

Specifying Systems (using TLA+)

Based on Leslie Lamport's book "Specifying Systems"

Definition: *State*

- Definition: A *state* is an assignment of values to (*all*) variables
- TLA+ notation: $[var_1 = value_1, var_2 = value_2, \dots]$
 - Meaning: a state in which var_1 has value $value_1$, \dots
 - Order is immaterial
- Example: $[hr = 3]$
 - Meaning: a state in which $hr = 3$
 - The values of other variables are not specified
 - There can be many infinitely many states in which $hr = 3$
 - e.g. $[hr = 3. temp = 62]$, $[hr = 3. temp = 68]$, ...
 - *Models* perhaps the hour hand being 3 on some hour clock HC

Definition: *Behavior*

- Definition 1: A *behavior* is a function of time to state

Computer systems can be thought of as executing in steps, so

- Definition 2: A *behavior* is a sequence of states
- Notation: $state_1 \rightarrow state_2 \rightarrow state_3 \rightarrow \dots$
- Example: $[hr = 11] \rightarrow [hr = 12] \rightarrow [hr = 1]$

Definition: *Step*

- Definition: A *step* consists of two consecutive states in a behavior
- aka *transition*
- Notation: $state_1 \rightarrow state_2$
- Example: $[hr = 3] \rightarrow [hr = 4]$

Definition: *Specification*

- A *specification* is a set of all possible behaviors
- Consists of two parts
 1. Set of all possible *initial states*
 2. A “*next-state*” relation that describes the ways a state may change in a step
 - i.e., the set of all possible pairs of states

Set of Initial States

- Example: $\text{HCini} \triangleq hr \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\}$
 - Or, informally, $\text{HCini} \triangleq hr \in \{1, \dots, 12\}$
 - HCini is simply a name given to the predicate
- A set of states can often be succinctly described by a predicate
 - Example: $\text{HCini} \triangleq hr \in \mathbb{N} \wedge 1 \leq hr \wedge hr \leq 12$
- Note again that these describe not 12 but an infinite set of states

Definition: *Next-State Relation*

- A *next-state relation* is a relation between pairs of successive states
 - $\{(state_1^{pre}, state_1^{post}), (state_2^{pre}, state_2^{post}), \dots\}$
- Example:
 - $HC_{nxt} \triangleq \{([hr = 11], [hr = 12]), ([hr = 12], [hr = 1]), \dots\}$

Definition: *Action*

- A next-state relation can often be more succinctly described by a predicate
- Definition 1: an *action* is a predicate over a pair of states
- Example: $\text{HCnxt} \triangleq hr' = hr \% 12 + 1$ (% is the “modulo” operator)
 - or, $\text{HCnxt}_2 \triangleq hr' = \text{IF } hr = 12 \text{ THEN } 1 \text{ ELSE } hr + 1$
 - But note that $\text{HCnxt}_2 \not\equiv \text{HCnxt}$
- hr' is the value of hr in the new state; hr is the value in the old state
- Definition 2: an *action* is a predicate containing both primed and unprimed variables
- An ordinary predicate and does not have to be of the form “ $x' = f(x)$ ”
 - Example: $\text{HCnxt} \triangleq hr' - hr = 1 \text{ mod } 12$

Steps versus Actions versus Execution

- A *step* is a pair of states
- An *action* \mathcal{A} is a predicate over steps
- We call a step that satisfies \mathcal{A} an \mathcal{A} step
 - Example: a step that satisfies HCnxt is an *HCnxt step*
- We sometimes informally say that HCnxt is *executed*

Example specification: hour clock (in complete isolation)

Module HourClock

Variable hr

- $HCini \triangleq hr \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\}$
- $HCnxt \triangleq hr' = hr \bmod 12 + 1$
- $HC \triangleq HCini \wedge \square HCnxt$

Temporal logic formula $\square P$ means that predicate P *always* holds
(thus $HCnxt$ is *invariant* in HC)

Note:

1. All three statements are definitions, but the last one happens to constitute the full specification of the hour clock)
2. There is no conventional naming in TLA+, so pick names that are descriptive

Definition: *Stuttering steps*

- Clocks are usually part of a larger system
- They have more state variables than just the hour hand of the clock
- State changes must allow for hour hand not to change
 - Example: $[hr = 3. temp = 62] \rightarrow [hr = 3. temp = 63]$
- This is called a *stuttering step* of the clock
 - i.e., $hr' = hr$

Final specification: hardware clock

Module HourClock

- Variable hr
- $HCini \triangleq hr \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\}$
- $HCnxt \triangleq hr' = hr \bmod 12 + 1$
- $HC \triangleq HCini \wedge \square(HCnxt \vee (hr' = hr))$

The latter can be abbreviated using the following TLA+ notation

$$HC \triangleq HCini \wedge \square[HCnxt]_{hr}$$

($[HCnxt]_{hr}$ is pronounced "square HCnxt sub hr")

Definition: *theorem*

- Definition: in TLA+, a *theorem* of a specification is a temporal formula that holds over every behavior of the specification
- Example: $HC \Rightarrow \Box hr \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\}$
 - That is, $HC \Rightarrow \Box HCini$
- Proof: by induction on #steps

