


Concurrency at Programming Language Level

Ch 2 [BenA 06]



Abstraction
Pseudo-language
BACI
Ada, Java, etc.

Levels of Abstraction

- Granularity of operations
 - Invoke a library module
 - Statement in high level programming language
 - Instruction in machine language
- Atomic statement
 - Anything that we can guarantee to be atomic
 - Executed completely “at once”
 - Always the same correct atomic result
 - Result does not depend on anybody else
 - Can be at any granularity
 - Can *trust* on that atomicity

Atomic Statement

- Atomicity guaranteed somehow

- Machine instruction: HW

- Memory bus transaction

Load R1, Y

Read mem(0x35FA8300)

- Programming language statement, set of statements, or set of machine instructions

- SW

- Manually coded
- Disable interrupts
- OS synchronization primitives



```
-- start atomic
Load R1, Y
Sub R1, =1
Jpos R1, Here
-- end atomic
```

- Library module

- SW

- Manually coded inside
- Provided automatically to the user by programming environment

Monitors
Ch 7 [BenA 06]

Concurrent Program


- Sequential process
 - Successive atomic statements

P: $p1 \rightarrow p2 \rightarrow p3 \rightarrow p4 \dots$

- Control pointer (= program counter)

- Concurrent program
 - Finite set of sequential processes working for same goal
 - Arbitrary interleaving of atomic statements in different processes

3 processes (P, R, Q) interleaved execution



P: $p1 \rightarrow p2$

Q: $q1 \rightarrow q2$

$p1 \rightarrow q1 \rightarrow p2 \rightarrow q2,$
 $p1 \rightarrow q1 \rightarrow q2 \rightarrow p2,$
 $p1 \rightarrow p2 \rightarrow q1 \rightarrow q2,$
 $q1 \rightarrow p1 \rightarrow q2 \rightarrow p2,$
 $q1 \rightarrow p1 \rightarrow p2 \rightarrow q2,$
 $q1 \rightarrow q2 \rightarrow p1 \rightarrow p2.$


?

~~$p1 \rightarrow q2 \rightarrow p2 \rightarrow q1$~~

Program State, Pseudo-language

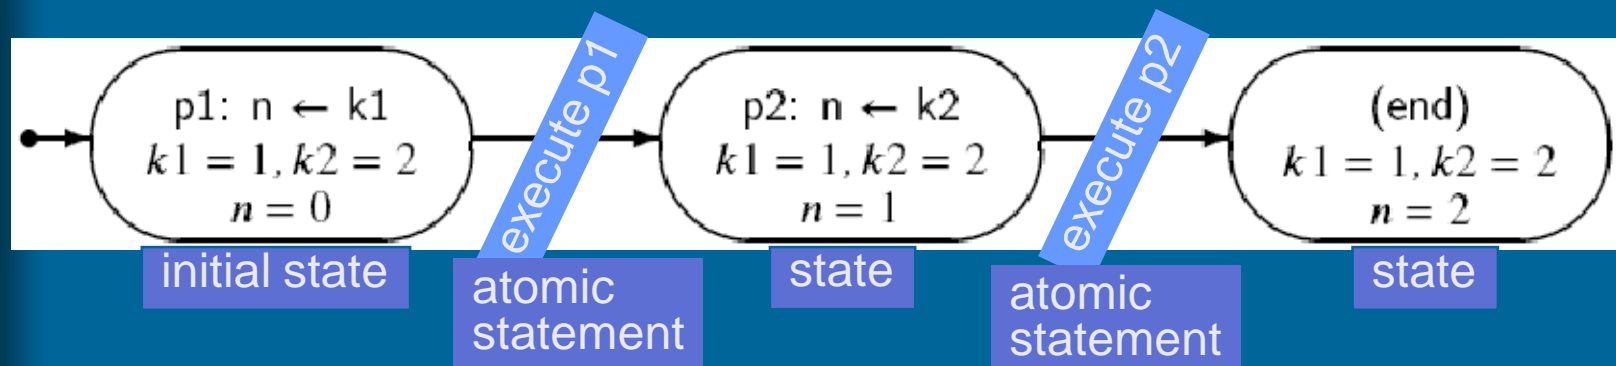
pseudo-kieli

- Sequential program



- State

- next statement to execute (cp, i.e., PC)
- variable values



(Global) Program State

- Concurrent program

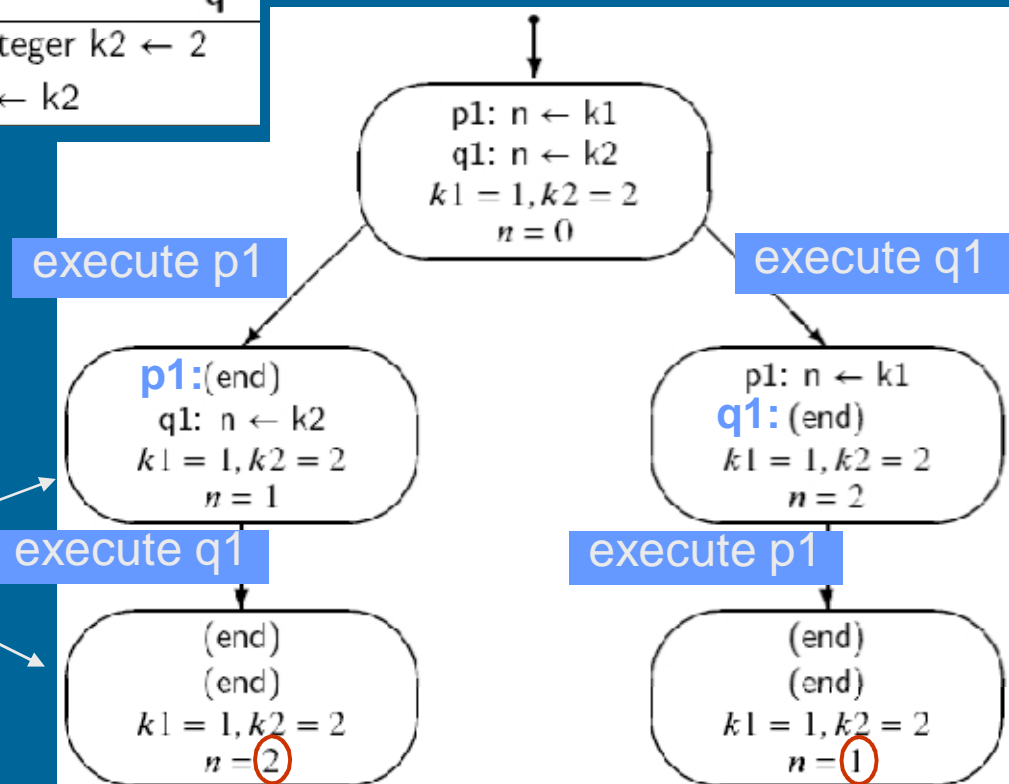
Algorithm 2.1: Trivial concurrent program	
integer $n \leftarrow 0$	
p	q
integer $k1 \leftarrow 1$	integer $k2 \leftarrow 2$
p1: $n \leftarrow k1$	q1: $n \leftarrow k2$

