



Non-determinism

Uwe R. Zimmer - The Australian National University

Non-determinism

Non-determinism by design

Dijkstra's guarded commands (non-deterministic case statements):

```

if
  x < y >> m := x
  x > y >> m := y
fi
  
```

Selection is non-deterministic (forky)

The programmer needs to design the alternatives as 'parallel' options: all cases need to be covered and overlapping conditions need to lead to the same result

All true case statements in any language are potentially concurrent and non-deterministic.

Numerical non-determinism in concurrent statements (Chapal):

```

wait50 (c reduce [1 in {0..99999}]);
wait50 (c reduce [1 in {0..99999}]);
  
```

Results may be non-deterministic depending on numeric type

The programmer needs to understand the numerical implications of out-of-order expressions.

Non-determinism

References for this chapter

- Robert Adami, *Ada Reference Manual - Language and Standard Libraries*, ISO/IEC 8652:2018 (E)
- Principles of Concurrent and Distributed Programming, 2nd edition, Addison-Wesley, Pearson Education, ISBN: 13374-8-321-3078-8, Harlow, England, 2006
- Barnes, John, *Programming in Ada 2005*, Addison-Wesley, Pearson Education, ISBN: 13374-8-321-3078-8, Harlow, England, 2006

Non-determinism

Non-determinism by design

Motivation for non-deterministic design

By explicitly leaving the sequence of evaluation or execution undetermined:

- The compiler / runtime environment can directly (i.e. without any abstractions) translate the source code into a concurrent implementation
- The implementation gains potentially significantly in performance
- The programmer does not need to handle any of the details of a concurrent implementation (access, locks, messages, synchronizations, ...)

A programming language which allows for those formulations is required!

current language support: Ada, X10, Chapel, Fortress, Haskell, OCaml, ...

Non-determinism

Non-determinism by interaction

Select function in POSIX

```

int select(int n, fd_set readfds, fd_set writefds, fd_set exceptfds,
          const struct timespec *timeout, sigset_t *sigmask);
  
```

with:

- n being one more than the maximum of any file descriptor in any of the sets.
- after return the sets will have been reduced to the channels which have been triggered.
- the return value is used as success / failure indicator.

The POSIX select function implements parts of general select-with-waiting:

- select returns if one or multiple I/O channels have been triggered or an error occurred.
- finishing into individual code sections is not provided.
- Guards are not provided.

After return it is required that the following code implement a sequential testing of all channels in the sets.

Non-determinism

Selective Synchronization

Basic forms of selective synchronization

```

select
  or
when condition >> accept
or
when condition >> accept
or
when condition >> accept
end select;
  
```

- If all conditions are 'true' => identical to the previous form.
- If some condition evaluate to 'true', conditions are tested like in the previous form
- If all conditions evaluate to 'false' => Program_Error is raised.

Since this is important that the set of conditions covers all possible states.

This form is identical to Dijkstra's guarded commands.

Non-determinism

Definitions

Non-determinism by design:

A property of a computation which may have more than one result.

Non-determinism by interaction:

A property of the operation environment which may lead to different sequences of (concurrent) stimuli.

Non-determinism

Non-determinism by interaction

Selective waiting in Occam2

ALT

```

Guard
Process1
Guard
Process2
...
  
```

- Guards are referring to boolean expressions and/or channel input operations.
- The boolean expressions are local expressions, i.e. if none of them evaluates to true the process is suspended until the next guard is evaluated.
- If all guarded channel input operations evaluate to false, the process is suspended until further activity on one of the named channels.
- Any Occam2 process can be employed in the ALT-statement
- The ALT-statement is non-deterministic (there is also a deterministic version: PRE-ALT).