A note on variables and types

- Variables in TLA+ are untyped
- However, if one can prove $\text{SPEC} \Rightarrow \Box v \in S$ for some variable v and constant set S , then one can call S the type of v in SPEC
- Example: the type of hr in HC is $\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\}$
- It is useful to specify the types in a specification
- Example: $\text{HCtypeInvariant} \triangleq hr \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\}$
- Note, in this case $\text{HCtypeInvariant} \equiv \text{HCini}$

A note on states and behaviors

- Recall
 - A state is an assignment of values to variables
 - A behavior is a sequence of states
- Thus
 - $[hr = 13]$ is still a state, and so is $[hr = \textit{blue}]$
 - $[hr = 4] \rightarrow [hr = 3]$ is still a behavior
- However, they are not in specification HC

Recall

- A *state* is an assignment of values to all variables
- A *step* is a pair of states
- A *stuttering step* wrt some variable leaves the variable unchanged
- An *action* is a predicate over a pair of states
 - If x is a variable in the old state, then x' is the same variable in the new state
- A *behavior* is an infinite sequence of states (with an initial state)
- A *specification* characterizes the initial state and actions

Spec that generates all prime numbers

MODULE *prime*

EXTENDS *Naturals*

VARIABLE *p*

$isPrime(q) \triangleq q > 1 \wedge \forall r \in 2 .. (q - 1) : q \% r \neq 0$

$TypeInvariant \triangleq isPrime(p)$

$Init \triangleq p = 2$

$Next \triangleq p' > p \wedge isPrime(p') \wedge \forall q \in (p + 1) .. (p' - 1) : \neg isPrime(q)$


$Spec \triangleq Init \wedge \square[Next]_p$

THEOREM $Spec \Rightarrow \square TypeInvariant$

Spec that generates all prime numbers

```
----- MODULE prime -----  
EXTENDS Naturals  
VARIABLE p  
  
isPrime(q) == q > 1  $\wedge \forall r \in 2..(q-1): q \% r \neq 0$   
  
TypeInvariant == isPrime(p)  
  
Init == p = 2  
Next == p' > p  $\wedge$  isPrime(p')  $\wedge \forall q \in (p+1)..(p'-1): \sim$ isPrime(q)  
  
Spec == Init  $\wedge []$  [Next]_p  
  
THEOREM Spec => []TypeInvariant
```

Asynchronous FIFO Channel Specification



$Send \triangleq \wedge rdy' = ack$
 $\wedge val' \in Data$
 $\wedge rdy' = 1 - rdy$
 $\wedge ack' = ack$

$Recv \triangleq \wedge rdy' \neq ack$
 $\wedge ack' = 1 - ack$
 $\wedge val' = val$
 $\wedge rdy' = rdy$

Asynchronous FIFO Channel Specification

$TypeInvariant \triangleq \wedge val \in Data$
 $\wedge rdy \in \{0, 1\}$
 $\wedge ack \in \{0, 1\}$

$Init \triangleq \wedge val \in Data$
 $\wedge rdy \in \{0, 1\}$
 $\wedge ack = rdy$

$Send \triangleq \wedge rdy = ack$
 $\wedge val' \in Data$
 $\wedge rdy' = 1 - rdy$
 $\wedge ack' = ack$

$Recv \triangleq \wedge rdy \neq ack$
 $\wedge ack' = 1 - ack$
 $\wedge val' = val$
 $\wedge rdy' = rdy$


$Next \triangleq Send \vee Recv$

$Spec \triangleq Init \wedge \square[Next]_{\langle rdy, ack, val \rangle}$

Asynchronous FIFO Channel Specification

introducing *operators with arguments*

$$\begin{aligned} \textit{Send} \triangleq & \ \wedge rdy = ack \\ & \ \wedge val' \in \textit{Data} \\ & \ \wedge rdy' = 1 - rdy \\ & \ \wedge ack' = ack \end{aligned}$$



$$\begin{aligned} \textit{Send}(d) \triangleq & \ \wedge rdy = ack \\ & \ \wedge val' = d \\ & \ \wedge rdy' = 1 - rdy \\ & \ \wedge ack' = ack \end{aligned}$$

$$\begin{aligned} \textit{Next} \triangleq & \ \vee \textit{Send} \\ & \ \vee \textit{Recv} \end{aligned}$$



$$\begin{aligned} \textit{Next} \triangleq & \ \vee \exists d \in \textit{Data}: \textit{Send}(d) \\ & \ \vee \textit{Recv} \end{aligned}$$

Asynchronous FIFO Channel Specification

introducing *records*

$TypeInvariant \triangleq chan \in [val: Data, rdy: \{0,1\}, ack: \{0,1\}]$

$Init \triangleq chan.val \in Data \wedge chan.rdy \in \{0,1\} \wedge chan.ack = chan.rdy$

$Send(d) \triangleq chan.rdy = chan.ack \wedge chan' =$
 $[val \mapsto d, rdy \mapsto 1 - chan.rdy, ack \mapsto chan.ack]$

$Recv \triangleq chan.rdy \neq chan.ack \wedge chan' =$
 $[val \mapsto chan.val, rdy \mapsto chan.rdy, ack \mapsto 1 - chan.ack]$

$Next \triangleq \exists d \in Data: Send(d) \vee Recv$

$Spec \triangleq Init \wedge \square [Next]_{chan}$

Some more terms

- A *state function* is an ordinary expression with (unprimed) variables
 - i.e., it is a function of a state to a value
 - note that a variable is a state function
- A *state predicate* is a Boolean state function
- A *temporal formula* is an assertion about behaviors
- A *theorem* of a specification is a temporal formula that holds over every behavior of the specification
- If S is a specification and I is a predicate and $S \Rightarrow \Box I$ is a theorem then we call I an *invariant* of S .