- Local state for each *process*:

- cp
- Variable values
 - Local & global

- Global state for *program*

- All cp's
- All local variables
- All global variables



Possible Program States

- List of processes in program

- List of values for each process

- cp

- value of each local/global/shared variable

```
p1: n ← k1
q1: n ← k2
k1 = 1, k2 = 2
n = 0
```

```
state: { { p1: n ← k1  – process p
          k1 = 1 }
         { q1: n ← k2  – process q
          k2 = 2 }
         n = 0    – shared variable
       }
```

- Nr of possible states

can be (very) large

- Not all states are reachable states!

(saavutettavissa, saavutettava tila)

- Different executions do not go through same states (even with same input)

unreachable

state:

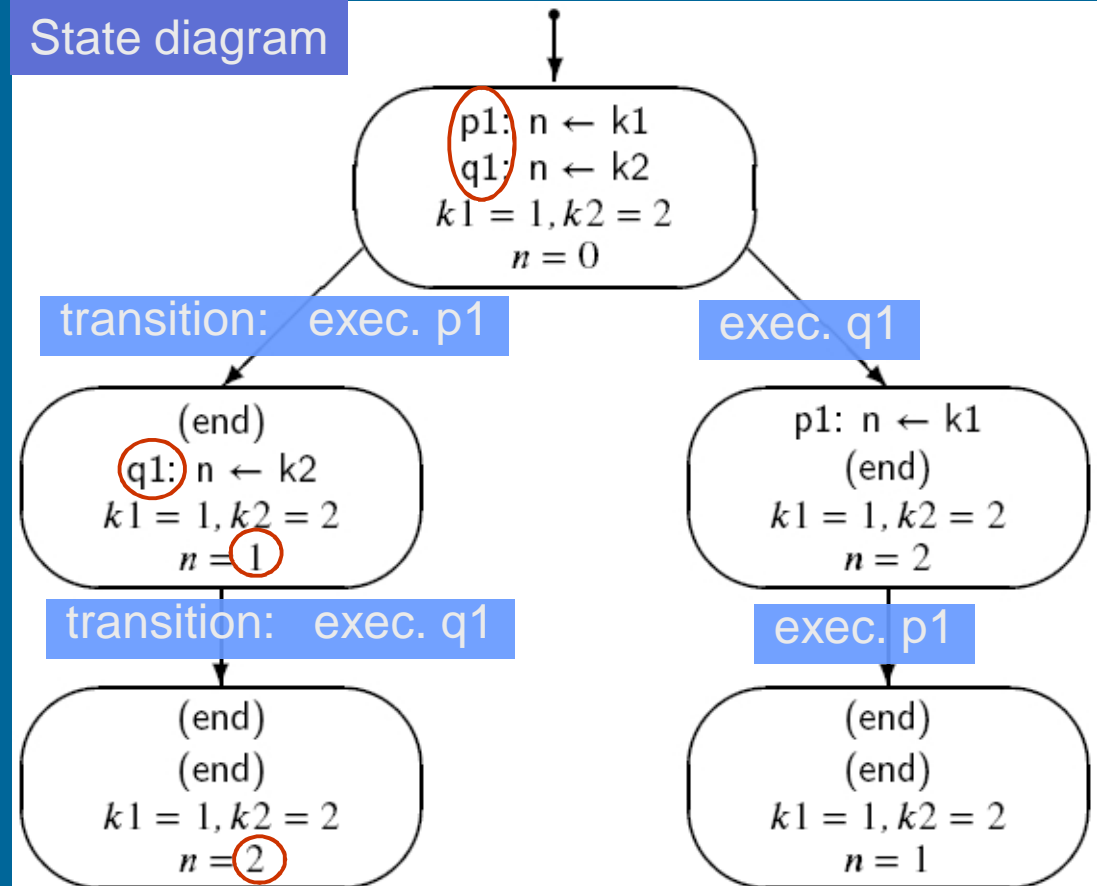
```
{ { p1: n ← k1
   k1 = 2 }
  { q1: n ← k2
   k2 = 1 }
  n = 3
}
```

State Diagram and Scenarios

Process p	Process q	n	k1	k2
p1: n ← k1	q1: n ← k2	0	1	2
(end)	q1: n ← k2	1	1	2
(end)	(end)	2	1	2

Scenario 1 (left side)

State diagram



- Transitions from one possible state to another
 - Executed statement must be one of those in the 1st state
- State diagram for concurrent program
 - Contains all reachable states and transitions
 - All possible executions are included, they are all correct!

