

COMP4161 T3/2019

Advanced Topics in Software Verification

Assignment 2

This assignment starts on **Friday 18th of October** and is due on **Friday 1st of November 6PM**. We will accept Isabelle theory (.thy) files.

The assignment is take-home. This does NOT mean you can work in groups. Each submission is personal. For more information, see the plagiarism policy: <https://student.unsw.edu.au/plagiarism>

Submit using `give` on a CSE machine:

```
give cs4161 a2 files ...
```

For example:

```
give cs4161 a2 a2.thy
```

Hint: there are hints at the end of this document.

1 A theory of network connectivity (25 marks)

1.1 Background

In this section, we will consider a simple model of connectivity between devices in a hierarchical network.

We will use natural numbers to stand for device identifiers, and will describe network topologies as a set of *connectivity conditions*. The conditions are of two kinds:

1. *Joinable* $n m$ means that the devices m and n are connected, in the sense that they can send messages to each other via some (unspecified) route.
2. *Above* $n m$, pronounced “ m is above n ”, means that device m is an uplink for n . In other words, m can forward messages to other devices in the network on behalf of n .

From an initial set A of conditions describing the network topology, we can construct a set *connection* A of all connections between devices that can be established in the network. *connection* A is defined inductively by the following rules:




Figure 1: An example topology.

1. I can forward messages on behalf of myself:

$$\frac{}{\textit{Above } a \in \textit{connection } A}$$

2. The connections in the initial topology are available:

$$\frac{\varphi \in A}{\varphi \in \textit{connection } A}$$

3. If you and I are connected, then I and you are connected:

$$\frac{\textit{Joinable } a \ b \in \textit{connection } A}{\textit{Joinable } b \ a \in \textit{connection } A}$$

4. My uplink's uplink is my (virtual) uplink:

$$\frac{\textit{Above } a \ b \in \textit{connection } A \quad \textit{Above } b \ c \in \textit{connection } A}{\textit{Above } a \ c \in \textit{connection } A}$$

5. We are connected if we share a common uplink:

$$\frac{\textit{Above } a \ b \in \textit{connection } A \quad \textit{Above } c \ b \in \textit{connection } A}{\textit{Joinable } a \ c \in \textit{connection } A}$$

6. If my uplink is connected to someone, then so am I:

$$\frac{\textit{Above } a \ b \in \textit{connection } A \quad \textit{Joinable } b \ c \in \textit{connection } A}{\textit{Joinable } a \ c \in \textit{connection } A}$$

As an example, consider the topology illustrated in Figure 1, where arrows represent the *Above* relationship. Here Computers 1 and 2 are joinable because of their common uplink (Router 1), and so are Computers 2 and 3. Computers 1 and 3 are *not* joinable because there is no common uplink to forward messages between the routers.

1.2 Questions

- (a) (3 marks) Implement the topology described in Figure 1 as a set of *Above* judgements, and prove that computers 1 and 2 are joinable.

Prove that the connectivity relation satisfies these properties:

- (b) (2 marks) $\varphi \in \text{connection } A \implies \varphi \in \text{connection } (A \cup B)$
- (c) (3 marks) $\text{connection } (\text{connection } A) = \text{connection } A$
- (d) (2 marks) $\varphi \in \text{connection } (A \cup B) \implies \exists C D. C \subseteq \text{connection } A \wedge D \subseteq \text{connection } B \wedge \varphi \in \text{connection } (C \cup D)$
- (e) (2 marks) $\varphi \in \text{connection } \emptyset \implies \text{is-refl } \varphi$
- (f) (2 marks) $\text{is-refl } \varphi \implies \varphi \in \text{connection } A$
- (g) (2 marks) $\varphi \in \text{connection } A \implies \varphi \in \text{connection } (A - \{\varphi \mid \text{is-refl } \varphi\})$
- (h) (3 marks) $\llbracket \text{Above } a \ b \in \text{connection } A; \wedge a \ b. \text{Above } a \ b \notin A \rrbracket \implies a = b$
- (i) (3 marks) $\llbracket \varphi \in \text{connection } (C \cup D); C \subseteq \text{connection } A; D \subseteq \text{connection } B \rrbracket \implies \varphi \in \text{connection } (A \cup B)$
- (j) (3 marks) $\text{connection } A = \text{connection } B \implies \text{connection } (A \cup C) = \text{connection } (B \cup C)$

2 A theory of communicating devices (15 marks)

2.1 Background

We will now extend our (static) model of a network to include the dynamic behaviour of processes communicating through the network. For simplicity we do not model data: communication is just a synchronous handshake with no information being exchanged. The state of the network is described by this datatype:

```
datatype process =  
  Cond condition  
  | Par process process  
  | Input nat process  
  | Output nat process  
  | Nil
```

A network may: contain connectivity assumptions (*Cond* φ); consist of multiple processes running in parallel, possibly on different devices (*Par* $P Q$); contain a process that does nothing (*Nil*); have device m ready to receive a message and then continue as P (*In* $m P$); or have a device ready to send a message (*Output* $m P$).