Temporal Formula

Based on Chapter 8 of Specifying Systems

- A *temporal formula* F assigns a Boolean value to a behavior σ
- $\sigma \models F$ means that F holds over σ
- F is a theorem if $\sigma \models F$ holds over all behaviors σ
- If P is a state predicate, then $\sigma \models P$ means that P holds over the first state in σ
- If A is an action, then $\sigma \models A$ means that A holds over the first two states in σ
 - i.e., the first step in σ is an A step
- If A is an action, then $\sigma \models [A]_v$ means that the first step in σ is an A step or a stuttering step with respect to v

□ Always

- $\sigma \models \Box F$ means that F holds over every suffix of σ
- More formally
 - Let σ^{+n} be σ with the first n states removed
 - Then $\sigma \models \Box F \triangleq \forall n \in \mathbb{N}: \sigma^{+n} \models F$

Boolean combinations of temporal formulas

- $\sigma \models (F \wedge G) \triangleq (\sigma \models F) \wedge (\sigma \models G)$
- $\sigma \models (F \vee G) \triangleq (\sigma \models F) \vee (\sigma \models G)$
- $\sigma \models \neg F \triangleq \neg (\sigma \models F)$
- $\sigma \models (F \Rightarrow G) \triangleq (\sigma \models F) \Rightarrow (\sigma \models G)$
- $\sigma \models (\exists r: F) \triangleq \exists r: \sigma \models F$
- $\sigma \models (\forall r \in S: F) \triangleq \forall r \in S: \sigma \models F$ // if S is a constant set

Example

What is the meaning of $\sigma \models \Box((x = 1) \Rightarrow \Box(y > 0))$?

$$\sigma \models \Box((x = 1) \Rightarrow \Box(y > 0))$$

$$\equiv \forall n \in \mathbb{N}: \sigma^{+n} \models ((x = 1) \Rightarrow \Box(y > 0))$$

$$\equiv \forall n \in \mathbb{N}: (\sigma^{+n} \models (x = 1)) \Rightarrow (\sigma^{+n} \models \Box(y > 0))$$

$$\equiv \forall n \in \mathbb{N}: (\sigma^{+n} \models (x = 1)) \Rightarrow (\forall m \in \mathbb{N}: (\sigma^{+n})^{+m} \models (y > 0))$$

If $x = 1$ in some state, then henceforth $y > 0$ in all subsequent states

Not: once $x = 1$, x will always be 1. That would be

$$\sigma \models \Box((x = 1) \Rightarrow \Box(x = 1))$$

Not every temporal formula is a TLA+ formula

- TLA+ formulas are temporal formulas that are *invariant under stuttering*
 - They hold even if you add or remove stuttering steps
- Examples
 - P if P is a state predicate
 - $\Box P$ if P is a state predicate
 - $\Box[A]_v$ if A is an action and v is a state variable (or even state function)
- But not
 - $x' = x + 1$ not satisfied by $[x = 1] \rightarrow [x = 1] \rightarrow [x = 2]$
 - $[x' = x + 1]_x$ satisfied by $[x = 1] \rightarrow [x = 1] \rightarrow [x = 3]$
but not by $[x = 1] \rightarrow [x = 3]$
- Yet $\Box[x' = x + 1]_x$ is a TLA+ formula!

Eventually F

$$\diamond F \triangleq \neg \Box \neg F$$

$$\sigma \models \diamond F$$

$$\equiv \sigma \models \neg \Box \neg F$$

$$\equiv \neg(\sigma \models \Box \neg F)$$

$$\equiv \neg(\forall n \in \mathbb{N}: \sigma^{+n} \models \neg F)$$

$$\equiv \neg(\forall n \in \mathbb{N}: \neg(\sigma^{+n} \models F))$$

$$\equiv \exists n \in \mathbb{N}: (\sigma^{+n} \models F)$$

Eventually an A step occurs that changes v ...

$$\diamond \langle A \rangle_v \triangleq \neg \Box [\neg A]_v$$

$$\sigma \models \diamond \langle A \rangle_v$$

$$\equiv \sigma \models \neg \Box [\neg A]_v$$

$$\equiv \neg(\sigma \models \Box [\neg A]_v)$$

$$\equiv \neg(\forall n \in \mathbb{N}: \sigma^{+n} \models [\neg A]_v)$$

$$\equiv \neg(\forall n \in \mathbb{N}: \sigma^{+n} \models (\neg A \vee v' = v))$$

$$\equiv \exists n \in \mathbb{N}: \sigma^{+n} \models (A \wedge v' \neq v)$$

HourClock revisited

Module HourClock

Variable *hr*

hr is a parameter of the specification HourClock

- $\text{HCini} \triangleq hr \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\}$
- $\text{HCnxt} \triangleq hr' = hr \bmod 12 + 1$
- $\text{HC} \triangleq \text{HCini} \wedge \square[\text{HCnxt}]_{hr}$

HourClock with *liveness* *clock that never stops*

Module HourClock

Variable hr

- $HCini \triangleq hr \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12\}$
- $HCnxt \triangleq hr' = hr \bmod 12 + 1$
- $HC \triangleq HCini \wedge \square[HCnxt]_{hr}$
- $LiveHC \triangleq HC \wedge \square(\diamond \langle HCnxt \rangle_{hr})$