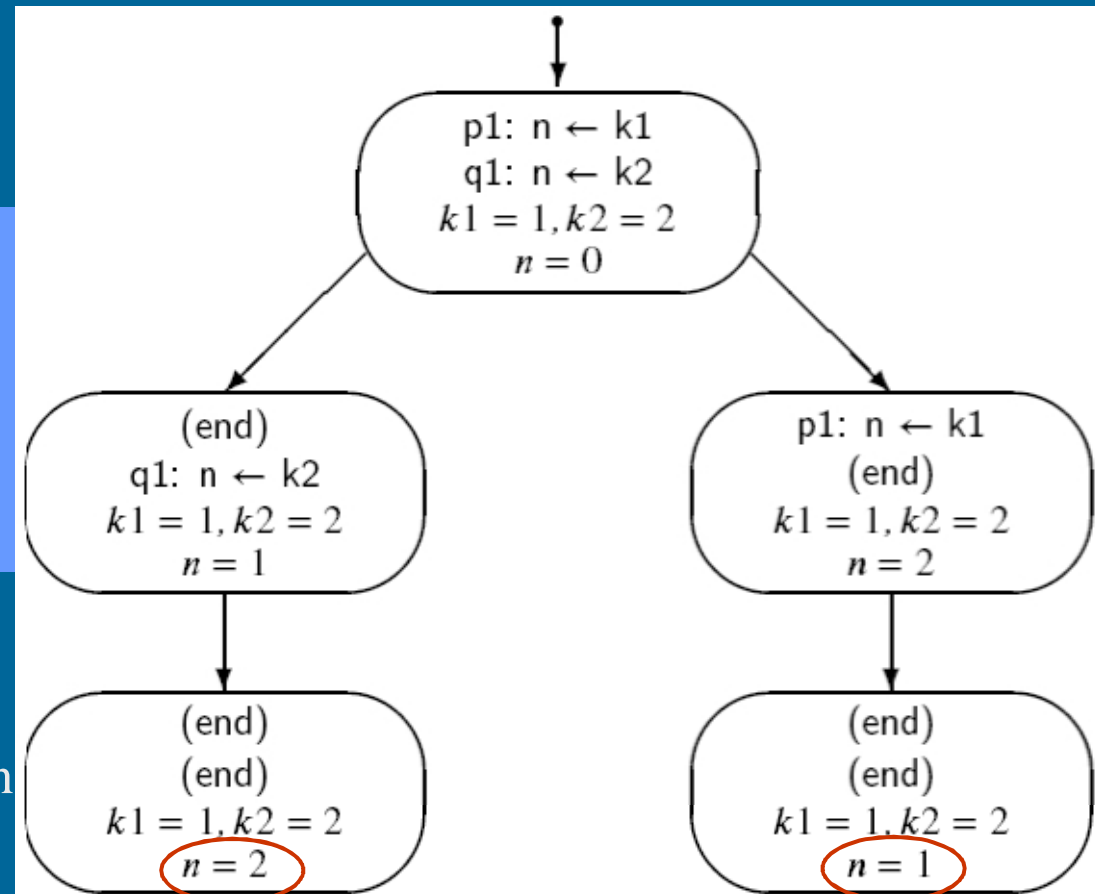
Atomic Statements

Algorithm 2.1: Trivial concurrent program	
integer $n \leftarrow 0$	
p	q
integer $k1 \leftarrow 1$ $p1: n \leftarrow k1$	integer $k2 \leftarrow 2$ $q1: n \leftarrow k2$

- Two scenarios
 - Both correct
 - Different result!

NO need to have the same result!
Statements do the same, but overall result may be different.
(see p. 19 [BenA 06])

- Atomic?
 - Assignment?
 - Boolean evaluation
 - Increment?



Algorithm 2.3: Atomic assignment statements

integer $n \leftarrow 0$

p	q
$p1: n \leftarrow n + 1$	$q1: n \leftarrow n + 1$

- Two scenarios for execution
 - Both correct
 - Both have the same result

P first, and then Q

Process p	Process q	n
p1: $n \leftarrow n + 1$	$q1: n \leftarrow n + 1$	0
(end)	q1: $n \leftarrow n + 1$	1
(end)	(end)	2

Q first, and then P

Process p	Process q	n
$p1: n \leftarrow n + 1$	q1: $n \leftarrow n + 1$	0
p1: $n \leftarrow n + 1$	(end)	1
(end)	(end)	2

Algorithm 2.3: Atomic assignment statements

integer $n \leftarrow 0$

p	q
p1: $n \leftarrow n + 1$	q1: $n \leftarrow n + 1$

Same statements with smaller atomic granularity:

Algorithm 2.4: Assignment statements with one global reference

integer $n \leftarrow 0$

p	q
integer temp	integer temp
p1: $temp \leftarrow n$	q1: $temp \leftarrow n$
p2: $n \leftarrow temp + 1$	q2: $n \leftarrow temp + 1$

Too Small Atomic Granularity

Algorithm 2.4: Assignment statements with one global reference

integer n ← 0	
p	q
integer temp p1: temp ← n p2: n ← temp + 1	integer temp q1: temp ← n q2: n ← temp + 1

- Scenario 1 →
– OK
- Scenario 2 →
– Bad result
- From now on
– Assignments and Boolean evaluations are atomic!