The behaviour of the network is captured with a small-step transition relation *semantics* $A P \alpha P'$, meaning that in a network with topology A , the process P can perform the action α and transition to the state P' . The actions can be input, output or internal (*LTau*):

datatype action = LInput nat | Output nat | LTau

The *frame* of a process collects all its top-level conditions, which taken together describe the current network topology. Conditions underneath *In* and *Out* prefixes represent ways the topology may change in the future.

An output or input prefix may send or receive a message, consuming the prefix:

$$\frac{}{\textit{semantics } A \textit{ (Input } n P \textit{) (LInput } n \textit{) } P}$$

$$\frac{}{\textit{semantics } A \textit{ (Output } n P \textit{) (LOutput } n \textit{) } P}$$

If a process P can perform an action, it may also perform that action with Q sitting inertly in parallel. The derivation of the transition from P may use the connectivity information present in *frame* Q :

$$\frac{\textit{semantics } (A \cup \textit{frame } Q) P \alpha P'}{\textit{semantics } A \textit{ (Par } P Q \textit{) } \alpha \textit{ (Par } P' Q \textit{)}}$$

If two parallel processes P and Q can perform an input and output, respectively, and if the devices are joinable, then P and Q may synchronise to perform a communication action (*LTau*):

$$\frac{\begin{array}{l} \textit{semantics } (A \cup \textit{frame } Q) P \textit{ (LOutput } n \textit{) } P' \\ \textit{semantics } (A \cup \textit{frame } P) Q \textit{ (LInput } m \textit{) } Q' \\ \textit{Joinable } n m \in \textit{connection } (A \cup \textit{frame } P \cup \textit{frame } Q) \end{array}}{\textit{semantics } A \textit{ (Par } P Q \textit{) } \textit{LTau } \textit{ (Par } P' Q' \textit{)}}$$

Symmetric versions of the rules for parallel and communication are elided.

2.2 Questions

Prove that the semantics satisfies the following:

- (a) (5 marks) $\text{semantics } A \ P \ \alpha \ Q \implies \text{semantics } (A \cup B) \ P \ \alpha \ Q$
- (b) (5 marks) $\text{semantics } A \ P \ \alpha \ Q \implies \exists \beta \ Q'. \text{ semantics } \emptyset \ P \ \beta \ Q'$
- (c) (5 marks) $\llbracket \text{semantics } A \ P \ \alpha \ Q; \text{ connection } A = \text{connection } B \rrbracket \implies \text{semantics } B \ P \ \alpha \ Q$

3 A theory of behavioural equivalence (60 marks)

3.1 Background

It is often useful to be able to say when two distinct processes have the same behaviour. For the purposes of this exercise, we will use *trace equivalence* for this purpose. Trace equivalence may be intuitively characterised as follows. Imagine you are presented with a log (aka trace) of all the events that transpired when a process ran to completion. You know that the log describes an execution of either P or Q , but you haven't been told which. By inspecting the log, can you figure out which process the log originates from? If the answer to this question is always no for every log, we say that P and Q are trace equivalent.

Let's formalise this intuitive understanding. A process is *stuck* if it has run to completion; that is, if it has no outgoing transitions:

$$\text{stuck } P \equiv \forall \alpha \ P'. \text{ semantics } \emptyset \ P \ \alpha \ P' \longrightarrow \text{False}$$

The relation *list-trans* lifts *semantics* from single actions to sequences of actions. A *trace* of a process is a sequence of actions it can take that lead to a stuck state. The set of all such traces is given by

$$\text{traces-of } P = \{tr \mid \exists Q. \text{ list-trans } \emptyset \ P \ tr \ Q \wedge \text{ stuck } Q\}$$

