

To prove equivalence (3.2), suppose first that a path satisfies $\phi U \psi$. Then, from clause 11, we have $i \geq 1$ such that $\pi^i \models \psi$ and for all $j = 1, \dots, i-1$ we have $\pi^j \models \phi$. From clause 12, this proves $\phi W \psi$, and from clause 10 it proves $F \psi$. Thus for all paths π , if $\pi \models \phi U \psi$ then $\pi \models \phi W \psi \wedge F \psi$. As an exercise, the reader can prove it the other way around.

Writing W in terms of U is also possible: W is like U but also allows the possibility of the eventuality never occurring:

$$\phi W \psi \equiv \phi U \psi \vee G \phi. \quad (3.3)$$

Inspection of clauses 12 and 13 reveals that R and W are rather similar. The differences are that they swap the roles of their arguments ϕ and ψ ; and the clause for W has an $i-1$ where R has i . Therefore, it is not surprising that they are expressible in terms of each other, as follows:

$$\phi W \psi \equiv \psi R (\phi \vee \psi) \quad (3.4)$$

$$\phi R \psi \equiv \psi W (\phi \wedge \psi). \quad (3.5)$$

3.2.5 Adequate sets of connectives for LTL

Recall that $\phi \equiv \psi$ holds iff any path in any transition system which satisfies ϕ also satisfies ψ , and vice versa. As in propositional logic, there is some redundancy among the connectives. For example, in Chapter 1 we saw that the set $\{\perp, \wedge, \neg\}$ forms an adequate set of connectives, since the other connectives \vee, \rightarrow, \top , etc., can be written in terms of those three.

Small adequate sets of connectives also exist in LTL. Here is a summary of the situation.

- X is completely orthogonal to the other connectives. That is to say, its presence doesn't help in defining any of the other ones in terms of each other. Moreover, X cannot be derived from any combination of the others.
- Each of the sets $\{U, X\}$, $\{R, X\}$, $\{W, X\}$ is adequate. To see this, we note that
 - R and W may be defined from U , by the duality $\phi R \psi \equiv \neg(\neg\phi U \neg\psi)$ and equivalence (3.4) followed by the duality, respectively.
 - U and W may be defined from R , by the duality $\phi U \psi \equiv \neg(\neg\phi R \neg\psi)$ and equivalence (3.4), respectively.
 - R and U may be defined from W , by equivalence (3.5) and the duality $\phi U \psi \equiv \neg(\neg\phi R \neg\psi)$ followed by equivalence (3.5).

Sometimes it is useful to look at adequate sets of connectives which do not rely on the availability of negation. That's because it is often convenient to assume formulas are written in negation-normal form, where all the negation symbols are applied to propositional atoms (i.e., they are near the leaves

of the parse tree). In this case, these sets are adequate for the fragment without X , and no strict subset is: $\{U, R\}$, $\{U, W\}$, $\{U, G\}$, $\{R, F\}$, $\{W, F\}$. But $\{R, G\}$ and $\{W, G\}$ are not adequate. Note that one cannot define G with $\{U, F\}$, and one cannot define F with $\{R, G\}$ or $\{W, G\}$.

We finally state and prove a useful equivalence about U .

Theorem 3.10 The equivalence $\phi U \psi \equiv \neg(\neg\psi U (\neg\phi \wedge \neg\psi)) \wedge F \psi$ holds for all LTL formulas ϕ and ψ .

PROOF: Take any path $s_0 \rightarrow s_1 \rightarrow s_2 \rightarrow \dots$ in any model.

First, suppose $s_0 \models \phi U \psi$ holds. Let n be the smallest number such that $s_n \models \psi$; such a number has to exist since $s_0 \models \phi U \psi$; then, for each $k < n$, $s_k \models \phi$. We immediately have $s_0 \models F \psi$, so it remains to show $s_0 \models \neg(\neg\psi U (\neg\phi \wedge \neg\psi))$, which, if we expand, means:

(*) for each $i > 0$, if $s_i \models \neg\phi \wedge \neg\psi$, then there is some $j < i$ with $s_j \models \psi$.

Take any $i > 0$ with $s_i \models \neg\phi \wedge \neg\psi$; $i > n$, so we can take $j \stackrel{\text{def}}{=} n$ and have $s_j \models \psi$.