Process p	Process q	n	p.temp	q.temp
p1: temp ← n	q1: temp ← n	0	?	?
p2: n ← temp + 1	q1: temp ← n	0	0	?
(end)	q1: temp ← n	1	0	?
(end)	q2: n ← temp + 1	1	0	1
(end)	(end)	2	0	1

Process p	Process q	n	p.temp	q.temp
p1: temp ← n	q1: temp ← n	0	?	?
p2: n ← temp + 1	q1: temp ← n	0	0	?
p2: n ← temp + 1	q2: n ← temp + 1	0	0	0
(end)	q2: n ← temp + 1	1	0	0
(end)	(end)	1	0	0

Correctness

- What is the correct answer?
- Usually clear for sequential programs
- Can be fuzzy for concurrent programs
 - Many correct answers?
 - What is intended semantics of the program?
 - Run programs 100 times, each time get different answer?
 - Each answer is correct, if program is correct!
 - Does not make debugging easier!
 - Usually can not test all possible scenarios (too many!)
 - How to define correctness for concurrent programs?
 - Safety properties = properties that are always true
 - Liveness properties = properties that eventually become true

“turvallisuus”

“elävyys”

Safety and Liveness

- Safety property safety-ominaisuus, turvallisuus
 - property must be true all the time (“bad” never happens)
 - “Identity” identiteetti, invariantti
 - $\text{memFree} + \text{memAllocated} = \text{memTotal}$
 - Mouse cursor is always displayed
 - System responds always to new commands
- Liveness property elävyys, liveness-ominaisuus
 - Property must eventually become true (“good” eventually happens)
 - Variable n value = 2
 - System prompt for next command is shown
 - Control will resume to calling program
 - Philosopher will get his turn to eat
 - Eventually the mouse cursor is not displayed
 - Program will terminate
- Duality of safety and liveness properties
 - $\{ P_i \text{ will get his turn to eat} \} \equiv \text{not } \{ P_i \text{ will never get his turn to eat} \}$
 - $\{ n \text{ value will become } 2 \} \equiv \text{not } \{ n \text{ value is always } \neq 2 \}$

Linear Temporal Logic (LTL)

(lineaarinen) temporaalilogiikka

- Define safety and liveness properties for certain state in some (arbitrary) scenario
 - Example of Modal Temporal Logic (MDL), logic on concepts like possibility, impossibility, and necessity
- **Alternative: Branching Temporal Logic (BTL)**
 - Properties true in some or all states starting from the given state
 - More complex, because all future states must be covered
 - **Common Temporal Logic (CTL)**
 - Can be checked automatically
 - Every time computation reaches given state
 - SMV model checker
 - NuSMV model checker

Fairness

reiluus

- (Weakly) fair scenario
 - Wanted condition eventually occurs
 - Nobody is locked out forever?
 - Will a philosopher ever get his turn to eat?
 - Will an algorithm eventually stop?
 - p and q are both scheduled to run eventually

Algorithm 2.5: Stop the loop A	
integer $n \leftarrow 0$ boolean flag \leftarrow false	
p	q
p1: while flag = false p2: $n \leftarrow 1 - n$	q1: flag \leftarrow true q2:

- All scenarios should be fair
 - One requirement in correct solution

Machine Language Code

- What is atomic and what is not?
 - Assignment? `X = Y;`
 - Increment? `X = X+1;`

Algorithm 2.6: Assignment statement for a register machine

integer $n \leftarrow 0$

p	q
p1: load R1,n	q1: load R1,n
p2: add R1,#1	q2: add R1,#1
p3: store R1,n	q3: store R1,n


Critical Reference

kriittinen viite

- Reference to variable v is critical reference, if ...

- Assigned value in P and read in Q

- Read directly or in a statement



- Program satisfies limited-critical-reference (LCR)

- Each statement has at most one critical reference

- Easier to analyze than without this property

- Each program is easy to transform into similar program with LCR

rajoitettu

kriittinen viite

	P	Q	
Not LCR:	$\underline{n} = \underline{n}+1;$	$\underline{n} = \underline{n}+1$	Bad
Not LCR:	$\underline{n} = \underline{m}+1;$	$\underline{m} = \underline{n}+1$	Bad
LCR:	$\text{tempP} = \underline{n}+1;$ $\underline{n} = \text{tempP};$	$\text{tempQ} = \underline{n}+1;$ $\underline{n} = \text{tempQ};$	Good

LCR vs. atomicity?
(ouch)

Volatile and non-atomic variables

riskialtis

- Volatile variable
 - Can be modified by many processes (must be in shared memory)
 - Advice for compiler (pragma)
 - Keep something in memory, not in register
 - Pseudocode – does not generate code
- Non-atomic variables
 - Multiword data structures: long ints, arrays, records, ...
 - Force access to be indivisible (atomic) in given order

What if compiler/hw decides to keep value of n in a register/cache?
 When is it stored back to memory?
 What if local1 & local2 were volatile?

Algorithm 2.8: Volatile variables	
integer $n \leftarrow 0$	
p	q
integer local1, local2	integer local
p1: $n \leftarrow \text{some expression}$	q1: local $\leftarrow n + 6$
p2: <i>computation not using n</i>	q2:
p3: local1 $\leftarrow (n + 5) * 7$	q3: which n?
p4: local2 $\leftarrow n + 5$	q4:
p5: $n \leftarrow \text{local1} * \text{local2}$	q5:

Annotations: Red arrows and text indicate execution order and storage questions. An arrow points from p1 to q1 with the text "store n?". Another arrow points from p3 to q1 with the text "exec. order?". A third arrow points from p5 to q1 with the text "store n?".

