

A Tutorial Introduction to TLA⁺

Stephan Merz

<http://www.loria.fr/~merz/>

INRIA Nancy & LORIA
Nancy, France



TLA⁺ Community Event, ABZ 2014
Toulouse, June 3, 2014

Objective


- Explain basic concepts of TLA⁺
 - ▶ modeling systems: static and dynamic aspects
 - ▶ existing tool support for modeling and analysis
PlusCal translator, TLC model checker, TLAPS proof platform
 - ▶ elementary aspects of system refinement

- Example-driven presentation, not trying to be exhaustive

Outline


- 1 Modeling Systems in TLA⁺
- 2 System Verification
- 3 The PlusCal Algorithm Language
- 4 Refinement in TLA⁺

Example: Distributed Termination Detection




- Nodes arranged on a ring perform some computation
 - ▶ nodes can be active (double circle) or inactive
 - ▶ how can node 0 (master node) detect when all nodes are inactive?

Example: Distributed Termination Detection




- Nodes arranged on a ring perform some computation
 - ▶ nodes can be active (double circle) or inactive
 - ▶ how can node 0 (master node) detect when all nodes are inactive?
- Token-based algorithm
 - ▶ initially: token at master node, who may pass it to its neighbor

Example: Distributed Termination Detection




- Nodes arranged on a ring perform some computation
 - ▶ nodes can be active (double circle) or inactive
 - ▶ how can node 0 (master node) detect when all nodes are inactive?
- Token-based algorithm
 - ▶ initially: token at master node, who may pass it to its neighbor
 - ▶ when a node is inactive, it passes on the token

Example: Distributed Termination Detection




- Nodes arranged on a ring perform some computation
 - ▶ nodes can be active (double circle) or inactive
 - ▶ how can node 0 (master node) detect when all nodes are inactive?
- Token-based algorithm
 - ▶ initially: token at master node, who may pass it to its neighbor
 - ▶ when a node is inactive, it passes on the token
 - ▶ **termination detected when token returns to inactive master node**

Example: Distributed Termination Detection




- Nodes arranged on a ring perform some computation
 - ▶ nodes can be active (double circle) or inactive
 - ▶ how can node 0 (master node) detect when all nodes are inactive?
- Token-based algorithm
 - ▶ initially: token at master node, who may pass it to its neighbor
 - ▶ when a node is inactive, it passes on the token
 - ▶ termination detected when token returns to inactive master node
- **Complication: nodes may send messages, activating receiver**

Dijkstra's Algorithm (EWD 840, 1983)



- Nodes and token colored black or white
 - ▶ master node initiates probe by sending white token


Dijkstra's Algorithm (EWD 840, 1983)



- Nodes and token colored black or white

- ▶ master node initiates probe by sending white token
- ▶ message to higher-numbered node stains sending node


Dijkstra's Algorithm (EWD 840, 1983)



- Nodes and token colored black or white

- ▶ master node initiates probe by sending white token
- ▶ message to higher-numbered node stains sending node
- ▶ **when passing the token, a black node stains the token**

Dijkstra's Algorithm (EWD 840, 1983)




- Nodes and token colored black or white

- ▶ master node initiates probe by sending white token
- ▶ message to higher-numbered node stains sending node
- ▶ when passing the token, a black node stains the token

- Termination detection by master node

- ▶ white token at inactive, white master node

Dijkstra's Algorithm (EWD 840, 1983)



- Nodes and token colored black or white

- ▶ master node initiates probe by sending white token
- ▶ message to higher-numbered node stains sending node
- ▶ when passing the token, a black node stains the token

- Termination detection by master node

- ▶ white token at inactive, white master node

- Required correctness properties

- ▶ **safety:** termination detected only if all nodes inactive
- ▶ **liveness:** when all nodes inactive, termination will be detected

TLA⁺ Specification of EWD 840: Data Model

MODULE *EWD840*

EXTENDS *Naturals*

CONSTANT *N*

ASSUME *NAssumption* $\triangleq N \in \text{Nat} \setminus \{0\}$

Nodes $\triangleq 0..N - 1$

Color $\triangleq \{\text{"white"}, \text{"black"}\}$

VARIABLES *tpos, tcolor, active, color*

TypeOK $\triangleq \wedge tpos \in \text{Nodes} \wedge tcolor \in \text{Color}$

$\wedge active \in [\text{Nodes} \rightarrow \text{BOOLEAN}] \wedge color \in [\text{Nodes} \rightarrow \text{Color}]$

- Declaration of parameters
- Definition of operators
 - ▶ sets *Nodes* and *Color*
 - ▶ *TypeOK* documents expected values of variables
 - ▶ *active* and *color* are arrays, i.e. functions

TLA⁺ Specification of EWD 840: Behavior (1)

$$\text{Init} \triangleq \wedge tpos \in \text{Nodes} \wedge tcolor = \text{"black"}$$
$$\wedge \text{active} \in [\text{Nodes} \rightarrow \text{BOOLEAN}] \wedge \text{color} \in [\text{Nodes} \rightarrow \text{Color}]$$

- **Initial condition:** any “type-correct” values; token should be black

TLA⁺ Specification of EWD 840: Behavior (1)

$Init \triangleq \wedge tpos \in Nodes \wedge tcolor = \text{"black"}$

$\wedge active \in [Nodes \rightarrow \text{BOOLEAN}] \wedge color \in [Nodes \rightarrow Color]$

$InitiateProbe \triangleq$

$\wedge tpos = 0 \wedge (tcolor = \text{"black"} \vee color[0] = \text{"black"})$

