

Modelling and validating distributed systems with TLA+


Carla Ferreira

29th April 2019

TLA+ specification language

- Formal language for describing and reasoning about distributed and concurrent systems.
- TLA+ is a model-oriented language:
 - based on mathematical logic and set theory plus temporal logic TLA (temporal logic of actions).
- Supported by the TLA Toolbox.
- References:
 - TLA+ Hyperbook (<http://research.microsoft.com/en-us/um/people/lamport/tla/hyperbook.html>)
 - TLA+ web page (<http://research.microsoft.com/en-us/um/people/lamport/tla/tla.html>)


LESLIE LAMPORT'S HOME PAGE



[NEW: TLA+ Use at Amazon](#)

[The TLA Web Page](#)

[My Collected Works](#)



Turing Award 2013

For fundamental contributions to the theory and practice of distributed and concurrent systems, notably the invention of concepts such as causality and logical clocks, safety and liveness, replicated state machines, and sequential consistency.

Use of TLA+ at Amazon

“We have used TLA+ on 10 large complex real-world systems. In every case TLA+ has added significant value, either finding subtle bugs that we are sure we would not have found by other means, or giving us enough understanding and confidence to make aggressive performance optimizations without sacrificing correctness.”

Use of TLA+ at Amazon

Applying TLA+ to some of our more complex systems

System	Components	Line count (excl. comments)	Benefit
S3	Fault-tolerant low-level network algorithm	804 PlusCal	Found 2 bugs. Found further bugs in proposed optimizations.
	Background redistribution of data	645 PlusCal	Found 1 bug, and found a bug in the first proposed fix.
DynamoDB	Replication & group-membership system	939 TLA+	Found 3 bugs, some requiring traces of 35 steps
EBS	Volume management	102 PlusCal	Found 3 bugs.
Internal distributed lock manager	Lock-free data structure	223 PlusCal	Improved confidence. Failed to find a liveness bug as we did not check liveness.
	Fault tolerant replication and reconfiguration algorithm	318 TLA+	Found 1 bug. Verified an aggressive optimization.

First TLA+ Example

1-bit Clock

- Clock's possible behaviours:

$b = 1 \longrightarrow b = 0 \longrightarrow b = 1 \longrightarrow b = 0 \longrightarrow \dots$

$b = 0 \longrightarrow b = 1 \longrightarrow b = 0 \longrightarrow b = 1 \longrightarrow \dots$

1-bit Clock

- State variable:

b

- Initial predicate:

$$b = 1 \vee b = 0$$

- Next-step action (b' is the variable at the next state):

$$\vee (b = 0) \wedge (b' = 1)$$

$$\vee (b = 1) \wedge (b' = 0)$$

The initial state and next-step action are formulas in TLA

1-bit Clock

- State variable:

b

- Initial predicate:

$$b = 1 \vee b = 0$$

- Next-step action (b' is the variable at the next state):

IF $b = 0$ THEN $b' = 1$
ELSE $b' = 0$

The initial state and next-step action are formulas in TLA

1-bit Clock: TLA specification

```
----- MODULE OneBitClock -----  
VARIABLE b  
  
Init == (b=0)  $\vee$  (b=1)  
  
Next ==  $\vee$  b = 0  $\wedge$  b' = 1  
        $\vee$  b = 1  $\wedge$  b' = 0
```

What about the clock properties?

System's properties

- Safety
 - Something bad never happens
 - E.g. system never deadlocks, the account balance is greater or equal to zero
- Liveness
 - Something good eventually happens
 - E.g. if a process request access to a critical region it will eventually be granted access, the light will eventually turn green

Let's ignore liveness properties for now

1-bit Clock: TLA specification

----- MODULE OneBitClock -----

VARIABLE b

Init == (b=0) \vee (b=1)

TypeInv == b \in {0,1}

Typing information (TLA+ is untyped)

Next == \vee b = 0 \wedge b' = 1
 \vee b = 1 \wedge b' = 0

Spec == Init \wedge \square [Next]_<>

THEOREM Spec \Rightarrow \square TypeInv
=====

1-bit Clock: TLA specification

----- MODULE OneBitClock -----

VARIABLE b

Init == (b=0) \vee (b=1)