Example Program with Volatile Variables

Algorithm 2.9: Concurrent counting algorithm	
integer n \leftarrow 0	
p	q
integer temp	integer temp
p1: do 10 times	q1: do 10 times
p2: temp \leftarrow n	q2: temp \leftarrow n
p3: n \leftarrow temp + 1	q3: n \leftarrow temp + 1

- Can implement it in any concurrent programming language
 - (Extended) Pascal and (Extended) C
 - BACI (Ben-Ari Concurrency Interpreter)
 - Code automatically compiled (from Extended Pascal or C)
 - Ada
 - Java

Discuss

Concurrent Program in Pascal

possibly volatile

```
1 program count;
2 var n: integer := 0;
3
4 procedure p;
5 var temp, i: integer;
6 begin
7   for i := 1 to 10 do
8     begin
9       temp := n;
10      n := temp + 1;
11    end
12 end;
```

n is volatile, because... it is assigned in one thread, and read in the other

```
16 procedure q;
17 var temp, i: integer;
18 begin
19   for i := 1 to 10 do
20     begin
21       temp := n;
22       n := temp + 1;
23     end
24 end;
25
26 begin { main program }
27 cobegin p; q coend;
28 writeln('The value of n is ', n);
29 end.
```

Concurrent Program in C (Ben-Ari Concurrent C, C--)

```
1  int n = 0;
2
3  void p() {
4      int temp, i;
5      for (i = 0; i < 10; i++) {
6          temp = n;
7          n = temp + 1;
8      }
9  }
10
```

possibly volatile, use carefully

(volatile, if critically referenced)

```
16 void q() {
17     int temp, i;
18     for (i = 0; i < 10; i++) {
19         temp = n;
20         n = temp + 1;
21     }
22 }
23
24 void main() {
25     cobegin { p(); q(); }
26     cout << "The value of n is " << n << "\n";
27 }
```

What if compiler optimized and kept n in a register?
Let's hope not!
(in ExtPascal or C--
global (volatile) variables are seemingly kept in memory by default)

Concurrent Program in Ada

```
1 with Ada.Text_IO; use Ada.Text_IO;
2 procedure Count is
3     N: Integer := 0;
4     pragma Volatile(N);
5
6     task type Count_Task;
7     task body Count_Task is
8         Temp: Integer;
9     begin
10        for I in 1..10 loop
11            Temp := N;
12            N := Temp + 1;
13        end loop;
14    end Count_Task;
15
```

advice compiler to keep N in memory

```
16 begin
17     declare
18         P, Q: Count_Task;
19     begin
20         null;
21     end;
22     Put_Line("The value of N is " & Integer'Image(N));
23 end Count;
```

Concurrent Program in Java

```
1 class Count extends Thread {
2     static volatile int n = 0;
3
4     public void run() {
5         int temp;
6         for (int i = 0; i < 10; i++) {
7             temp = n;
8             n = temp + 1;
9         }
10    }
```

Thread.yield(); // force?

```
> javac Adder8.java
> java Adder8
```

```
16     public static void main(String[] args) {
17         Count p = new Count();
18         Count q = new Count();
19         p.start ();
20         q.start ();
```

How many threads
really in parallel?
• how to control it?

```
21     try {
22         p.join ();
23         q.join ();
24     }
25     catch (InterruptedException e) { }
26     System.out.println ("The value of n is " + n);
27 }
28 }
```

Execute on 8-processor vera.cs.helsinki.fi?

<http://www.cs.helsinki.fi/u/kerola/rio/Java/examples/Adder8.java>

<http://www.cs.helsinki.fi/u/kerola/rio/Java/examples/Adder8b.java>

Run Multi-threaded Java

Execute on 8-processor vera.cs.helsinki.fi?

<http://www.cs.helsinki.fi/u/kerola/rio/Java/examples/Adder8.java>

```
kerola@vera:~/public_html/rio/Java/examples$ javac  
Adder8.java
```

```
kerola@vera:~/public_html/rio/Java/examples$ java Adder8
```

```
finally n = 80000 = 37358
```

```
kerola@vera:~/public_html/rio/Java/examples$ java Adder8
```

```
finally n = 80000 = 34464
```

- Why different result?
- What is correct result?

Run them your self?
(Copy source code in
your own directory)

