T–79.4301 Parallel and Distributed Systems (4 ECTS)

T–79.4301 Rinnakkaiset ja hajautetut järjestelmät (4 op)

Lecture 8

17th of March 2008

Keijo Heljanko

Keijo.Heljanko@tkk.fi

Commonly used Abstractions

We can for example use the following predicates:

- x is even or odd,
- x is zero or non-zero,
- x is `NULL`; and,
- x is positive, zero, or negative.

Predicate Abstraction

We can use abstractions which use boolean variables to record relationships between variables:

- $x_equals_z = (x == z)$, and
- $x_is_less_than_z = (x < z)$.

The last two talk about two variables, and thus can actually replace both x and z with a single boolean variable (the predicate) which is changed in the abstract program whenever either x or z is changed in the concrete program.

This is called predicate abstraction.

Abstraction (cnt.)

- Similar abstraction methods can also be used when manually modelling a system.
- The main goal is to be able to show that when an operation of the concrete program is performed, the abstract program can always simulate it by the abstract version of the operation.
- If this is systematically done, then the traces of the abstract program will always be a superset of the traces of the concrete program.

Abstraction (cnt.)

- Note that the data abstraction method described in the previous slides does not preserve bisimulation.
- However, because simulation is guaranteed, all traces are preserved, and thus the abstraction method can be used for checking trace containment.
- We have only scratched the surface of abstraction methods available.

Building Verification Models

When creating a verification model of a system the following aspects need to be considered:

- Which aspects of the design are important and need verification:
  - What are the correctness requirements of the design?
  - The requirements often tell a lot about which parts of the system to model in order to capture the relevant details. The modelled parts should be exactly those of the required system relevant to the property being verified.
  - Other detail should be discarded.

Building Verification Models (cnt.)

- Model checking is usually best for verifying control flow properties. Data manipulation not necessary for control flow should usually be checked by other means. Discard data in verification models when possible.
- We are looking for the smallest sufficient model to allow verification of the requirement at hand.

Smallest Sufficient Model

- State explosion is usually the main problem in any model checking effort.
- Abstraction is the most efficient way of alleviating the state explosion problem.
- Different properties of the system might in the end need different verification models. (This can lead to version control problems between different models!)
- One should try to avoid unnecessary redundancy in the model.

Avoiding Redundancy

Be careful with unnecessary data. Some examples:

- Temporary variables: get rid of values stored in temporaries not needed for the control flow.
- Generating dummy data that is never used: a sender in a protocol might generate data messages which are never read in the model, just passed along to be discarded in the end. Clearly such data should either be removed or be finally checked in the receiver.
- Sinks, sources, and filters: if model maintainability (readability) allows, processes that only generate, consume, or trivially filter messages can usually be gotten rid of by modifying the model slightly.

Verification Flow

1. Formalize the critical properties to be verified.
2. Construct the smallest sufficient model(s) for the verification task. Use abstraction: watch out for too much detail/concrete model - state explosion might occur.
3. Do the verification by choosing among different alternative options available in the model checker.
4. If a counterexample is found: Either refine the model to remove spurious counterexamples created by abstraction, or modify the concrete system design being modelled to meet its requirements if a real problem with the design shows up.

Verification Flow (cnt.)

- Phase 1. specifying the properties can be a major headache for many systems. If requirements exist, they are usually not formalized, so formalizing the requirements is usually a major undertaking.
- Experience has shown that actually most of the bugs in real designs are found in phase 2. of the verification flow even without starting the model checker. (Modelling is efficient design/code review.)
- Phase 3. can mean modelling changes if the capacity of the available model checker is exceeded.
- Phase 4. is in principle straightforward. However, quite often several iterations are needed.