Module Channel with Liveness

Constant *Data*

Variable *chan*

TypeInvariant $\triangleq \text{chan} \in [\text{val}: \text{Data}, \text{rdy}: \{0,1\}, \text{ack}: \{0,1\}]$

Init $\triangleq \text{chan.val} \in \text{Data} \wedge \text{chan.rdy} \in \{0,1\} \wedge \text{chan.ack} = \text{chan.rdy}$

Send(*d*) $\triangleq \text{chan.rdy} = \text{chan.ack} \wedge \text{chan}' =$
 $[\text{val} \mapsto d, \text{rdy} \mapsto 1 - \text{chan.rdy}, \text{ack} \mapsto \text{chan.ack}]$

Recv $\triangleq \text{chan.rdy} \neq \text{chan.ack} \wedge \text{chan}' =$
 $[\text{val} \mapsto \text{chan.val}, \text{rdy} \mapsto \text{chan.rdy}, \text{ack} \mapsto 1 - \text{chan.ack}]$

Next $\triangleq \exists d \in \text{Data}: \text{Send}(d) \vee \text{Recv}$

Spec $\triangleq \text{Init} \wedge \square[\text{Next}]_{\text{chan}}$

LiveSpec $\triangleq \text{Spec} \wedge \square(\diamond \langle \text{Next} \rangle_{\text{chan}}) \text{ ???}$

Module Channel with Liveness

Constant *Data*

Variable *chan*

TypeInvariant $\triangleq \text{chan} \in [\text{val}: \text{Data}, \text{rdy}: \{0,1\}, \text{ack}: \{0,1\}]$

Init $\triangleq \text{chan.val} \in \text{Data} \wedge \text{chan.rdy} \in \{0,1\} \wedge \text{chan.ack} = \text{chan.rdy}$

Send(*d*) $\triangleq \text{chan.rdy} = \text{chan.ack} \wedge \text{chan}' =$
 $[\text{val} \mapsto d, \text{rdy} \mapsto 1 - \text{chan.rdy}, \text{ack} \mapsto \text{chan.ack}]$

Recv $\triangleq \text{chan.rdy} \neq \text{chan.ack} \wedge \text{chan}' =$
 $[\text{val} \mapsto \text{chan.val}, \text{rdy} \mapsto \text{chan.rdy}, \text{ack} \mapsto \text{chan.ack}]$

Too Strong --- If nothing
to send that should be ok

Next $\triangleq \exists d \in \text{Data}: \text{Send}(d) \vee \text{Recv}$

Spec $\triangleq \text{Init} \wedge \square[\text{Next}]_{\text{chan}}$

LiveSpec $\triangleq \text{Spec} \wedge \square(\diamond \langle \text{Next} \rangle_{\text{chan}}) \text{ ???}$

Module Channel with Liveness

Constant *Data*

Variable *chan*

TypeInvariant $\triangleq \text{chan} \in [\text{val}: \text{Data}, \text{rdy}: \{0,1\}, \text{ack}: \{0,1\}]$

Init $\triangleq \text{chan.val} \in \text{Data} \wedge \text{chan.rdy} \in \{0,1\} \wedge \text{chan.ack} = \text{chan.rdy}$

Send(*d*) $\triangleq \text{chan.rdy} = \text{chan.ack} \wedge \text{chan}' =$
 $[\text{val} \mapsto d, \text{rdy} \mapsto 1 - \text{chan.rdy}, \text{ack} \mapsto \text{chan.ack}]$

Recv $\triangleq \text{chan.rdy} \neq \text{chan.ack} \wedge \text{chan}' =$
 $[\text{val} \mapsto \text{chan.val}, \text{rdy} \mapsto \text{chan.rdy}, \text{ack} \mapsto 1 - \text{chan.ack}]$

Next $\triangleq \exists d \in \text{Data}: \text{Send}(d) \vee \text{Recv}$

Spec $\triangleq \text{Init} \wedge \square[\text{Next}]_{\text{chan}}$

LiveSpec $\triangleq \text{Spec} \wedge \square(\text{chan.rdy} \neq \text{chan.ack} \Rightarrow \diamond \langle \text{Recv} \rangle_{\text{chan}})$

Weak Fairness as a liveness condition

- **ENABLED** $\langle A \rangle_v$ means action A is possible in some state
 - State predicate conjuncts all hold *and* some next state must exist
- **WF_v**(A) $\triangleq \Box(\Box_{\text{ENABLED}} \langle A \rangle_v \Rightarrow \Diamond \langle A \rangle_v)$
- HourClock: $WF_{hr}(HCnext)$
- Channel: $WF_{chan}(Recv)$

(surprising) Weak Fairness equivalence

- $WF_v(A) \triangleq \Box(\Box_{\text{ENABLED}} \langle A \rangle_v \Rightarrow \Diamond \langle A \rangle_v)$
 $\equiv \Box \Diamond (\neg_{\text{ENABLED}} \langle A \rangle_v) \vee \Box \Diamond \langle A \rangle_v$
 $\equiv \Diamond \Box (\text{ENABLED} \langle A \rangle_v) \Rightarrow \Box \Diamond \langle A \rangle_v$