Non-determinism

Selective Synchronization

Message-based selective synchronization in Ada

```

selective_accept ::= select
  [ or Guard selective_accept_alternative ]
end select;

selective_accept_alternative ::= accept_statement |
  [ or Guard selective_accept_alternative ]
  terminate_alternative;

accept_statement ::= accept_statement [ sequence_of_statements ]
terminate_alternative ::= terminate;
  
```

underlying concept: Dijkstra's guarded commands

- wait for more than a single rendezvous at any one time
- time-out if no rendezvous is forthcoming within a specified time
- withdrew is able to communicate if no rendezvous is available immediately
- terminate if no clients can possibly call its entries

Non-determinism

Selective Synchronization

Basic forms of selective synchronization

```

select
  or
when condition >> accept
or
when condition >> accept
or
when condition >> accept
or
when condition >> delay [unit]
or
when condition >> delay [unit]
end select;
  
```

- If none of the open entries have waiting calls before the deadline specified by the entry's open delay alternative then the current delay alternative is chosen and the entry is selected.
- Otherwise one of the open entries with waiting calls is chosen as above.

This enables a task to withdraw its offer to accept a set of calls if no other task is calling after some time.

Non-determinism

Non-determinism by design

Dijkstra's guarded commands (non-deterministic case statements):

```

if
  x < y >> m := x
  x > y >> m := y
fi
  
```

Selection is non-deterministic (forky)

The programmer needs to design the alternatives as 'parallel' options: all cases need to be covered and overlapping conditions need to lead to the same result

All true case statements in any language are potentially concurrent and non-deterministic.

Non-determinism

Non-determinism by interaction

Selective waiting in Occam2

ALT

```

Number_of_buffers := 5; Size & Guard; Top := Top + 1; Request;
Number_of_buffers := Number_of_buffers + 1;
Take 1 Buffer (Base);
Base := (Base + 1); Request;
Number_of_buffers := Number_of_buffers - 1;
Base := (Base - 1); Request;
  
```

- Synchronization on input-channels only (channels are directed in Occam2); a request need to be made first which triggers the condition. (Base := ? ANY)
- CSO (Control Semantics Operator) (see also the book by Hoare, 1978) also supports non-deterministic selective waiting.

Non-determinism

Selective Synchronization

Basic forms of selective synchronization

```

select
  or
accept
or
accept
or
accept
end select;
  
```

- If none of the entries have waiting calls => the process is suspended
- If exactly one of the entries has waiting calls => this entry is selected
- If multiple entries have waiting calls => one of those is selected non-deterministically. The selection can be prioritized by the order of the entries in the select-statement.

The code following the select-statement is executed and the select-statement completes.

Non-determinism

Selective Synchronization

Basic forms of selective synchronization

```

select
  or
when condition >> accept
or
when condition >> accept
or
when condition >> accept
or
when condition >> terminate;
end select;
  
```

- If none of the open entries have waiting calls and none of them can ever be called => the process is suspended
- If the condition alternative is chosen, i.e. the task is terminated. This situation occurs if:
 - ... all tasks which can possibly call on any of the open entries are terminated.
 - or ... all remaining tasks which can possibly call on select-terminate-statement alternatives and none of their open entries can be called either. In this case all those waiting for termination tasks are terminated as well.

terminate is called with a list of delay mixed with a list of delay

Non-determinism
Selective Synchronization
Message-based selective synchronization in Ada

Forms of selective waiting:

```
select_statement ::= selective_accept
                  | conditional_entry_call
                  | timed_entry_call
                  | asynchronous_select
```

... underlying concept: Dijkstra's guarded commands

```
conditional_entry_call and timed_entry_call implements ...
... the possibility to withdraw an outgoing call.
... this might be restricted if calls have already been partly processed.
```

Non-determinism
Selective Synchronization
Conditional entry-calls

```
conditional_entry_call ::=
  entry_call_statement
  | sequence_of_statements
  | sequence_of_statements
  and select
```

Example:

```
Light_monitor_wait_for_Light;
if Lux := True;
then
  ...
else
  Lux := False;
end;
```

- If the call is not accepted immediately the call alternative is chosen. This is e.g. useful to make the state of a server before committing to a potentially blocking call. Even though it is tempting to use this statement in a busy-waiting semaphore, we need to be careful because the state is not stable. There is only one entry call and one alternative.

Non-determinism
Selective Synchronization
Timed entry-calls

```
timed_entry_call ::=
  entry_call_statement
  | sequence_of_statements
  | delay_alternative
  and select
```

Example:

```
Control_request (Some_Ten);
accept priority_wait
or
  delay 45.8; ..... seconds
  try something else
end select;
```

- If the call is not accepted before the delay line specified by the delay alternative in the delay alternative is chosen. This is e.g. used to withdraw an entry call after some specified times out. There is only one entry call and one delay alternative.