Safety Properties

- Safety properties are properties of systems that are characterized by the intuitive formulation: “nothing bad happens”.
- Another intuition is the following: If some execution σ of the system breaks a safety property, then also all longer executions of the system which begin with σ break the safety property.
- More formally, given an alphabet Σ, a safety property S is a language $S \subseteq \Sigma^*$ such that: for all words $\sigma \notin S$ it holds that $\sigma\sigma' \notin S$ for all $\sigma' \in \Sigma^*$.

Safety - Examples

Examples of safety properties are:

- It is always the case that two processes are never at the same time in the critical section.
- It is always the case that if the system reboots the reset button has been pressed in the past.
- It is always the case that if a message "`ack0`" arrives from the receiver, then a messages "`data0`" has been sent in the past from the sender.
- It is always the case that a multiply command is followed in three clock cycles by the multiply data.

Safety

- Now given an LTS L_S such that $traces(L_S) = S$, we can check whether for an implementation I it holds that $I \leq_{tr} L_S$ as shown earlier in this course.
- The main problem with this approach is that it requires one to compute the deterministic automaton $\overline{det(S)}$ accepting the complement language of S, which might be of exponential size in S.
- This worst-case blow-up happens very seldom, and thus the approach is actually quite useful in practice.

Run-time Verification

- Safety properties have also another nice feature: They can also be observed during the concrete system runtime.
- Thus synchronizing with a safety property observer $\overline{det(S)}$ can be easily simulated in the final implementation of the system to either shutdown a malfunctioning system or just to log (less crucial) violations of safety properties.
- This is called run-time verification and it can only be done for safety properties.

History-variables Method

- Another possibility for run-time verification is to use the so called history-variables method.
- This approach is based on using a temporal logic which is capable of only specifying safety properties of the system.
- For simplicity we use a state-based version of the logic and assume we are dealing with a Kripke structure (see Lecture 1) with a set of atomic propositions AP.
- The implementation needs to be able to evaluate for each reachable state s whether an atomic proposition $p \in AP$ holds in s or not.

Past Formulas

The logic we use is a (proper) subset of the temporal logic PLTL (linear temporal logic with past) and will be defined using the following syntax:

- $p \in AP$ is a past formula,
- if ψ_1 is a past formula, then $\neg\psi_1$, and $\mathbf{Y}\,\psi_1$ ("yesterday") are past formulas,
- if ψ_1, ψ_2 are past formulas, then $\psi_1 \vee \psi_2$ and $\psi_1\,\mathbf{S}\,\psi_2$ ("since") are past formulas.

Shorthands

We will define the following shorthands:

- $\top = p \vee \neg p$ ("true") for an arbitrary $p \in AP$,
- $\bot = \neg\top$ ("false"),
- $\psi_1 \wedge \psi_2 = \neg((\neg\psi_1) \vee (\neg\psi_2))$,
- $\mathbf{Z}\,\psi_1 = \neg(\mathbf{Y}\,(\neg\psi_1))$ ("weak yesterday"),
- $\mathbf{O}\,\psi_1 = \top\,\mathbf{S}\,\psi_1$ ("once"),
- $\psi_1\,\mathbf{T}\,\psi_2 = \neg((\neg\psi_1)\,\mathbf{S}\,(\neg\psi_2))$ ("trigger"), and
- $\mathbf{H}\,\psi_1 = \bot\,\mathbf{T}\,\psi_1$ ("historically").

Semantics of Past Formulas

The semantics of past formulas is defined at each index i in a word $\pi \in (2^{AP})^*$ such that $\pi = x_0 x_1 x_2 \ldots x_n$ as follows:

$$
\begin{array}{lcl}
\pi^i \models p & \Leftrightarrow & p \in x_i \text{ (i.e., } p \text{ holds in } x_i) \text{ for } p \in AP. \\
\pi^i \models \neg\psi_1 & \Leftrightarrow & \pi^i \not\models \psi_1. \\
\pi^i \models \mathbf{Y}\,\psi_1 & \Leftrightarrow & i > 0 \text{ and } \pi^{i-1} \models \psi_1. \\
\pi^i \models \psi_1 \vee \psi_2 & \Leftrightarrow & \pi^i \models \psi_1 \text{ or } \pi^i \models \psi_2. \\
\pi^i \models \psi_1\,\mathbf{S}\,\psi_2 & \Leftrightarrow & \exists\, 0 \le j \le i \text{ such that } \pi^j \models \psi_2 \text{ and} \\
 & & \pi^n \models \psi_1 \text{ for all } j < n \le i.
\end{array}
$$