- Always, if A is enabled forever, then an A step eventually occurs
- A is infinitely often disabled or infinitely many A steps occur
- If A is eventually enabled forever then infinitely many A steps occur

Strong Fairness

- $SF_v(A) \triangleq \diamond \square (\neg \text{ENABLED } \langle A \rangle_v) \vee \square \diamond \langle A \rangle_v$
 $\equiv \square \diamond (\text{ENABLED } \langle A \rangle_v) \Rightarrow \square \diamond \langle A \rangle_v$

- A is eventually disabled forever or infinitely many A steps occur
- If A is infinitely often enabled then infinitely many A steps occur

$SF_v(A)$: an A step must occur if A is **continually** enabled

$WF_v(A)$: an A step must occur if A is **continuously** enabled

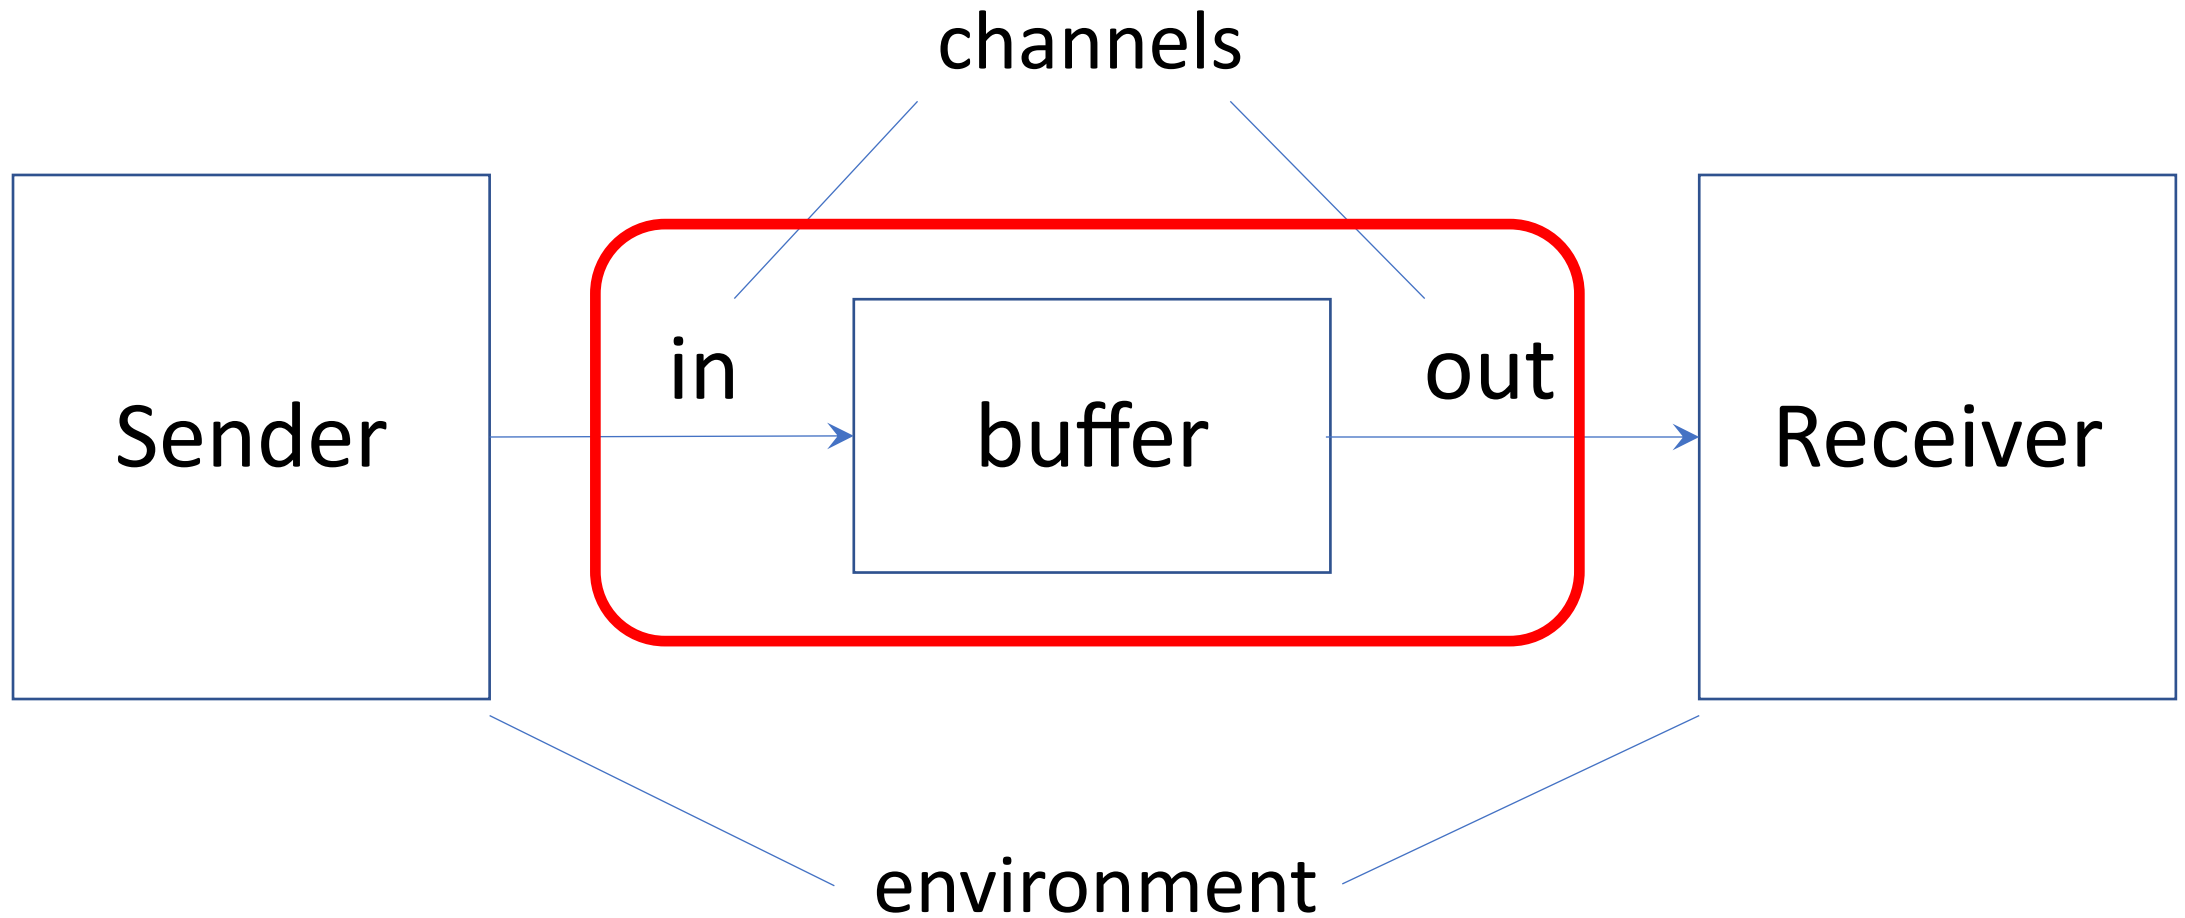
As always, better to make the weaker assumption if you can