Non-determinism
Selective Synchronization
Message-based selective synchronization in Ada

```
select_statement ::= selective_accept
                  | conditional_entry_call
                  | timed_entry_call
                  | asynchronous_select
```

... underlying concept: Dijkstra's guarded commands

asynchronous_select implements ...

... the possibility to escape a running code block due to an event from outside this task. Inside the scope of this routine or thread (Real Time Systems).

Non-determinism
Sources of Non-determinism

As concurrent entities are not in "lock-step" synchronization, they "overtake" each other and arrive at synchronization points in non-deterministic order, due to (just a few):

- Operating systems / runtimes environments.
- Schedulers are often non-deterministic.
- System load will have an influence on concurrent execution.
- Message passing systems react load dependent.
- Networks & communication systems.
- Traffic will arrive in an unpredictable way (non-deterministic).
- Communication systems congestions are generally unpredictable.
- Computing hardware (e.g. cache) can have parallelities.
- Processors have out-of-order units.
- Basically: Physical systems and computer systems connected to the physical world are intrinsically non-deterministic.

Non-determinism
Correctness of non-deterministic programs

Partial correctness:
 $(P() \wedge \text{terminates}(\text{Program}(I, O))) \Rightarrow Q(I, O)$

Total correctness:
 $P() \Rightarrow (\text{terminates}(\text{Program}(I, O)) \wedge Q(I, O))$

Safety properties:
 $(P() \wedge \text{Processes}(I, S)) \Rightarrow \square Q(I, S)$
 where $\square Q$ means that Q does always hold

Liveness properties:
 $(P() \wedge \text{Processes}(I, S)) \Rightarrow \diamond Q(I, S)$
 where $\diamond Q$ means that Q does eventually hold (and will then stay true) and S is the current state of the concurrent system

Non-determinism
Correctness of non-deterministic programs

or

Correctness predicates need to hold true irrespective of the actual sequence of interaction points.

Correctness predicates need to hold true for all possible sequences of interaction points.

Therefore correctness predicates need to be based on **invariants**, i.e. invariant predicates which are independent of the potential execution sequences, yet support the overall correctness predicates.

Non-determinism
Correctness of non-deterministic programs

For example (in verbal form):

- "Mutual exclusion accessing a specific resource holds true, for all possible numbers, sequences or interleavings of requests to it"

An invariant would for instance be that the number of writing tasks inside a protected object is less or equal to one.

- Those invariants are the only practical way to guarantee (in a logical sense) correctness in concurrent / non-deterministic systems. (as enumerating all possible cases and proving them individually is, in general, not feasible)

Non-determinism
Correctness of non-deterministic programs

```
select
  when <cond1(t) > => accept --
or
  when <cond2(t) > => accept --
or
  when <cond3(t) > => accept --
end select;
```

Generate:

- Every time you formulate a non-deterministic program, you need to formulate an invariant which holds true whichever alternative will actually be chosen. This is very similar to finding loop invariants in sequential programs

Non-determinism
Summary

- Non-determinism by design:
 - Benefits & considerations
- Non-determinism by interaction:
 - Selective synchronization
 - Selective accepts
 - Selective calls
- Correctness of non-deterministic programs:
 - Partial correctness
 - Total correctness
 - Predicates & invariants

Non-determinism
Correctness of non-deterministic programs

Correctness predicates need to hold true irrespective of the actual sequence of interaction points.

Correctness predicates need to hold true for all possible sequences of interaction points.

Therefore correctness predicates need to be based on invariants, i.e. invariant predicates which are independent of the potential execution sequences, yet support the overall correctness predicates.

Non-determinism
Correctness of non-deterministic programs

For example (in verbal form):

- "Mutual exclusion accessing a specific resource holds true, for all possible numbers, sequences or interleavings of requests to it"

An invariant would for instance be that the number of writing tasks inside a protected object is less or equal to one.

- Those invariants are the only practical way to guarantee (in a logical sense) correctness in concurrent / non-deterministic systems. (as enumerating all possible cases and proving them individually is, in general, not feasible)