$\wedge tpos' = N - 1 \wedge tcolor' = \text{"white"}$

$\wedge color' = [color \text{ EXCEPT } ![0] = \text{"white"}]$

$\wedge active' = active$

$PassToken(i) \triangleq$

$\wedge tpos = i \wedge \neg active[i]$

$\wedge tpos' = i - 1$

$\wedge tcolor' = \text{IF } color[i] = \text{"black"} \text{ THEN "black" ELSE } tcolor$

$\wedge color' = [color \text{ EXCEPT } ![i] = \text{"white"}]$

$\wedge active' = active$

- **Initial condition:** any “type-correct” values; token should be black
- **Action definitions:** describe transitions of the algorithm

TLA⁺ Specification of EWD 840: Behavior (2)

$$\begin{aligned} \text{SendMsg}(i) &\triangleq \\ &\wedge \text{active}[i] \\ &\wedge \exists j \in \text{Nodes} \setminus \{i\} : \\ &\quad \wedge \text{active}' = [\text{active EXCEPT } ![j] = \text{TRUE}] \\ &\quad \wedge \text{color}' = [\text{color EXCEPT } ![i] = \text{IF } j > i \text{ THEN "black" ELSE @}] \\ &\wedge \text{UNCHANGED } \langle tpos, tcolor \rangle \\ \text{Deactivate}(i) &\triangleq \\ &\wedge \text{active}[i] \wedge \text{active}' = [\text{active EXCEPT } ![i] = \text{FALSE}] \\ &\wedge \text{UNCHANGED } \langle \text{color}, tpos, tcolor \rangle \end{aligned}$$

- Definition of remaining actions

TLA⁺ Specification of EWD 840: Behavior (2)

$SendMsg(i) \triangleq$

$\wedge active[i]$

$\wedge \exists j \in Nodes \setminus \{i\} :$

$\wedge active' = [active \text{ EXCEPT } ![j] = TRUE]$

$\wedge color' = [color \text{ EXCEPT } ![i] = IF j > i \text{ THEN "black" ELSE @}]$

$\wedge UNCHANGED \langle tpos, tcolor \rangle$

$Deactivate(i) \triangleq$

$\wedge active[i] \wedge active' = [active \text{ EXCEPT } ![i] = FALSE]$

$\wedge UNCHANGED \langle color, tpos, tcolor \rangle$

$Next \triangleq$

$\vee InitiateProbe \vee \exists i \in Nodes \setminus \{0\} : PassToken(i)$

$\vee \exists i \in Nodes : SendMsg(i) \vee Deactivate(i)$

$vars \triangleq \langle tpos, tcolor, active, color \rangle$

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

- Definition of remaining actions
- Possible executions: initial condition, interleaving of transitions

Modeling a System in TLA⁺

1 Describe the system configurations

- ▶ represent the state of the system by state variables
- ▶ mathematical abstractions: numbers, sets, functions, tuples, ...

Modeling a System in TLA⁺

1 Describe the system configurations

- ▶ represent the state of the system by state variables
- ▶ mathematical abstractions: numbers, sets, functions, tuples, ...

2 Specify system behavior as a state machine $Init \wedge \square[Next]_v$

- ▶ initial condition: **state formula** identifies initial states
- ▶ next-state relation: **action formula** constrains allowed transitions
- ▶ overall spec: **temporal formula** defines system executions
- ▶ $\square[Next]_v$ every transition satisfies *Next* or leaves *v* unchanged

Modeling a System in TLA⁺

1 Describe the system configurations

- ▶ represent the state of the system by state variables
- ▶ mathematical abstractions: numbers, sets, functions, tuples, ...

2 Specify system behavior as a state machine $Init \wedge \Box[Next]_v$

- ▶ initial condition: **state formula** identifies initial states
- ▶ next-state relation: **action formula** constrains allowed transitions
- ▶ overall spec: **temporal formula** defines system executions
- ▶ $\Box[Next]_v$ every transition satisfies *Next* or leaves *v* unchanged

• Specifications (and properties) expressed in mathematical logic

- ▶ formally, specify a universe that contains the modeled system
- ▶ use the power of mathematical logic to decompose the specification

Outline

- 1 Modeling Systems in TLA⁺
- 2 System Verification**
 - Safety Properties
 - Liveness Properties
- 3 The PlusCal Algorithm Language
- 4 Refinement in TLA⁺

Outline

- 1 Modeling Systems in TLA⁺
- 2 **System Verification**
 - Safety Properties
 - Liveness Properties
- 3 The PlusCal Algorithm Language
- 4 Refinement in TLA⁺

Safety Properties in TLA⁺

1 Prove type correctness

- ▶ invariant of the specification:
- ▶ asserts that *TypeOK* is always true during any execution of *Spec*

THEOREM $Spec \Rightarrow \Box TypeOK$

Safety Properties in TLA⁺

1 Prove type correctness

- ▶ invariant of the specification: THEOREM $Spec \Rightarrow \Box TypeOK$
- ▶ asserts that *TypeOK* is always true during any execution of *Spec*

2 Termination detection implies that all nodes are inactive

- ▶ termination detected when white token at inactive, white node 0

$terminationDetected \triangleq$
 $tpos = 0 \wedge tcolor = \text{"white"} \wedge \neg active[0] \wedge color[0] = \text{"white"}$
 $TerminationDetection \triangleq$
 $terminationDetected \Rightarrow \forall i \in Nodes : \neg active[i]$
THEOREM $Spec \Rightarrow \Box TerminationDetection$

- ▶ formally again expressed as an invariant

Model Checking Using TLC

- Define a model: finite instance of TLA⁺ specification
 - ▶ instantiate system parameters by concrete values
for example, create instance for $N = 5$
 - ▶ indicate operator corresponding to system specification
in our example, *Spec*
 - ▶ indicate invariants to verify
formulas *TypeOK* and *TerminationDetection*
 - ▶ run TLC on this model and for these properties
- Fully integrated into the TLA⁺ Toolbox
 - ▶ Eclipse IDE for developing and analyzing TLA⁺ specifications
- Use TLC also for validating the specification

EWD840.tla

Model_1

Model Overview | Advanced Options | Model Checking Results

Model Overview


 What is the behavior spec?