How important is liveness?

- Liveness rules out behaviors that have only stuttering steps
 - Add non-triviality of a specification
- In practice, “eventual” is often not good enough
- Instead, need to specify performance requirements
 - Service Level Objectives (SLOs)
 - Usually done quite informally

A “FIFO” (async buffered FIFO channel)

Chapter 4 from Specifying Systems



Module Channel

Constant *Data*

Variable *chan*

TypeInvariant $\triangleq \text{chan} \in [\text{val}: \text{Data}, \text{rdy}: \{0,1\}, \text{ack}: \{0,1\}]$

Init $\triangleq \text{chan.val} \in \text{Data} \wedge \text{chan.rdy} \in \{0,1\} \wedge \text{chan.ack} = \text{chan.rdy}$

Send(d) $\triangleq \text{chan.rdy} = \text{chan.ack} \wedge \text{chan}' =$
 $[\text{val} \mapsto d, \text{rdy} \mapsto 1 - \text{chan.rdy}, \text{ack} \mapsto \text{chan.ack}]$

Recv $\triangleq \text{chan.rdy} \neq \text{chan.ack} \wedge \text{chan}' =$
 $[\text{val} \mapsto \text{chan.val}, \text{rdy} \mapsto \text{chan.rdy}, \text{ack} \mapsto 1 - \text{chan.ack}]$

Next $\triangleq \exists d \in \text{Data}: \text{Send}(d) \vee \text{Recv}$

Spec $\triangleq \text{Init} \wedge \square[\text{Next}]_{\text{chan}}$

Instantiating a Channel

$InChan \triangleq \text{INSTANCE Channel WITH } Data \leftarrow Message, chan \leftarrow in$

$TypeInvariant \triangleq chan \in [val: Data, rdy: \{0,1\}, ack: \{0,1\}]$



$InChan!TypeInvariant \equiv in \in [val: Message, rdy: \{0,1\}, ack: \{0,1\}]$

Instantiation is Substitution!

EXTENDS *Naturals, Sequences*

CONSTANT *Message*

VARIABLES *in, out, q*

InChan \triangleq INSTANCE *Channel* WITH *Data* \leftarrow *Message*, *chan* \leftarrow *in*

OutChan \triangleq INSTANCE *Channel* WITH *Data* \leftarrow *Message*, *chan* \leftarrow *out*

Init \triangleq \wedge *InChan!Init*
 \wedge *OutChan!Init*
 \wedge *q* = $\langle \rangle$

TypeInvariant \triangleq \wedge *InChan!TypeInvariant*
 \wedge *OutChan!TypeInvariant*
 \wedge *q* \in *Seq(Message)*

$SSend(msg) \triangleq \wedge InChan!Send(msg)$ Send msg on channel in .
 $\wedge UNCHANGED \langle out, q \rangle$

$BufRcv \triangleq \wedge InChan!Rcv$ Receive message from channel in
 $\wedge q' = Append(q, in.val)$ and append it to tail of q .
 $\wedge UNCHANGED out$

$BufSend \triangleq \wedge q \neq \langle \rangle$ Enabled only if q is nonempty.
 $\wedge OutChan!Send(Head(q))$ Send $Head(q)$ on channel out
 $\wedge q' = Tail(q)$ and remove it from q .
 $\wedge UNCHANGED in$