Conversely, suppose $s_0 \models \neg(\neg\psi U (\neg\phi \wedge \neg\psi)) \wedge F \psi$ holds; we prove $s_0 \models \phi U \psi$. Since $s_0 \models F \psi$, we have a minimal n as before. We show that, for any $i < n$, $s_i \models \phi$. Suppose $s_i \models \neg\phi$; since n is minimal, we know $s_i \models \neg\psi$, so by (*) there is some $j < i < n$ with $s_j \models \psi$, contradicting the minimality of n . \square

3.3 Model checking: systems, tools, properties

3.3.1 Example: mutual exclusion

Let us now look at a larger example of verification using LTL, having to do with *mutual exclusion*. When concurrent processes share a resource (such as a file on a disk or a database entry), it may be necessary to ensure that they do not have access to it at the same time. Several processes simultaneously editing the same file would not be desirable.

We therefore identify certain *critical sections* of each process' code and arrange that only one process can be in its critical section at a time. The critical section should include all the access to the shared resource (though it should be as small as possible so that no unnecessary exclusion takes place). The problem we are faced with is to find a *protocol* for determining which process is allowed to enter its critical section at which time. Once we have found one which we think works, we verify our solution by checking that it has some expected properties, such as the following ones:

Safety: Only one process is in its critical section at any time.

```

MODULE one-bit-chan(input)
VAR
  output : boolean;
  forget : boolean;
ASSIGN
  next(output) := case
    forget : output;
    1:      input;
  esac;
FAIRNESS running
FAIRNESS input & !forget
FAIRNESS !input & !forget

MODULE two-bit-chan(input1,input2)
VAR
  forget : boolean;
  output1 : boolean;
  output2 : boolean;
ASSIGN
  next(output1) := case
    forget : output1;
    1:      input1;
  esac;
  next(output2) := case
    forget : output2;
    1:      input2;
  esac;
FAIRNESS running
FAIRNESS input1 & !forget
FAIRNESS !input1 & !forget
FAIRNESS input2 & !forget
FAIRNESS !input2 & !forget

```

Figure 3.16. The two modules for the two ABP channels in SMV.

constraints of the form ‘infinitely often p implies infinitely often q ’, which would be more satisfactory here, but is not allowed by SMV.

Finally, we tie it all together with the module `main` (Figure 3.17). Its role is to connect together the components of the system, and giving them initial values of their parameters. Since the first control bit is 0, we also initialise the receiver to expect a 0. The receiver should start off by sending 1 as its

```

MODULE main
VAR
  s : process sender(ack_chan.output);
  r : process receiver(msg_chan.output1,msg_chan.output2);
  msg_chan : process two-bit-chan(s.message1,s.message2);
  ack_chan : process one-bit-chan(r.ack);
ASSIGN
  init(s.message2) := 0;
  init(r.expected) := 0;
  init(r.ack)      := 1;
  init(msg_chan.output2) := 1;
  init(ack_chan.output) := 1;

LTLSPEC G (s.st=sent & s.message1=1 -> msg_chan.output1=1)

```

Figure 3.17. The main ABP module.

acknowledgement, so that `sender` does not think that its very first message is being acknowledged before anything has happened. For the same reason, the output of the channels is initialised to 1.

The specifications for ABP. Our SMV program satisfies the following specifications:

Safety: If the message bit 1 has been sent and the correct acknowledgement has been returned, then a 1 was indeed received by the receiver: $G (S.st=sent \ \& \ S.message1=1 \rightarrow msg_chan.output1=1)$.

Liveness: Messages get through eventually. Thus, for any state there is inevitably a future state in which the current message has got through. In the module `sender`, we specified $G \ F \ st=sent$. (This specification could equivalently have been written in the main module, as $G \ F \ S.st=sent$.) Similarly, acknowledgements get through eventually. In the module `receiver`, we write $G \ F \ st=received$.

3.4 Branching-time logic

In our analysis of LTL (*linear-time temporal logic*) in the preceding sections, we noted that LTL formulas are evaluated on *paths*. We defined that a *state* of a system satisfies an LTL formula if *all paths* from the given state satisfy it. Thus, LTL implicitly quantifies universally over paths. Therefore, properties which assert the existence of a path cannot be expressed in LTL. This problem can partly be alleviated by considering the negation of the property in question, and interpreting the result accordingly. To check whether there

exists a path from s satisfying the LTL formula ϕ , we check whether all paths satisfy $\neg\phi$; a positive answer to this is a negative answer to our original question, and vice versa. We used this approach when analysing the ferryman puzzle in the previous section. However, as already noted, properties which *mix* universal and existential path quantifiers cannot in general be model checked using this approach, because the complement formula still has a mix.