 Initial predicate and next-state relation
Init: Next:
 Temporal formula
Spec
 No Behavior Spec

 What to check?

 Deadlock

 Invariants

Formulas true in every reachable state.

-
- TypeOK
-
-
- TerminationDetection

Add

Edit

Remove

 What is the model?

Specify the values of declared constants.

N <- 5

[Advanced parts of the State constraints.](#)
[Additional definitions, Action constraints.](#)
[Definition override, Additional model values.](#)
 How to run?

47M of 81M

Spec Status : parsed

EWD840.tla Model_1

Model Overview Advanced Options Model Checking Results

Model Checking Results



General

Start time: Sun May 11 18:40:55 CEST 2014

End time: Sun May 11 18:40:59 CEST 2014

Last checkpoint time:

Current status: Not running

Errors detected: No errors

Fingerprint collision probability: calculated: 1.8E-11, observed: 5.9E-14

Statistics

State space progress (click column header for graph)

Time	Diameter	States Found	Distinct States	Queue Size
2014-05-11 18:40:59	17	68267	5166	0

Coverage at 2014-05-11 18:40:59

Module	Location	Count
EWD840	line 33, col 6 to line 33, col 16	896
EWD840	line 34, col 6 to line 34, col 22	896
EWD840	line 35, col 6 to line 35, col 23	896
EWD840	line 36, col 6 to line 36, col 45	896
EWD840	line 45, col 6 to line 45, col 16	1942
EWD840	line 46, col 6 to line 46, col 62	1942
EWD840	line 47, col 6 to line 47, col 23	1942

55M of 81M

Spec Status : **parsed**

Using TLAPS to Prove Safety of EWD 840

- TLAPS: proof assistant for verifying TLA⁺ specifications
 - ▶ interesting specifications cannot be verified fully automatically (for arbitrary instances)
 - ▶ user interaction guides verification
 - ▶ automatic back-end proves discharge leaf obligations

Using TLAPS to Prove Safety of EWD 840

- TLAPS: proof assistant for verifying TLA^+ specifications
 - ▶ interesting specifications cannot be verified fully automatically (for arbitrary instances)
 - ▶ user interaction guides verification
 - ▶ automatic back-end proves discharge leaf obligations
- Proving a simple invariant in TLAPS

THEOREM $TypeOK_inv \stackrel{\Delta}{=} Spec \Rightarrow \Box TypeOK$
 $\langle 1 \rangle 1. Init \Rightarrow TypeOK$
 $\langle 1 \rangle 2. TypeOK \wedge [Next]_{vars} \Rightarrow TypeOK'$
 $\langle 1 \rangle 3. QED \quad \text{BY } \langle 1 \rangle 1, \langle 1 \rangle 2, PTL \text{ DEF } Spec$

- ▶ hierarchical proof language represents proof tree
- ▶ steps can be proved in any order: usually start with QED step
- ▶ invariant follows from steps $\langle 1 \rangle 1$ and $\langle 1 \rangle 2$ by temporal logic

Simple Proofs

- Prove that *Init* implies *TypeOK*

⟨1⟩1. *Init* \Rightarrow *TypeOK*

BY *NAssumption* DEFS *Init, TypeOK, Node, Color*

- ▶ definitions and facts must be cited explicitly
- ▶ this helps manage the size of the search space for backend provers

Simple Proofs

- Prove that *Init* implies *TypeOK*

⟨1⟩1. *Init* \Rightarrow *TypeOK*

BY *NAssumption* DEFS *Init, TypeOK, Node, Color*

- ▶ definitions and facts must be cited explicitly
- ▶ this helps manage the size of the search space for backend provers

- Attempt similar proof for step ⟨1⟩2

⟨1⟩2. *TypeOK* \wedge $[Next]_{vars} \Rightarrow TypeOK'$

BY *NAssumption* DEFS *TypeOK, Next, vars, InitiateProbe, ...*

- ▶ back-end provers don't prove this automatically
- ▶ use TLC to ensure that *TypeOK* is an inductive invariant

TypeOK \wedge $\square [Next]_{vars} \Rightarrow \square TypeOK$

- ▶ decompose proof obligation into simpler steps

Hierarchical Proofs

```
⟨1⟩2.  $TypeOK \wedge [Next]_{vars} \Rightarrow TypeOK'$   
  ⟨2⟩ USE DEF  $TypeOK, Node, Color$   
  ⟨2⟩ SUFFICES ASSUME  $TypeOK, Next$   
      PROVE  $TypeOK'$   
      BY DEFS  $TypeOK, vars$   
  ⟨2⟩1. CASE  $InitiateProbe$   
      BY ⟨2⟩1 DEF  $InitiateProbe$   
  ⟨2⟩2. ASSUME NEW  $i \in Node \setminus \{0\}, PassToken(i)$   
      PROVE  $TypeOK'$   
      BY ⟨2⟩2,  $NAssumption$  DEF  $PassToken$   
  ... similar for remaining actions ...  
  ⟨2⟩ QED BY ⟨2⟩1, ⟨2⟩2, ... DEF  $Next$ 
```