$RRcv \triangleq \wedge OutChan!Rcv$ Receive message from channel out .
 $\wedge UNCHANGED \langle in, q \rangle$

$$\begin{aligned} Next &\triangleq \bigvee \exists msg \in Message : SSend(msg) \\ &\quad \bigvee BufRcv \\ &\quad \bigvee BufSend \\ &\quad \bigvee RRcv \end{aligned}$$

$$Spec \triangleq Init \wedge \square[Next]_{\langle in, out, q \rangle}$$

THEOREM $Spec \Rightarrow \square TypeInvariant$

Parametrized Instantiation

$InChan \triangleq \text{INSTANCE Channel WITH Data} \leftarrow \text{Message}, \text{chan} \leftarrow in$



$Chan(ch) \triangleq \text{INSTANCE Channel WITH Data} \leftarrow \text{Message}, \text{chan} \leftarrow ch$

$TypeInvariant \triangleq \text{chan} \in [val: \text{Data}, rdy: \{0,1\}, ack: \{0,1\}]$



$Chan(in)!TypeInvariant \equiv in \in [val: \text{Message}, rdy: \{0,1\}, ack: \{0,1\}]$

Internal (= Non-Interface) Variables



There is no *q* here

MODULE *InnerFIFO*

EXTENDS *Naturals, Sequences*

CONSTANT *Message*

VARIABLES *in, out, q*

But there is a *q* here

Not incorrect, but don't really want *q* to be a specification parameter

Hiding Internal Variables

MODULE *FIFO*

CONSTANT *Message*

VARIABLES *in, out*

$Inner(q) \triangleq$ INSTANCE *InnerFIFO*

$Spec \triangleq \exists q : Inner(q)!Spec$

Hiding Internal Variables

MODULE *FIFO*

CONSTANT *Message*

VARIABLES *in, out*

Inner(*q*) \triangleq INSTANCE *InnerFIFO*

Spec \triangleq $\exists q$: *Inner*(*q*)!*Spec*

Not the normal existential quantifier!!!

In temporal logic, this means that for every state in a behavior, there is a value for *q* that makes *Inner*(*q*)!*Spec* true

Pretty. Now for something cool!

- Suppose we wanted to implemented a bounded buffer
- That is, $\square len(q) \leq N$ for some constant $N > 0$
- The only place where q is extended is in *BufRcv*

$$\begin{aligned} BufRcv &\triangleq \wedge InChan!Rcv \\ &\wedge q' = Append(q, in.val) \\ &\wedge UNCHANGED out \end{aligned}$$

Pretty. Now for something cool!

- Suppose we wanted to implemented a bounded buffer
- That is, $\square len(q) \leq N$ for some constant $N > 0$
- The only place where q is extended is in *BufRcv*

$$\begin{aligned} BufRcv &\triangleq \wedge InChan!Rcv \\ &\wedge q' = Append(q, in.val) \\ &\wedge UNCHANGED out \\ &\wedge len(q) < N \end{aligned}$$

Even cooler (but tricky)

MODULE *BoundedFIFO*

EXTENDS *Naturals, Sequences*

VARIABLES *in, out*

CONSTANT *Message, N*

ASSUME $(N \in \text{Nat}) \wedge (N > 0)$

$\text{Inner}(q) \triangleq \text{INSTANCE } \text{InnerFIFO}$

$\text{BNext}(q) \triangleq \wedge \text{Inner}(q)!\text{Next}$
 $\wedge \text{Inner}(q)!\text{BufRcv} \Rightarrow (\text{Len}(q) < N)$

$\text{Spec} \triangleq \exists q : \text{Inner}(q)!\text{Init} \wedge \square[\text{BNext}(q)]_{\langle in, out, q \rangle}$

If it is a *BufRcv* step,
then $\text{len}(q) < N$

Even cooler (but tricky)

MODULE *BoundedFIFO*

EXTENDS *Naturals, Sequences*

VARIABLES *in, out*

CONSTANT *Message, N*

ASSUME $(N \in \text{Nat}) \wedge (N > 0)$

$\text{Inner}(q) \triangleq \text{INSTANCE } \text{InnerFIFO}$


$\text{BNext}(q) \triangleq \wedge \text{Inner}(q)!\text{Next}$
 $\wedge \text{Inner}(q)!\text{BufRcv} \Rightarrow (\text{Len}(q) < N)$

$\text{Spec} \triangleq \exists q : \text{Inner}(q)!\text{Init} \wedge \square[\text{BNext}(q)]_{\langle in, out, q \rangle}$

Refinement


Based on material from Section 10.8, Specifying Systems by Leslie Lamport

You ask for:




Specification

You ask for:




Specification

You get:




Implementation

You ask for:



Specification

You get:




Implementation


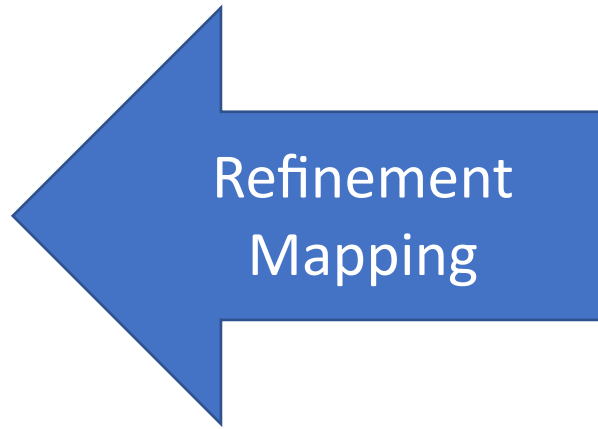
Is every behavior of the implementation also a behavior of the specification?