TypeInv == b \in {0,1}

Next == \vee b = 0 \wedge b' = 1
 \vee b = 1 \wedge b' = 0

Spec == Init \wedge \square [Next]_b

THEOREM Spec \Rightarrow \square TypeInv

$$[A]_{\langle\langle f \rangle\rangle} == A \vee (f' = f)$$

The initial state satisfies *Init*
Every transition satisfies *Next* or leaves *b* unchanged

1-bit Clock: TLA specification

```
----- MODULE OneBitClock -----  
VARIABLE b  
  
Init == (b=0)  $\vee$  (b=1)  
  
TypeInv == b  $\in$  {0,1}  
  
Next ==  $\vee$  b = 0  $\wedge$  b' = 1  
         $\vee$  b = 1  $\wedge$  b' = 0  
  
Spec == Init  $\wedge$   $\square$ [Next]_<<b>>
```

THEOREM Spec \Rightarrow \square TypeInv

Theorem specifies an invariant property

TLC model checker

- Exhaustive breath-first search of all reachable states
- Finds (one of) the shortest path to the property violation

Computing all possible behaviours

- State graph is a directed graph G
 1. Put into G to the set of all initial states
 2. For every state s in G compute all possible states t such that $s \rightarrow t$ can be a step in a behaviour
 3. For every state t found in step 2 not in G , draw an edge from s to t
 4. Repeat the previous steps until no new states or edges can be added to G

TLC: state space progress

- Diameter
 - Number of states in the longest path of G with no repeated states
- States found
 - Total number of states it examined in step 1 and 2
- Distinct states
 - Number of states that form the set of nodes of G
- Queue size
 - Number of states s in G for which step 2 has not yet been done

1-bit Clock: Model checking

- Checking the 1-bit clock with TLC model checker (demo)

Model Checking Results

▶ ⚙️ 🛑 🔒 🔑 📄

General

Start time: Sun Sep 18 15:47:11 BST 2016

End time: Sun Sep 18 15:47:11 BST 2016

Last checkpoint time:

Current status: Not running

Errors detected: No errors

Fingerprint collision probability: calculated: 2.2E-19, observed: 3.7E

Statistics

State space progress (click column header for graph)

Time	Diameter	States Found	Distinct States	Queue Size
2016-09-18 15:47...	1	4	2	0

Exercise 1

- Define a TLA+ specification of an hour clock
- Check with TLC the typing invariant

TLA+ Overview

TLA+ Module

```
----- MODULE M -----  
EXTENDS M1, ..., Mn  
\* Incorporates the declarations, definitions, assumptions, and theorems from  
\* the modules named M1, ..., Mn into the current module.  
  
CONSTANTS C1, ..., Cn \* Declares the C1, ..., Cn to be constant parameters.  
  
ASSUME P \* Asserts P as an assumption.  
  
VARIABLES x1, ..., xn \* Declares x1, ..., xn as variables.  
  
TypeInv == exp \* Declares the types of variables x1, ..., xn.  
  
Init == exp \* Initializes variables x1, ..., xn.  
  
F(x1, ..., xn) == exp  
\* Defines F to be an operator such that  
\* F(e1, ..., en) equals exp with each identifier xk replaced by ek.  
  
f[x \in S] == exp  
\* Defines f to be the function with domain S such that f[x] = exp  
\* for all x in S.  
\* The symbol f may occur in exp, allowing a recursive definition.  
  
THEOREM P  
\* Asserts that P can be proved from the definitions and assumptions of the  
\* current module.  
=====
```

TLA+ syntax and semantics

- Logic
- Sets
- Functions
- Tuples, sequences and records
- EXCEPT, UNION, and CHOOSE operators

Logic

$\wedge \vee \neg \Rightarrow \equiv$
TRUE FALSE BOOLEAN [the set {TRUE, FALSE}]
 $\forall x \in S : p$ $\exists x \in S : p$