- SUFFICES steps represent backward chaining
- trivial case UNCHANGED $vars$ handled during decomposition
- Toolbox IDE helps with hierarchical decomposition

Verifying Safety: Summing Up

- Model checking for finite instances

- ▶ TLC computes reachable state graph, checking invariants on the fly
- ▶ for termination, the set of reachable states must be finite
- ▶ updates of state variables: $v' = e$ or $v' \in E$
for “computable” expressions e, E (E must evaluate to finite set)
- ▶ finite bounds for quantifiers, set comprehensions, function domains

Verifying Safety: Summing Up

- Model checking for finite instances

- ▶ TLC computes reachable state graph, checking invariants on the fly
- ▶ for termination, the set of reachable states must be finite
- ▶ updates of state variables: $v' = e$ or $v' \in E$
for “computable” expressions e, E (E must evaluate to finite set)
- ▶ finite bounds for quantifiers, set comprehensions, function domains

- Theorem proving for arbitrary instances

- ▶ explicit, hierarchical proofs
- ▶ inductive invariant must be provided by the user
- ▶ main proof effort at action level: supported by automatic backends
- ▶ PTL decision procedure for simple temporal reasoning

Outline

- 1 Modeling Systems in TLA⁺
- 2 System Verification
 - Safety Properties
 - Liveness Properties
- 3 The PlusCal Algorithm Language
- 4 Refinement in TLA⁺

Verifying Liveness (1)

- When all nodes are inactive, termination will be detected
 - ▶ expressed in TLA⁺ using a **leadsto**-formula $F \leadsto G \triangleq \Box(F \Rightarrow \Diamond G)$

$Liveness \triangleq (\forall i \in Nodes : \neg active[i]) \leadsto terminationDetected$

THEOREM $Spec \Rightarrow Liveness$

- ▶ verification using TLC, for 5 nodes

Verifying Liveness (1)

- When all nodes are inactive, termination will be detected

- expressed in TLA⁺ using a **leadsto**-formula $F \leadsto G \triangleq \Box(F \Rightarrow \Diamond G)$

$Liveness \triangleq (\forall i \in Nodes : \neg active[i]) \leadsto terminationDetected$

THEOREM $Spec \Rightarrow Liveness$

- verification using TLC, for 5 nodes
- TLC produces a counter-example that ends in infinite **stuttering**
- $\Box[Next]_{vars}$ allows for steps that do not change *vars*

- Use fairness constraint to ensure progress

- fairness: action will be taken, provided it is long enough enabled

$Spec \triangleq Init \wedge \Box[Next]_{vars} \wedge WF_{vars}(Next)$

- don't stutter indefinitely as long as some transition is possible**

Verifying Liveness (2)

- Verify liveness property for redefined specification
 - ▶ TLC verifies that the property is now satisfied

Verifying Liveness (2)

- Verify liveness property for redefined specification
 - ▶ TLC verifies that the property is now satisfied
- TLC also verifies an undesired liveness property

$$\begin{aligned} \text{AllNodesTerminateIfNoMessages} &\triangleq \\ &\diamond \square [\neg \exists i \in \text{Nodes} : \text{SendMsg}(i)]_{\text{vars}} \Rightarrow \diamond (\forall i \in \text{Nodes} : \neg \text{active}[i]) \end{aligned}$$

- ▶ our fairness requirement is too strong!
- ▶ weaken fairness constraint: only ensure termination detection

$$\begin{aligned} \text{System} &\triangleq \text{InitiateProbe} \vee \exists i \in \text{Nodes} \setminus \{0\} : \text{PassToken}(i) \\ \text{Spec} &\triangleq \text{Init} \wedge \square [\text{Next}]_{\text{vars}} \wedge \text{WF}_{\text{vars}}(\text{System}) \end{aligned}$$

- ▶ TLC now verifies liveness, but no longer termination

- **Too strong fairness constraints are a frequent specification error!**

Verifying Liveness: Summing Up

- Specification of fair state machine $Init \wedge \Box[Next]_v \wedge F$
 - ▶ F can be a conjunction of weak or strong fairness conditions

$ENABLED A \triangleq \exists var' : A$ (var' contains all primed variables in A)

$WF_v(A) \triangleq \Box(\Box ENABLED \langle A \rangle_v \Rightarrow \Diamond \langle A \rangle_v)$

$SF_v(A) \triangleq \Box(\Box \Diamond ENABLED \langle A \rangle_v \Rightarrow \Diamond \langle A \rangle_v)$

- ▶ typically: $Next$ is a disjunction, fairness assumed for some disjuncts
- ▶ don't assume too strong fairness conditions: use TLC for validation

Verifying Liveness: Summing Up

- Specification of fair state machine $Init \wedge \Box[Next]_v \wedge F$
 - ▶ F can be a conjunction of weak or strong fairness conditions

$ENABLED A \triangleq \exists var' : A$ (var' contains all primed variables in A)

$WF_v(A) \triangleq \Box(\Box ENABLED \langle A \rangle_v \Rightarrow \Diamond \langle A \rangle_v)$