BACI

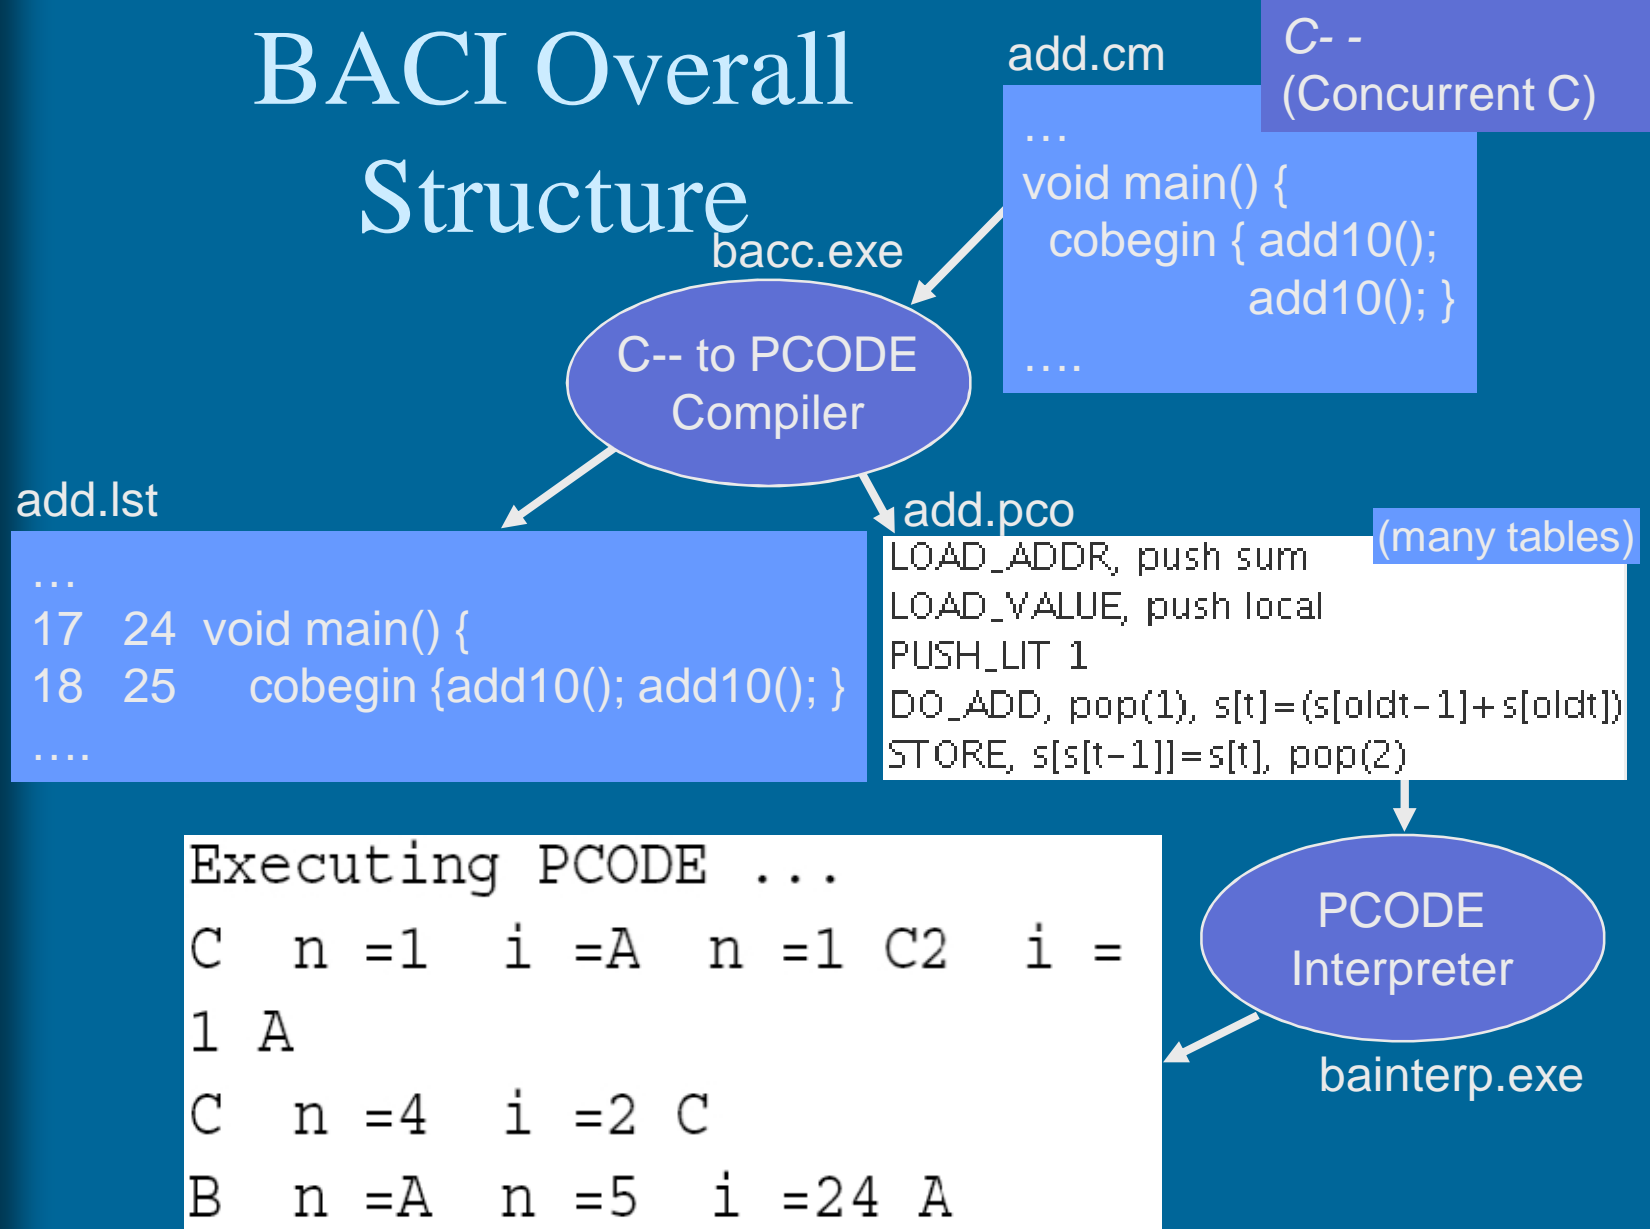
<http://inside.mines.edu/~tcamp/baci/baci.html>

- Ben-Ari Concurrency Interpreter
 - Write concurrent programs with
 - *C--* or *Ben-Ari Concurrent Pascal* (.cm and .pm suffixes)
 - Compile and run in BACI
 - GUI for Unix/Linux
- jBACI <http://stwww.weizmann.ac.il/g-cs/benari/jbaci/>
 - Just like BACI
 - GUI for Windows
- Installation <http://stwww.weizmann.ac.il/g-cs/benari/jbaci/jbaci1-4-5.zip>
 - load version 1.4.5 jBACI executable files and example programs, unzip, edit config.cfg to have correct paths to bin/bacc.exe and bin/bapas.exe translators, click run.bat
- Use in class, homeworks and in project