$\sim(\text{TRUE} \wedge b)$

$a \Rightarrow b$

$\text{Next} == b' = 0$

$b \in \text{BOOLEAN}$

$x \notin S$

$\forall x \in \{1, 2, 3, 4, 5\} : x \geq 0$

$\exists x \in \{1, 2, 3, 4, 5\} : x \% 2 = 0$

Sets

$=$	\neq	\in	\notin	\cup	\cap	\subseteq	\setminus [set difference]
$\{e_1, \dots, e_n\}$							[Set consisting of elements e_i]
$\{x \in S : p\}$	⁽²⁾						Set of elements x in S satisfying p]
$\{e : x \in S\}$	⁽¹⁾						Set of elements e such that x in S]
SUBSET	S						Set of subsets of S]
UNION	S						[Union of all elements of S]

$S = \{1, 2, 3\}$

$S \neq \{1, 2, 3\}$ $S \neq \{1, 2, 3\}$

$x \in S$

$x \notin S$

$S \cup \{1, 2, 3\}$

$\{n \in \{1, 2, 3, 4, 5\} : n \% 2 \neq 0\} = \{1, 3, 5\}$

$\{2*n+1 : n \in \{1, 2, 3, 4, 5\}\} = \{3, 5, 7, 9, 11\}$

UNION $\{\{1, 2\}, \{2, 3\}, \{3, 4\}\} = \{1, 2, 3, 4\}$

SUBSET $\{1, 2\} = \{\{\}, \{1\}, \{2\}, \{1, 2\}\}$

CHOOSE

CHOOSE $x \in S : P(x)$

* Equals some value v in S such that $P(v)$ equals true, if such a value exists.

* Its value is unspecified if no such v exists

CHOOSE $x \in \{1, 2, 3, 4, 5\} : \text{TRUE}$

CHOOSE $x \in \{1, 2, 3, 4, 5\} : x \% 2 = 0$

CHOOSE is deterministic!

CHOICE vs. non-determinism

```
removeOneDet ==  
IF procs \= {}  
THEN procs' =  
    procs \ {CHOOSE t \in procs : TRUE}  
ELSE UNCHANGED procs
```

Deterministic

a single successor state

```
removeOneNonDet ==  
IF procs \= {}  
THEN \E x \in procs : procs' = procs \ {x}  
ELSE UNCHANGED waiting
```

Non-deterministic

many of successor states

Functions

$f[e]$	[Function application]
DOMAIN f	[Domain of function f]
$[x \in S \mapsto e]$	[Function f such that $f[x] = e$ for $x \in S$]
$[S \rightarrow T]$	[Set of functions f with $f[x] \in T$ for $x \in S$]
$[f \text{ EXCEPT } ![e_1] = e_2]$	[Function \hat{f} equal to f except $\hat{f}[e_1] = e_2$. An @ in e_2 equals $f[e_1]$.]

$[i \ \backslash \text{in } \{2,3,5,9\} \mapsto i - 7] = (2 \text{ :> } -5 \ @\@ \ 3 \text{ :> } -4 \ @\@ \ 5 \text{ :> } -2 \ @\@ \ 9 \text{ :> } 2)$

DOMAIN $[i \ \backslash \text{in } \{2,3,5,9\} \mapsto i - 7] = \{2, 3, 5, 9\}$

$[[i \ \backslash \text{in } \{2,3,5,9\} \mapsto i - 7][3] = -4$

$[\{2,4\} \rightarrow \{ "a", "b" \}] = \{ (2 \text{ :> } "a" \ @\@ \ 4 \text{ :> } "a"), (2 \text{ :> } "a" \ @\@ \ 4 \text{ :> } "b"),$
 $(2 \text{ :> } "b" \ @\@ \ 4 \text{ :> } "a"), (2 \text{ :> } "b" \ @\@ \ 4 \text{ :> } "b") \}$

$[[i \ \backslash \text{in } \{2,3,5,9\} \mapsto i - 7] \text{ EXCEPT } ![2]= 12] =$
 $(2 \text{ :> } 12 \ @\@ \ 3 \text{ :> } -4 \ @\@ \ 5 \text{ :> } -2 \ @\@ \ 9 \text{ :> } 2)$

Records

$e.h$	[The h -component of record e]
$[h_1 \mapsto e_1, \dots, h_n \mapsto e_n]$	[The record whose h_i component is e_i]
$[h_1 : S_1, \dots, h_n : S_n]$	[Set of all records with h_i component in S_i]
$[r \text{ EXCEPT } !.h = e]$ ³⁾	[Record \hat{r} equal to r except $\hat{r}.h = e$. An @ in e equals $r.h$.]