$SF_v(A) \triangleq \Box(\Box \Diamond ENABLED \langle A \rangle_v \Rightarrow \Diamond \langle A \rangle_v)$

- ▶ typically: $Next$ is a disjunction, fairness assumed for some disjuncts
- ▶ don't assume too strong fairness conditions: use TLC for validation
- Verifying liveness through model checking
 - ▶ temporal properties such as $\Box \Diamond P, \Diamond \Box P, P \rightsquigarrow Q$ and combinations
 - ▶ verification using TLC is more complex, but still automatic
- Verification through theorem proving
 - ▶ not yet supported by TLAPS: needs quantified temporal logic

Outline

- 1 Modeling Systems in TLA⁺
- 2 System Verification
- 3 The PlusCal Algorithm Language**
- 4 Refinement in TLA⁺

Languages for Describing Algorithms

- TLA⁺: algorithms specified by logical formulas
 - ▶ set-theoretical language for modeling data
 - ▶ fair state machine specified in temporal logic

Languages for Describing Algorithms

- TLA⁺: algorithms specified by logical formulas
 - ▶ set-theoretical language for modeling data
 - ▶ fair state machine specified in temporal logic
- Conventional descriptions of algorithms by pseudo-code
 - ▶ familiar presentations, using imperative-style language
 - ▶ (obviously) effective for conveying algorithmic ideas
 - ▶ neither executable nor mathematically precise
- PlusCal: pseudo-code flavor, but precise and more expressive

PlusCal: Elements of an Algorithm Language

- Language for modeling, not programming concurrent algorithms
- High-level abstractions, precise semantics
- Familiar control structure + non-determinism
- Concurrency: indicate grain of atomicity

PlusCal: Elements of an Algorithm Language

- Language for modeling, not programming concurrent algorithms
- High-level abstractions, precise semantics
 - ▶ use TLA⁺ expressions for modeling data
 - ▶ simple translation of PlusCal to TLA⁺ specification
- Familiar control structure + non-determinism
 - ▶ flavor of imperative language: assignment, loop, conditional, ...
 - ▶ special constructs for non-deterministic choice

either { A } **or** { B } **with** $x \in S$ { A }

- Concurrency: indicate grain of atomicity

PlusCal: Elements of an Algorithm Language

- Language for modeling, not programming concurrent algorithms
- High-level abstractions, precise semantics
 - ▶ use TLA⁺ expressions for modeling data
 - ▶ simple translation of PlusCal to TLA⁺ specification
- Familiar control structure + non-determinism
 - ▶ flavor of imperative language: assignment, loop, conditional, ...
 - ▶ special constructs for non-deterministic choice

either { A } **or** { B } **with** $x \in S$ { A }

- Concurrency: indicate grain of atomicity

- ▶ statements may be labeled req: `try[self] := TRUE;`
- ▶ statements from one label to the next one are executed atomically

Example: Alternating-Bit Protocol in PlusCal

MODULE *AlternatingBit*

EXTENDS *Naturals, Sequences*

CONSTANT *Data*

noData \triangleq CHOOSE $x : x \notin \text{Data}$

(****

--algorithm *AlternatingBit* {

 variables *sndC* = $\langle \rangle$, *ackC* = $\langle \rangle$;

 process (*send* = "sender")

 ...

 process (*rcv* = "receiver")

 ...

 process (*err* = "error")

 ...

}

****)

* BEGIN TRANSLATION

* END TRANSLATION

Example: Alternating-Bit Protocol in PlusCal

MODULE *AlternatingBit*

EXTENDS *Naturals, Sequences*

CONSTANT *Data*

noData \triangleq CHOOSE $x : x \notin \text{Data}$

(****

--algorithm *AlternatingBit* {

 variables *sndC* = $\langle \rangle$, *ackC* = $\langle \rangle$;

 process (*send* = "sender")

 ...

 process (*rcv* = "receiver")

 ...

 process (*err* = "error")

 ...

}

****)

* BEGIN TRANSLATION

* END TRANSLATION

*PlusCal algorithm embedded
within TLA+ module*

Example: Alternating-Bit Protocol in PlusCal

MODULE *AlternatingBit*

EXTENDS *Naturals, Sequences*

CONSTANT *Data*

noData \triangleq CHOOSE $x : x \notin \text{Data}$

(****

--algorithm *AlternatingBit* {

 variables *sndC* = $\langle \rangle$, *ackC* = $\langle \rangle$;

 process (*send* = "sender")

 ...

 process (*rcv* = "receiver")

 ...

 process (*err* = "error")

 ...

}

****)

* BEGIN TRANSLATION

* END TRANSLATION

*PlusCal algorithm embedded
within TLA+ module*

global variable declarations

Example: Alternating-Bit Protocol in PlusCal

MODULE *AlternatingBit*

EXTENDS *Naturals, Sequences*

CONSTANT *Data*

noData \triangleq CHOOSE $x : x \notin \text{Data}$

(****

--*algorithm AlternatingBit* {

variables *sndC* = $\langle \rangle$, *ackC* = $\langle \rangle$;

process (*send* = "sender")

...

process (*rcv* = "receiver")

...

process (*err* = "error")

...

}

****)

* BEGIN TRANSLATION

* END TRANSLATION

*PlusCal algorithm embedded
within TLA+ module*

global variable declarations

*three parallel processes
code to be filled in*

Example: Alternating-Bit Protocol in PlusCal

MODULE *AlternatingBit*

EXTENDS *Naturals, Sequences*

CONSTANT *Data*

noData \triangleq CHOOSE $x : x \notin \text{Data}$

(****

--*algorithm AlternatingBit* {

variables *sndC* = $\langle \rangle$, *ackC* = $\langle \rangle$;

process (*send* = "sender")

...

process (*rcv* = "receiver")

...

process (*err* = "error")

...