Finally, two processes are trace equivalent if they have the same sets of traces—or in other words, if for every trace we cannot tell which process it originated from:

$$\text{trace-eq } P \ Q = (\text{traces-of } P = \text{traces-of } Q)$$

3.2 Questions

- (a) (3 marks) Show that trace equivalence is indeed an equivalence relation:

$$\text{trace-eq } P \ P$$

$$\text{trace-eq } P \ Q \implies \text{trace-eq } Q \ P$$

$$\llbracket \text{trace-eq } P \ Q; \text{ trace-eq } Q \ R \rrbracket \implies \text{trace-eq } P \ R$$

- (b) (4 marks) $\llbracket \text{stuck } P; \text{stuck } Q \rrbracket \implies \text{trace-eq } P \ Q$
- (c) (8 marks) $\text{stuck } R \implies \text{traces-of } P \subseteq \text{traces-of } (\text{Par } P \ R)$
- (d) (4 marks) The previous question used set inclusion instead of trace equivalence because inclusion does not hold in the other direction: it is possible to introduce more traces by adding a stuck parallel component. Give an example of concrete P and R such that the following proposition is *false*:
- $$\text{stuck } R \implies \text{traces-of } (\text{Par } P \ R) \subseteq \text{traces-of } P$$
- (e) (7 marks) Use your counterexample from the previous question to formally prove
- $$(\bigwedge P \ R. \text{stuck } R \implies \text{traces-of } (\text{Par } P \ R) \subseteq \text{traces-of } P) \implies \text{False}$$

Fill in the gaps to complete a proof that parallel composition is associative up to trace equivalence:

- (f) (6 marks) $\text{semantics } A \ (\text{Par } P \ (\text{Par } Q \ S)) \ \alpha \ R \implies \exists P' \ Q' \ S'. R = \text{Par } P' \ (\text{Par } Q' \ S') \wedge \text{semantics } A \ (\text{Par } (\text{Par } P \ Q) \ S) \ \alpha \ (\text{Par } (\text{Par } P' \ Q') \ S')$
- (g) (2 marks) $\text{list-trans } A \ (\text{Par } (\text{Par } P \ Q) \ S) \ \text{tr } (\text{Par } (\text{Par } P' \ Q') \ S') = \text{list-trans } A \ (\text{Par } P \ (\text{Par } Q \ S)) \ \text{tr } (\text{Par } P' \ (\text{Par } Q' \ S'))$
- (h) (4 marks) $\text{stuck } (\text{Par } P \ Q) = (\text{stuck } P \ \wedge \ \text{stuck } Q)$
- (i) (4 marks) $\text{trace-eq } (\text{Par } (\text{Par } P \ Q) \ S) \ (\text{Par } P \ (\text{Par } Q \ S))$

Finally, prove that Nil is the unit of parallel composition, and that parallel composition commutes:

- (j) (9 marks) $\text{trace-eq } P \ (\text{Par } P \ \text{Nil})$
- (k) (9 marks) $\text{trace-eq } (\text{Par } P \ Q) \ (\text{Par } Q \ P)$

4 Hints

- You are allowed to use sledgehammer as much as you like.
- You are allowed—encouraged, in fact—to use solutions to questions as lemmas in the answers to other questions.
- Most proofs will require induction of one kind or the other. The relevant induction rules are `connection.induct`, `semantics.induct` and `list.induct`.

- The assumptions of theorems stated in the `assumes-show` format are accessible via the fact named `assms`. For example, you can do `simp add: assms` or `rule assms(1)`. The assumptions can be added directly to the goal state by beginning your proof with the command `using assms`; that is usually what you want to do before starting an induction.
- The equivalent of `spec` for the meta-logic universal quantifier is called `meta_spec`.
- `semantics.simps` is sometimes useful for doing case analysis on the derivation of a transition, but is loop-prone. It is often more useful to do this case analysis using specialised elimination rules like `par` and `nil`. Another option is to specialise `semantics.simps` to the pattern you want to decompose before adding it to the simpset. For example, if your assumptions mention a transition from a process `Par P Q`, adding the following will decompose it into assumptions about transitions from `P` and `Q`.

```
simp add: semantics.simps[where ?a2.0="Par P Q" for P
                             Q,simplified]
```

- Some inductive proofs will require you to strengthen the induction hypothesis in order to close the proof. Think about which variables should be `arbitrary`.
- For the later exercises, your life will be much easier if you decompose your proofs into auxiliary lemmas. For example, proofs of trace equivalence naturally decompose into separate proofs about `stuck` and `list-trans`. For proofs about `list-trans`, a good strategy is to first prove a similar lemma for single-step transitions (`semantics`), then lift them to `list-trans` by induction on the action sequence.