Branching-time logics solve this problem by allowing us to quantify explicitly over paths. We will examine a logic known as *Computation Tree Logic*, or CTL. In CTL, as well as the temporal operators U, F, G and X of LTL we also have quantifiers A and E which express 'all paths' and 'exists a path', respectively. For example, we can write:

- There is a reachable state satisfying q : this is written $EF q$.
- From all reachable states satisfying p , it is possible to maintain p continuously until reaching a state satisfying q : this is written $AG(p \rightarrow E[p U q])$.
- Whenever a state satisfying p is reached, the system can exhibit q continuously forevermore: $AG(p \rightarrow EG q)$.
- There is a reachable state from which all reachable states satisfy p : $EF AG p$.

3.4.1 Syntax of CTL

Computation Tree Logic, or CTL for short, is a *branching-time* logic, meaning that its model of time is a tree-like structure in which the future is not determined; there are different paths in the future, any one of which might be the 'actual' path that is realised.

As before, we work with a fixed set of atomic formulas/descriptions (such as p, q, r, \dots , or p_1, p_2, \dots).

Definition 3.12 We define CTL formulas inductively via a Backus Naur form as done for LTL:

$$\begin{aligned} \phi ::= & \perp \mid \top \mid p \mid (\neg\phi) \mid (\phi \wedge \phi) \mid (\phi \vee \phi) \mid (\phi \rightarrow \phi) \mid AX\phi \mid EX\phi \mid \\ & AF\phi \mid EF\phi \mid AG\phi \mid EG\phi \mid A[\phi U \phi] \mid E[\phi U \phi] \end{aligned}$$

where p ranges over a set of atomic formulas.

Notice that each of the CTL temporal connectives is a pair of symbols. The first of the pair is one of A and E. A means 'along All paths' (*inevitably*) and E means 'along at least (there Exists) one path' (*possibly*). The second one of the pair is X, F, G, or U, meaning 'neXt state,' 'some Future state,' 'all future states (Globally)' and Until, respectively. The pair of symbols in $E[\phi_1 U \phi_2]$, for example, is EU. In CTL, pairs of symbols like EU are

indivisible. Notice that AU and EU are binary. The symbols X, F, G and U cannot occur without being preceded by an A or an E; similarly, every A or E must have one of X, F, G and U to accompany it.

Usually weak-until (W) and release (R) are not included in CTL, but they are derivable (see Section 3.4.5).

Convention 3.13 We assume similar binding priorities for the CTL connectives to what we did for propositional and predicate logic. The unary connectives (consisting of \neg and the temporal connectives AG, EG, AF, EF, AX and EX) bind most tightly. Next in the order come \wedge and \vee ; and after that come \rightarrow , AU and EU.

Naturally, we can use brackets in order to override these priorities. Let us see some examples of well-formed CTL formulas and some examples which are not well-formed, in order to understand the syntax. Suppose that p, q and r are atomic formulas. The following are well-formed CTL formulas:

- $AG(q \rightarrow EGr)$, note that this is not the same as $AG q \rightarrow EGr$, for according to Convention 3.13, the latter formula means $(AG q) \rightarrow (EG r)$
- $EFE[r U q]$
- $A[p U EF r]$
- $EF EG p \rightarrow AF r$, again, note that this binds as $(EF EG p) \rightarrow AF r$, not $EF(EG p \rightarrow AF r)$ or $EF EG(p \rightarrow AF r)$
- $A[p_1 U A[p_2 U p_3]]$
- $E[A[p_1 U p_2] U p_3]$
- $AG(p \rightarrow A[p U (\neg p \wedge A[\neg p U q])])$.

It is worth spending some time seeing how the syntax rules allow us to construct each of these. The following are not well-formed formulas:

- $EFGr$
- $A \neg G \neg p$
- $F[r U q]$
- $EF(r U q)$
- $AEF r$
- $A[(r U q) \wedge (p U r)]$.

It is especially worth understanding why the syntax rules don't allow us to construct these. For example, take $EF(r U q)$. The problem with this string is that U can occur only when paired with an A or an E. The E we have is paired with the F. To make this into a well-formed CTL formula, we would have to write $EFE[r U q]$ or $EFA[r U q]$.