You ask for:

You get:



Specification




Implementation

Is every behavior of the implementation also a behavior of the specification?

External/internal variables of a state

- A specification has certain *external variables* that can be observed and/or manipulated
- It may also have *internal variables* that are used to describe behaviors but that cannot be observed
- Example: FIFO
 - External variables: in, out
 - Internal variable: buffer.q



Externally visible vs complete behavior

A system may exhibit externally visible behavior

$$e_1 \rightarrow e_2 \rightarrow e_3 \rightarrow e_4 \rightarrow \dots$$

if there exists a complete behavior

$$(e_1, y_1) \rightarrow (e_2, y_2) \rightarrow (e_3, y_3) \rightarrow (e_4, y_4) \rightarrow$$

that is allowed by the specification

Here e_i is some externally visible state (for example, in and out channels) and y_i is internal state (for example, the buffer)

Stuttering Steps

A specification should allow changes to the internal state that does not change the externally visible state.

For example:

$$(e_1, y_1) \rightarrow (e_2, y_2) \rightarrow (e_2, y'_2) \rightarrow (e_3, y_3) \rightarrow (e_4, y_4) \rightarrow$$

leads to external behavior

$$e_1 \rightarrow e_2 \rightarrow e_2 \rightarrow e_3 \rightarrow e_4 \rightarrow \dots$$

which should be identical (up to stuttering) to

$$e_1 \rightarrow e_2 \rightarrow e_3 \rightarrow e_4 \rightarrow \dots$$

Proving that an implementation meets the specification

- First note that an implementation is just a specification
- We call the implementation the “lower-level” specification

We need to prove that if an implementation allows the complete behavior

$$(e_1, z_1) \rightarrow (e_2, z_2) \rightarrow (e_3, z_3) \rightarrow (e_4, z_4) \rightarrow$$

then there exists a complete behavior

$$(e_1, y_1) \rightarrow (e_2, y_2) \rightarrow (e_3, y_3) \rightarrow (e_4, y_4) \rightarrow$$

allowed by the specification

A mapping from low-level complete behaviors to high-level complete behaviors is called a “**refinement mapping**”


Note, there may be multiple possible refinement mappings---you only need to show one

Refinement mapping

1. Map the state of the implementation to the state of the specification
 - Using some function of your choice
2. Show that the initial states of the implementation are also initial states of the specification
 - Using the function above
3. Show that each step of the implementation corresponds to either
 - a state changing step of the specification
 - leaves the specification state unchanged (stuttering step of the spec)

It's not always possible to get a refinement 😞

Binary Consensus, Specification



Paxos

- Value is chosen if a majority of proposers have all accepted the value on the same ballot
- This suggest an easy mapping of the Paxos state to the consensus state

Problem 1: lack of history

- Unfortunately, Paxos acceptors only remember the latest value they accepted
- So, while there may exist a majority that have all accepted the value at time t , that majority may no longer exist at time $t+1$
 - Even though it is guaranteed that no other value will ever be chosen
- But specification does not allow decided state to revert

Fix 1: add history variables

- We can add a “ghost variable” to each acceptor that remembers all (value, ballot) pairs it has ever accepted
 - “ghost” means that it does not actually have to be realized
- With this “history variable”, we can exhibit a state mapping

Problem 2: outrunning the specification

- A refinement mapping maps each step of the low-level specification to either one step of the high-level specification or a stuttering step of the high-level specification
- In Paxos, when $f=1$ and $n=3$, the following scenario is possible:
 - Leader proposes a (value, ballot)
 - Some acceptor accepts (value, ballot)
 - In that one step:
 - The value is chosen
 - The acceptor learns that the value is chosen (decided)
- However, our high-level consensus spec requires two steps:
 - From undecided to chosen and from chosen to learned

Fix 2: two possibilities

- Change the high-level spec to include a “choose + learn” step
 - i.e., speed up the high-level spec
 - complicates the high-level specification
 - changing the specification may not be allowed
- Add a ghost “prophecy variable” to the low-level specification
 - slow down the low-level spec
 - artificially insert a step between accepting and learning by changing the prophecy variable
 - does not change either the implementation or the high-level spec

Completeness

- If $S1$ implements $S2$ then, possibly by adding history and prophecy variables, there exists a refinement mapping from $S2$ to $S1$ (under certain reasonable assumptions)

See Martin Abadi and Leslie Lamport, “The Existence of Refinement Mappings”

Writing Specs

Why specify?

- To avoid errors!
 - writing a spec identifies corner cases
 - allows automated checking
 - model checking / verification
- To clarify communication between designers and builders
 - which avoids errors too...

What to specify?

- Start with the most difficult pieces
 - those pieces that are most likely to have errors in it
- Grain of atomicity
 - Too coarse may fail to reveal important details
 - Too fine may make the spec unwieldy

When to specify

- Ideally before system is implemented
 - find errors early!
- In reality, often implementation provides additional insights that may require changes to the specification
 - try to minimize this---changing the spec a lot wastes dollars and can even kill entire projects
- In practice, not unusual to write spec after an implementation is completed
 - because specs make good documentation

General hints

- Keep it simple, stupid (KISS principle)
 - spec must be clear
- Don't be too abstract
 - may overlook details that are important in a real system
- Write comments