}

****)

* *BEGIN TRANSLATION*

* *END TRANSLATION*

PlusCal algorithm embedded within TLA+ module

global variable declarations

three parallel processes code to be filled in

PlusCal translator inserts TLA+ between these lines

PlusCal Code of Processes

```
process (send = "sender")
  variables sending = noData, sBit = 0, lastAck = 0; {
s0:  while (TRUE) {
      with ( $d \in \text{Data}$ ) { sending := d; sBit := 1 - sBit };
s1:  while (lastAck  $\neq$  sBit) {
      either {
        sndC := Append(sndC,  $\langle$  sending, sBit  $\rangle$ );
      } or {
        await (Len(ackC) > 0);
        lastAck := Head(ackC); ackC := Tail(ackC);
      } } }
}  \* end process send
```

PlusCal Code of Processes

```
process (send = "sender")
  variables sending = noData, sBit = 0, lastAck = 0; {
s0:  while (TRUE) {
      with ( $d \in \text{Data}$ ) { sending := d; sBit := 1 - sBit };
s1:  while (lastAck  $\neq$  sBit) {
      either {
        sndC := Append(sndC,  $\langle$  sending, sBit  $\rangle$ );
      } or {
        await (Len(ackC) > 0);
        lastAck := Head(ackC); ackC := Tail(ackC);
      } } }
}  \* end process send
```

initialize local variables

PlusCal Code of Processes

```
process (send = "sender")
  variables sending = noData, sBit = 0, lastAck = 0; {
s0:  while (TRUE) {
      with ( $d \in \text{Data}$ ) { sending := d; sBit := 1 - sBit };
s1:  while (lastAck  $\neq$  sBit) {
      either {
        sndC := Append(sndC,  $\langle$  sending, sBit  $\rangle$ );
      } or {
        await (Len(ackC) > 0);
        lastAck := Head(ackC); ackC := Tail(ackC);
      } } }
}  \* end process send
```

initialize local variables

prepare new data

PlusCal Code of Processes

```
process (send = "sender")
  variables sending = noData, sBit = 0, lastAck = 0; {
s0:  while (TRUE) {
      with ( $d \in \text{Data}$ ) { sending := d; sBit := 1 - sBit };
s1:  while (lastAck  $\neq$  sBit) {
      either {
        sndC := Append(sndC,  $\langle$  sending, sBit  $\rangle$ );
      } or {
        await (Len(ackC) > 0);
        lastAck := Head(ackC); ackC := Tail(ackC);
      } } }
}  \* end process send
```

initialize local variables

prepare new data

while not acknowledged,
either (re)send data or
receive acknowledgement

PlusCal Code of Processes

```
process (send = "sender")
  variables sending = noData, sBit = 0, lastAck = 0; {
s0:  while (TRUE) {
      with (d ∈ Data) { sending := d; sBit := 1 - sBit };
s1:  while (lastAck ≠ sBit) {
      either {
        sndC := Append(sndC, ⟨sending, sBit⟩);
      } or {
        await (Len(ackC) > 0);
        lastAck := Head(ackC); ackC := Tail(ackC);
      } } }
}  \* end process send
```

initialize local variables

prepare new data

*while not acknowledged,
either (re)send data or
receive acknowledgement*

- Code of the two other processes is similar
- Familiar “look and feel” of imperative code

Translation to TLA⁺: System State

- TLA⁺ variables

- ▶ variables corresponding to those declared in PlusCal algorithm
- ▶ “program counter” stores current point of program execution

VARIABLES *sndC*, *ackC*, *pc*, *sending*, *sBit*, *lastAck*, *rcvd*, *rBit*

ProcSet \triangleq {"sender"} \cup {"receiver"} \cup {"error"}

Init \triangleq

\wedge *sndC* = $\langle \rangle$ \wedge *ackC* = $\langle \rangle$

\wedge *sending* = *noData* \wedge *sBit* = 0 \wedge *lastAck* = 0

\wedge *rcvd* = *noData* \wedge *rBit* = 0

\wedge *pc* = [*self* \in *ProcSet* \mapsto CASE *self* = "sender" \rightarrow "s0"

□ *self* = "receiver" \rightarrow "r0"

□ *self* = "error" \rightarrow "e0"]

Translation to TLA⁺: Transitions

$s1 \stackrel{\Delta}{=}$

```
s1: while (lastAck  $\neq$  sBit) {  
  either {  
    sndC := Append(sndC,  $\langle$ sending, sBit $\rangle$ );  
  } or {  
    await (Len(ackC) > 0);  
    lastAck := Head(ackC); ackC := Tail(ackC);  
  }  
}
```

Translation to TLA⁺: Transitions

```
s1: while (lastAck ≠ sBit) {  
    either {  
        sndC := Append(sndC, ⟨sending, sBit⟩);  
    } or {  
        await (Len(ackC) > 0);  
        lastAck := Head(ackC); ackC := Tail(ackC);  
    } }
```

```
s1  $\triangleq$   
  ∧ pc["sender"] = "s1"  
  ∧ IF lastAck ≠ sBit  
    THEN ∧ ∨ ∧ sndC' = Append(sndC, ⟨sending, sBit⟩)  
          ∧ UNCHANGED ⟨ackC, lastAck⟩  
          ∨ ∧ Len(ackC) > 0  
            ∧ lastAck' = Head(ackC)  
            ∧ ackC' = Tail(ackC)  
            ∧ sndC' = sndC  
          ∧ pc' = [pc EXCEPT !["sender"] = "s1"]  
    ELSE ∧ pc' = [pc EXCEPT !["sender"] = "s0"]  
          ∧ UNCHANGED ⟨sndC, ackC, lastAck⟩  
  ∧ UNCHANGED ⟨sending, sBit, rcvd, rBit⟩
```