Alternative Semantic Definition

We can alternatively define the semantics of $\pi^i \models \mathbf{Y}\,\psi_1$ and $\pi^i \models \psi_1\,\mathbf{S}\,\psi_2$ recursively as follows:

- $i = 0$:
  - $\pi^0 \not\models \mathbf{Y}\,\psi_1$
  - $\pi^0 \models \psi_1\,\mathbf{S}\,\psi_2 \;\Leftrightarrow\; \pi^0 \models \psi_2$
- $i > 0$:
  - $\pi^i \models \mathbf{Y}\,\psi_1 \;\Leftrightarrow\; \pi^{i-1} \models \psi_1$
  - $\pi^i \models \psi_1\,\mathbf{S}\,\psi_2 \;\Leftrightarrow\; \pi^i \models \psi_2 \vee (\psi_1 \wedge \mathbf{Y}\,(\psi_1\,\mathbf{S}\,\psi_2))$

De Morgan Rules

The De Morgan rules are as follows:

$$
\begin{array}{lcl}
\neg(\neg\psi_1) & \Leftrightarrow & \psi_1 \\
\neg(\psi_1 \vee \psi_2) & \Leftrightarrow & (\neg\psi_1) \wedge (\neg\psi_2) \\
\neg(\mathbf{Y}\,\psi_1) & \Leftrightarrow & \mathbf{Z}\,(\neg\psi_1) \\
\neg(\mathbf{O}\,\psi_1) & \Leftrightarrow & \mathbf{H}\,(\neg\psi_1) \\
\neg(\psi_1\,\mathbf{S}\,\psi_2) & \Leftrightarrow & (\neg\psi_1)\,\mathbf{T}\,(\neg\psi_2)
\end{array}
$$

We also have the duals of the De Morgan rules above, e.g., $\neg(\mathbf{Z}\,\psi_1) \Leftrightarrow \mathbf{Y}\,\neg\psi_1$.

Semantics in a Path

A formula $\mathbf{G}\,(\varphi)$ ("always" φ), where φ is a past formula is called a *past safety formula*. The semantics in a path $\pi = x_0 x_1 x_2 \ldots x_n$ is defined as follows:

- $\pi \models \mathbf{G}\,(\varphi)$ iff for all indexes $0 \leq i \leq n$ it holds that $\pi^i \models \varphi$.

or alternatively:

- $\pi \not\models \mathbf{G}\,(\varphi)$ iff there is an index $0 \leq i \leq n$ such that $\pi^i \models \neg\varphi$.

Semantics in a Kripke Structure

- Recall the definition of a Kripke structure $M = (S, s^0, R, L)$ from Lecture 1.
- An execution σ of M is a sequence of states $\sigma = s_0 s_1 \ldots s_n$ such that $s_0 = s^0$ (starts from the initial state), and $(s_{i-1}, s_i) \in R$ for all $1 \leq i \leq n$ (follows the arcs of the Kripke structure).
- An execution path π of M is a sequence of labels $\pi = x_0 x_1 \ldots x_n$, such that $x_i = L(s_i)$ for some execution $\sigma = s_0 s_1 \ldots s_n$ of M.

Semantics in a Kripke Structure (cnt.)