See BACI instructions

<http://www.cs.helsinki.fi/u/kerola/rio/ohjeet/ohjeet.html>

BACI Overall Structure



<http://www.cs.helsinki.fi/u/kerola/rio/BACI/baci-c.pdf>

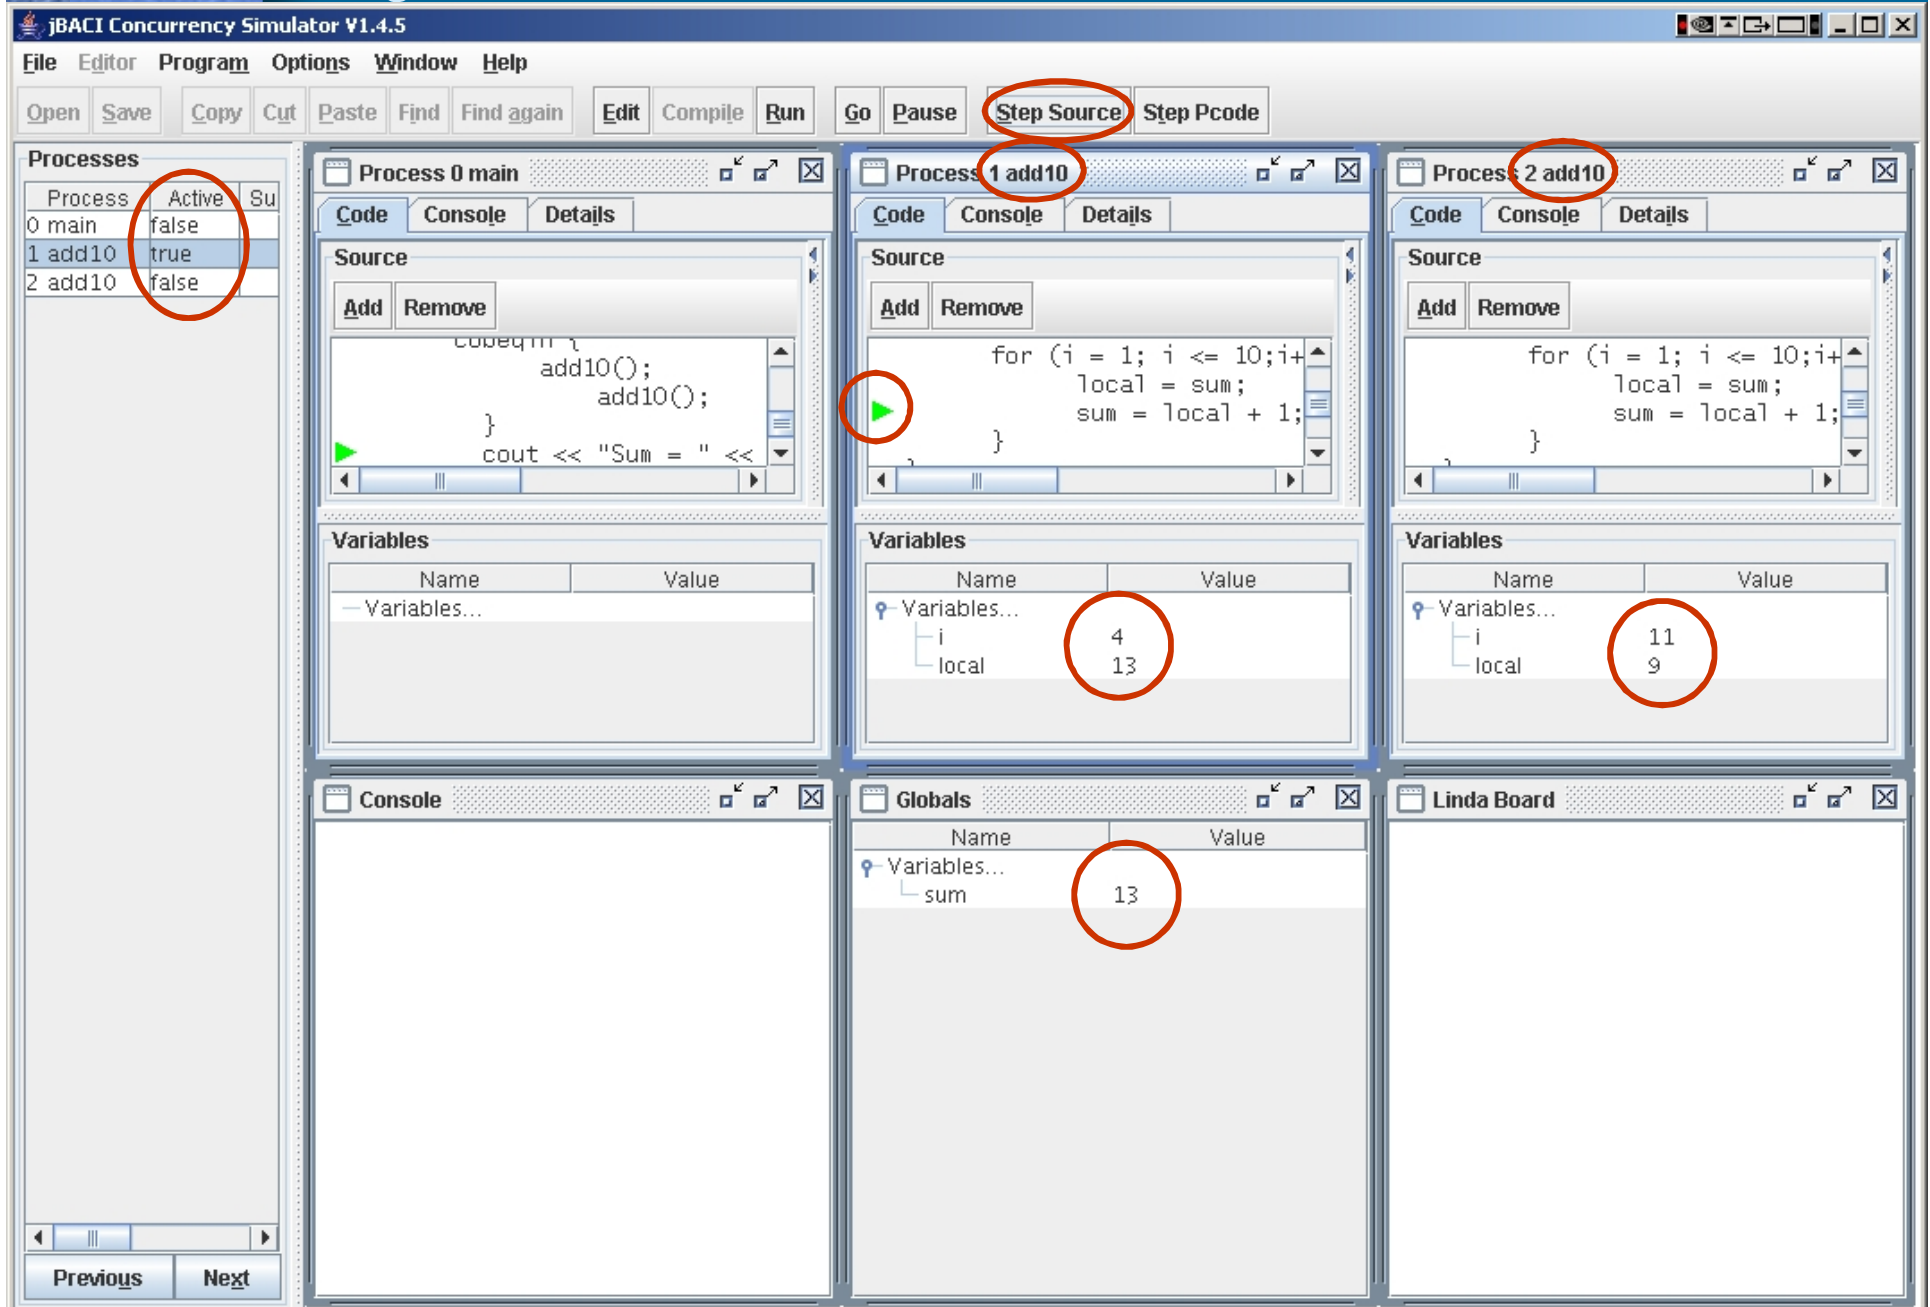
jBACI

Just like
BACI, but
with Java

- requires
Java v. 1.4
(SDK or
JRE)
- Built-in
compiler
and
interpreter
- edit state
- run state

<http://www.cs.helsinki.fi/u/kerola/rio/BACI/jbaci.pdf>

jBACI IDE (integrated development environment)



The screenshot displays the jBACI IDE interface with three process windows and a console window. The 'Step Source' button in the toolbar is circled in red. The 'Processes' list on the left has the 'Active' column circled in red, with 'Process 1 add10' selected. The source code for Process 1 is shown with a green play button circled in red. The variables for Process 1 are shown with 'i' and 'local' values circled in red. The variables for Process 2 are shown with 'i' and 'local' values circled in red. The console window is empty. The Globals window shows the 'sum' variable with a value of 13 circled in red.

Processes List:

Process	Active	Sum
0 main	false	
1 add10	true	
2 add10	false	

Process 0 main Source:

```
codeqmqm {
    add10();
    add10();
}
cout << "Sum = " <<
```

Process 1 add10 Source:

```
for (i = 1; i <= 10; i++)
    local = sum;
    sum = local + 1;
}
```

Process 1 Variables:

Name	Value
Variables...	
i	4
local	13

Process 2 add10 Source:

```
for (i = 1; i <= 10; i++)
    local = sum;
    sum = local + 1;
}
```

Process 2 Variables:

Name	Value
Variables...	
i	11
local	9

Console:

Globals:

Name	Value
Variables...	
sum	13

Linda Board:

jBACI IDE (integrated development environment)

The screenshot displays the jBACI IDE interface with several key components:

- Process List:** A table showing the status of three processes. Process 2 is highlighted with a red circle around its 'Finish' column.
- Code Editor:** Shows the source code for 'Process 1 add10'. A red circle highlights the 'Add' button for breakpoints, and another red circle highlights the assembly code in the PCode window.
- Stack Window:** Shows the stack for 'Process 0 main' with a red circle around the 'Details' tab.
- History Window:** Shows a list of recent instructions with a red circle around the 'History' title bar.
- Process Control:** Buttons for 'Process 0 m...', 'Process 2 ad...', and 'Globals' are visible at the bottom.

Process List:

Process	Active	Suspend	Finish	Monitor	Pri	Atm
0 main	false		false		0	0
1 add10	true		false		0	0
2 add10	false		true		0	0

Code Editor (Process 1 add10):

```
int sum = 0;

void add10() {
    int i;
    int local;
    for (i = 1; i <= 10; i++) {
        local = sum;
        sum = local + 1;
    }
}

void main() {
    cobegin {
        add10();
        add10();
    }
}
```

PCode (Process 1 add10):

```
LOAD_ADDR, push local
LOAD_VALUE, push sum
STORE, s[s[t-1]]=s[t], pop(2)
```

Stack (Process 0 main):

Index	Value
0	16
1	10
2	20
3	Address[0,0]
4	Address[0,-1]
5	6

History (Most recent instructions):

Proc	Line	Source	File
1	12	sum = local + 1;	add.cm
1	13	}	add.cm
1	10	for (i = 1; i <= 10; i++)	add.cm
1	11	local = sum;	add.cm
1	12	sum = local + 1;	add.cm
1	13	}	add.cm
1	10	for (i = 1; i <= 10; i++)	add.cm
1	11	local = sum;	add.cm

Process Control:

- Process 0 m...
- Process 2 ad...
- Globals
- Linda Board
- Console

Buttons: Previous, Next

Summary

- Abstraction, atomicity
- Concurrent program, program state
- Pseudo-language algorithms
- High level language algorithms
- BACI