Fairly direct translation from PlusCal block to TLA⁺ action

Translation to TLA⁺: Tying It All Together

- Define the transition relation of the algorithm
 - ▶ transition relation of process: disjunction of individual transitions
 - ▶ overall next-state relation: disjunction of processes
 - ▶ generalizes to multiple instances of same process type

$$\begin{array}{lll} \mathit{send} \stackrel{\Delta}{=} s0 \vee s1 & \mathit{rcv} \stackrel{\Delta}{=} r0 \vee r1 & \mathit{err} \stackrel{\Delta}{=} e0 \\ \mathit{Next} \stackrel{\Delta}{=} \mathit{send} \vee \mathit{rcv} \vee \mathit{err} & & \end{array}$$

Translation to TLA⁺: Tying It All Together

- Define the transition relation of the algorithm

- ▶ transition relation of process: disjunction of individual transitions
- ▶ overall next-state relation: disjunction of processes
- ▶ generalizes to multiple instances of same process type

$$\begin{aligned} send &\stackrel{\Delta}{=} s0 \vee s1 & rcv &\stackrel{\Delta}{=} r0 \vee r1 & err &\stackrel{\Delta}{=} e0 \\ Next &\stackrel{\Delta}{=} send \vee rcv \vee err \end{aligned}$$

- Define the overall TLA⁺ specification

$$Spec \stackrel{\Delta}{=} Init \wedge \square [Next]_{vars}$$

Translation to TLA⁺: Tying It All Together

- Define the transition relation of the algorithm
 - ▶ transition relation of process: disjunction of individual transitions
 - ▶ overall next-state relation: disjunction of processes
 - ▶ generalizes to multiple instances of same process type

$$\begin{aligned} send &\stackrel{\Delta}{=} s0 \vee s1 & rcv &\stackrel{\Delta}{=} r0 \vee r1 & err &\stackrel{\Delta}{=} e0 \\ Next &\stackrel{\Delta}{=} send \vee rcv \vee err \end{aligned}$$

- Define the overall TLA⁺ specification

$$Spec \stackrel{\Delta}{=} Init \wedge \Box [Next]_{vars}$$

- Extension: fairness conditions per process or label

$$\begin{aligned} \textit{fair process} (send = \textit{"sender"}) & & Spec &\stackrel{\Delta}{=} \dots \wedge WF_{vars}(send) \\ s1:+ \textit{ while} (lastAck \neq sBit) \dots & & Spec &\stackrel{\Delta}{=} \dots \wedge SF_{vars}(s1) \end{aligned}$$

PlusCal: Summing Up

- A gateway drug for programmers (C. Newcombe, Amazon)
 - ▶ retain familiar look and feel of pseudo-code
 - ▶ abstractness and expressiveness through embedded TLA⁺
 - ▶ precision through simple translation to TLA⁺
 - ▶ formal verification via standard TLA⁺ tool set

PlusCal: Summing Up

- A gateway drug for programmers (C. Newcombe, Amazon)
 - ▶ retain familiar look and feel of pseudo-code
 - ▶ abstractness and expressiveness through embedded TLA⁺
 - ▶ precision through simple translation to TLA⁺
 - ▶ formal verification via standard TLA⁺ tool set
- Simplicity of translation induces some limitations
 - ▶ single level of processes can make modeling unnatural
 - ▶ translation dictates rules on where labels must and cannot go
 - ▶ properties must be written in TLA⁺ (probably a feature)

PlusCal: Summing Up

- A gateway drug for programmers (C. Newcombe, Amazon)
 - ▶ retain familiar look and feel of pseudo-code
 - ▶ abstractness and expressiveness through embedded TLA⁺
 - ▶ precision through simple translation to TLA⁺
 - ▶ formal verification via standard TLA⁺ tool set
- Simplicity of translation induces some limitations
 - ▶ single level of processes can make modeling unnatural
 - ▶ translation dictates rules on where labels must and cannot go
 - ▶ properties must be written in TLA⁺ (probably a feature)
- Algorithm language: much better than pseudo-code

Outline

- 1 Modeling Systems in TLA⁺
- 2 System Verification
- 3 The PlusCal Algorithm Language
- 4 Refinement in TLA⁺**

Refinement of System Specifications

- Refining (implementing) specification *Spec* by *Impl*
 - ▶ every behavior allowed by *Impl* is a possible execution of *Spec*
 - ▶ TLA⁺ formalization $Impl \Rightarrow Spec$
 - ▶ systems and properties represented as formulas

Refinement of System Specifications

- Refining (implementing) specification *Spec* by *Impl*
 - ▶ every behavior allowed by *Impl* is a possible execution of *Spec*
 - ▶ TLA⁺ formalization $Impl \Rightarrow Spec$
 - ▶ systems and properties represented as formulas
- Problem: *Impl* will include detail not present in *Spec*
 - ▶ additional variables + extra transitions
 - ▶ many steps of *Impl* will be meaningless for *Spec*