- The formula φ holds in M, denoted $M \models \varphi$ iff $\pi \models \varphi$ holds for every execution path π of M.
- Or alternatively: the formula φ does not hold in M, denoted $M \not\models \varphi$ iff there is an execution path $\pi = x_0 x_1 \ldots x_n$ such that $\pi \models \neg\varphi$.
  - Such a path φ is called a *counterexample* to property φ, and the corresponding execution σ is called the counterexample execution.

Examples

- $\mathbf{G}\,(\neg(cr_0 \wedge cr_1))$: processes 0 and 1 are never at the same time in the critical section.
- $\mathbf{G}\,(starts \Rightarrow \mathbf{O}\,(ignition))$: if the car starts the ignition key has been turned in the past.
- $\mathbf{G}\,(alarm \Rightarrow \mathbf{O}\,(crash))$: an alarm is given only if the system has crashed in the past.
- $\mathbf{G}\,(alarm \Rightarrow (\neg reset\,\mathbf{S}\,crash))$: an alarm is given only if the system has crashed in the past and no reset has been applied since.
- $\mathbf{G}\,(alarm \Rightarrow \mathbf{Y}\,(crash))$: if an alarm is given, the system crashed at the previous time step.

Implementing the semantics

- To find a safety violation, we need to observe the system state after each step it makes, and report an error at the first index i such that $\pi^i \models \neg\varphi$.
- We do this by using two boolean variables for each subformula ψ. One bit to store the current value of ψ and another bit to remember the value of ψ at the previous time step, denoted by ψ'.
- We can do the calculation of the new values for all the bits as shown in the following slides.
- If after running the system for i steps the top-level formula $\neg\varphi$ evaluates to true we need report that the past safety formula $\mathbf{G}\,(\varphi)$ is violated.

mplementing the semantics (cnt.)

- We will now evaluate the subformula value ψ in bottom-up order. Namely, the evaluation order must be such that both subformulas ψ_1 and ψ_2 of ψ have been evaluated at the current state s_i before ψ is evaluated.
- Each subformula ψ must also be evaluated exactly once at each s_i.
- The implementation is based on the alternative recursive semantic definition.
- To know the contents of the next two slides will not be part of the exam requirements.

The Translation at $i = 0$

Formula ψ	Update rules at $i = 0$
$\psi \in AP$	$\psi = evaluate(s_i, \psi)$
$\neg\psi_1$	$\psi = \neg\psi_1$
$\psi_1 \vee \psi_2$	$\psi = \psi_1 \vee \psi_2$
$\mathbf{Y}\,\psi_1$	$\psi = \bot$ (false)
$\psi_1\,\mathbf{S}\,\psi_2$	$\psi = \psi_2$

Where $evaluate(s_i, \psi)$ evaluates the atomic proposition ψ in the current state s_i.

The Translation at $i > 0$

Formula ψ	Update rules at $i > 0$
$\psi \in AP$	$\psi' = \psi; \psi = evaluate(s_i, \psi)$
$\neg\psi_1$	$\psi' = \psi; \psi = \neg\psi_1$
$\psi_1 \vee \psi_2$	$\psi' = \psi; \psi = \psi_1 \vee \psi_2$
$\mathbf{Y}\,\psi_1$	$\psi' = \psi; \psi = \psi'_1$
$\psi_1\,\mathbf{S}\,\psi_2$	$\psi' = \psi; \psi = \psi_2 \vee (\psi_1 \wedge \psi')$

Where ψ'_1 (ψ') is the value of ψ_1 (ψ) at the previous time step, and $evaluate(s_i, \psi)$ evaluates the atomic proposition ψ in the current state s_i.

History-variables Implementation

- The implementation of the history variables method can be made extremely fast.
- The memory overhead is tiny, just two bits per subformula, out of which the ψ' variables are just temporaries needed to evaluate the new ψ variables.
- It can be used as a fast, low-overhead runtime verification observer for safety properties. The same observer can also be used in combination with a model checker to check safety properties.
- Unfortunately the procedure is not implemented in most model checkers, so it has to be usually implemented by hand.