[node \mapsto "n1", edge \mapsto "e1"]

[node \mapsto "n1", edge \mapsto "e1"].edge = "e1"

[nodes : {"n1", "n2"}, edges : {"e1", "e2"}]

[node \mapsto "n1", edge \mapsto "e1" EXCEPT !.edge = "xpto"] =

[node \mapsto "n1", edge \mapsto "xpto"]

Tuples

$e[i]$	[The i^{th} component of tuple e]
$\langle e_1, \dots, e_n \rangle$	[The n -tuple whose i^{th} component is e_i]
$S_1 \times \dots \times S_n$	[The set of all n -tuples with i^{th} component in S_i]

$\langle\langle\text{"ana"}, 32, 37495\rangle\rangle$

$\langle\langle\text{"ana"}, 32\rangle\rangle[2] = 32$

$\langle\langle\text{"ana"}, 32\rangle\rangle[1] = \text{"ana"}$

$\{1, 2, 3\} \times \{\text{"a"}, \text{"b"}\} = \{ \langle\langle 1, \text{"a"}\rangle\rangle, \langle\langle 1, \text{"b"}\rangle\rangle, \langle\langle 1, \text{"c"}\rangle\rangle, \\ \langle\langle 2, \text{"a"}\rangle\rangle, \langle\langle 2, \text{"b"}\rangle\rangle, \langle\langle 2, \text{"c"}\rangle\rangle, \\ \langle\langle 3, \text{"a"}\rangle\rangle, \langle\langle 3, \text{"b"}\rangle\rangle, \langle\langle 3, \text{"c"}\rangle\rangle \}$

Sequences

```
----- MODULE Sequences -----  
LOCAL INSTANCE Naturals  
  
Seq(S) == UNION {[1..n -> S] : n \in Nat}  
  
Len(s) == CHOOSE n \in Nat : DOMAIN s = 1..n  
  
s \o t == [i \in 1..(Len(s) + Len(t)) |-> IF i \leq Len(s) THEN s[i]  
                                           ELSE t[i-Len(s)]]  
  
Append(s, e) == s \o <<e>>  
  
Head(s) == s[1]  
  
Tail(s) == [i \in 1..(Len(s)-1) |-> s[i+1]]  
  
SubSeq(s, m, n) == [i \in 1..(1+n-m) |-> s[i+m-1]]  
=====
```


Crossing the river



- A farmer is on one shore of a river and has with him a fox, a chicken, and a sack of grain.
- He has a boat that fits one item besides himself.
- In the presence of the farmer nothing gets eaten, but if left without the farmer, the fox will eat the chicken, and the chicken will eat the grain.
- How can the farmer get all three items across the river safely?

Exercise: Crossing the river

- Define a TLA+ specification for this problem.
- Check with TLC the typing invariant.
- Add an invariant stating that is not possible to get all three items across the river.
- Use TLC to find a solution to this problem.

Only allow safe operations

Exercise: Crossing the river

- Some help:

```
----- MODULE CrossingRiver -----  
EXTENDS Integers  
  
CONSTANTS Farmer, Fox, Chicken, Grain  
  
Items == {Fox, Chicken, Grain}  
  
safe(S) == ~({Fox, Chicken} \subseteq S \vee {Chicken, Grain} \subseteq S)  
  
VARIABLES onLeftShore, onRightShore  
  
TypeInv ==  
  /\ onLeftShore \in SUBSET (Items \union {Farmer})  
  /\ onRightShore \in SUBSET (Items \union {Farmer})
```

Crossing the river: Solution

