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

Lecture Parallel Programming WS 2014/2015 / Slide 15f

Objectives:

Learn to find WPs

In the lecture:

The topics of the slide are explained:

- The formula.
- If $i \leq 0$, then the execution of $i := i + 1$ terminates with $i \leq 1$, while if $i > 0$ the execution of S cannot make $i \leq 1$.
- Execution of S always sets z to $\max(x, y)$.
- Execution of S beginning with $y \geq x$; sets z to y and execution of S beginning with $y < x$ sets z to x , which is $\max(x, y)$.
- There is no start state for S such that it can set z less than y .
- Only if $x = y + 1$ holds when execution of S begins, it will set z to $y + 1$.
- Clear.

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**:

```

account = 50;
      a           b           c
-----
Process1: t1 = account; t1 = t1 + 10; account = t1;
      d           e           f
-----
Process2: t2 = account; t2 = t2 - 5; account = t2;

```

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.

Lecture Parallel Programming WS 2014/2015 / Slide 17a

Objectives:

Motivation of the execution model

In the lecture:

- Explain the notion of atomic operations.
- Scheduling strategies are discussed later.
- Processes interfere via common, global variables.
- The desired result of a program execution may not depend on unjustified assumptions on the interleaving.
- Check all results the example may yield.

Questions:

- Which results may the example yield?
- Declare atomic statement sequences such that any interleaved execution yields the same result.

Atomic actions

PPJ-17b

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.

© 2008 bei Prof. Dr. Uwe Kastens

Lecture Parallel Programming WS 2014/2015 / Slide 17b

Objectives:

Notion of atomic actions in the interleaving model

In the lecture:

- Explanation and examples for AMO.
- The example of PPJ-15 is varied.

Questions:

- Explain the AMO property using the terms "observable states" and "interleaved execution".

Atomic by AMO

PPJ-17c

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 $\langle \dots \rangle$:

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

interleaving actions of two processes:

(1)	<p style="text-align: center; margin: 0;">a</p> <p>p1: $\langle t = 0; t = t + 1; \rangle$</p> <p>p2: $\langle s = 0; s = s + 1; \rangle$</p> <p style="text-align: center; margin: 0;">b</p>	(2)
(3)	<p style="text-align: center; margin: 0;">b a</p> <p>p1: $x = \langle y + 1 \rangle;$</p> <p>p2: $y = \langle x + 1 \rangle;$</p> <p style="text-align: center; margin: 0;">d c</p>	(4)
	<p style="text-align: center; margin: 0;">c a b</p> <p>p1: $x = \langle y \rangle + \langle z \rangle;$</p> <p>p2: $\langle y = 1; \rangle \langle z = 2; \rangle;$</p> <p style="text-align: center; margin: 0;">d e</p>	

© 2005 bei Prof. Dr. Uwe Kastens

Lecture Parallel Programming WS 2014/2015 / Slide 17c

Objectives:

Understand: AMO constructs can be considered atomic

In the lecture:


The examples are explained using the definition of AMO.

Questions:

Which states can the processes in (1) to (4) reach depending on the execution order of the atomic actions?

Non-interference checks

```
x := 0; y := 0;
{ x = 0 ∧ y = 0 }
co {x+1 = 1} x := x+1 {x=1} //
```




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

```
oc
{ x = 1 ∧ y = 1 } => {x+y = 2}
z := x+y
{z = 2}
```

```
{ x+1 = 1 ∧ y+1 = 1 } y:= y+1 {x+1 = 1 ∧ y = 1} =>
c {x+1 = 1}
```

```
x := 0; y := 0;
{ x = 0 ∧ y = 0 }
co {y+1 = 1} x := y+1 {x=1} //
```



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

```
oc
{ x = 1 ∧ y = 1 } => {x+y = 2}
z := x+y
{z = 2}
```

```
{ y+1 = 1 ∧ x+1 = 1 } y:= x+1 {y+1 = 1}
is not correct
```

is not correct

Objectives:

Apply interference checks

In the lecture:

The interference checks of the examples are explained.

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.

$$\frac{\{P_i\} S_i \{Q_i\}, 1 \leq i \leq n, \text{ are interference-free theorems}}{\{P_1 \wedge \dots \wedge P_n\} \text{co } S_1 // \dots // S_n \text{ oc } \{Q_1 \wedge \dots \wedge Q_n\}}$$

Objectives:

Understand the inference rules

In the lecture:

The two statements and their inference rules are explained.

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 $x:=e$

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

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

$q:: \dots S1 \{C\} S2 \dots$

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

$q:: \dots S1 \{C\} S2 \dots$

Objectives:

Techniques to reduce interference

In the lecture:

The techniques are explained using small examples.

- (4): Show that the await statement causes NI($x:=e, C$) to hold.