Refinement of System Specifications

- Refining (implementing) specification *Spec* by *Impl*
 - ▶ every behavior allowed by *Impl* is a possible execution of *Spec*
 - ▶ TLA⁺ formalization $Impl \Rightarrow Spec$
 - ▶ systems and properties represented as formulas
- Problem: *Impl* will include detail not present in *Spec*
 - ▶ additional variables + extra transitions
 - ▶ many steps of *Impl* will be meaningless for *Spec*
- Stuttering invariance to the rescue!
 - ▶ extra transitions of *Impl* stutter w.r.t. the variables in *Spec*
 - ▶ stuttering steps are allowed by next-state relation $\square[Next]_{vars}$
 - ▶ infinite stuttering ruled out by fairness conditions

Example: Specifying Data Transmission

MODULE *Transmission*

CONSTANT *Data*

$noData \triangleq \text{CHOOSE } x : x \notin Data$

VARIABLES *sending, rcvd*

$Init \triangleq sending = noData \wedge rcvd = noData$

$Send \triangleq \wedge rcvd = sending$
 $\wedge sending' \in Data$
 $\wedge rcvd' = rcvd$

$Receive \triangleq \wedge rcvd \neq sending$
 $\wedge rcvd' = sending$
 $\wedge sending' = sending$

$vars \triangleq \langle sending, rcvd \rangle$

$Spec \triangleq Init \wedge \square [Send \vee Receive]_{vars}$

• Simple handshake protocol specified as a state machine

- ▶ new data may be sent when previous one has been received
- ▶ no explicit mechanism for transferring data

Implementation Through Alternating Bit Protocol

MODULE *AlternatingBit*

...

INSTANCE *Transmission*

THEOREM *Spec* \Rightarrow *Transmission!Spec*

- Implementation checked by TLC

- ▶ for fixed instances of *Data* and constraints on channel size
- ▶ **exercise:** extend this to fair data transmission
hint: ensure that the channels do not lose all data or acknowledgements

- Stuttering invariance is essential here

- ▶ most transitions of alternating bit protocol stutter on $\langle \textit{sending}, \textit{rcvd} \rangle$

Information Hiding

- TLA⁺ specifications describe state machines
 - ▶ often introduce “implementation detail” for controlling transitions
 - ▶ example: program counter generated by PlusCal translator
 - ▶ internal detail should be hidden from “interface”

Information Hiding

- TLA⁺ specifications describe state machines
 - ▶ often introduce “implementation detail” for controlling transitions
 - ▶ example: program counter generated by PlusCal translator
 - ▶ internal detail should be hidden from “interface”
- In logic, hiding corresponds to existential quantification

$$\begin{aligned} \text{Inner} &\triangleq \text{Init} \wedge \square[\text{Next}]_{\text{vars}} \wedge F \\ \text{Spec} &\triangleq \exists x : \text{Inner} \end{aligned}$$

- ▶ behaves like inner specification, but with variables x hidden

Information Hiding

- TLA⁺ specifications describe state machines
 - ▶ often introduce “implementation detail” for controlling transitions
 - ▶ example: program counter generated by PlusCal translator
 - ▶ internal detail should be hidden from “interface”
- In logic, hiding corresponds to existential quantification

$$\begin{aligned} \text{Inner} &\triangleq \text{Init} \wedge \square[\text{Next}]_{\text{vars}} \wedge F \\ \text{Spec} &\triangleq \exists x : \text{Inner} \end{aligned}$$

- ▶ behaves like inner specification, but with variables x hidden
- Refinement under information hiding
 - ▶ prove $\text{Impl} \Rightarrow \text{Inner}[t/x]$ for showing $\text{Impl} \Rightarrow \text{Spec}$
 - ▶ refinement mapping t : computed from implementation variables

Summing Up

- TLA⁺: Specify systems in logic, from first principles
 - ▶ describe system behavior at appropriate level of abstraction
 - ▶ mathematical logic is flexible and expressive
 - ▶ set theory plus state machine plus temporal logic
 - ▶ no formal distinction between systems and properties
 - ▶ experience shows that this approach scales to practical systems
- Support tools
 - ▶ TLA⁺ Toolbox: editor, syntax/semantic analysis, pretty printer
 - ▶ TLC: explicit-state model checker, checkpointing, parallelization
 - ▶ TLAPS: interactive proof platform with powerful theorem provers
 - ▶ PlusCal translator for generating TLA⁺ specification
- Community: Google group, this workshop!

Going Further

- The TLA⁺ Web page

<http://research.microsoft.com/en-us/um/people/lamport/tla/tla.html>

- Detailed presentations

The Hyperbook

Principles and Specifications of Concurrent Systems

Leslie Lamport
Version of 24 March 2014

The *Principles* and *Specification* Tracks

1 Introduction

- 1.1 Concurrent Computation
- 1.2 Modeling Computation
- 1.3 Specification
- 1.4 Systems and Languages

If you are just starting to read this hyperbook, click here.

Sections colored like this have not yet been written.

2 The One-Bit Clock

- 2.1 The Clock's Behaviors
- 2.2 Describing the Behaviors
- 2.3 Writing the Specification
- 2.4 The Pretty-Printed Version of Your Spec
- 2.5 Checking the Specification
- 2.6 Computing the Behaviors from the Specification
- 2.7 Other Ways of Writing the Behavior Specification
- 2.8 Specifying the Clock in PlusCal

3 The Die Hard Problem

- 3.1 Representing the Problem in TLA⁺
- 3.2 Applying TLC
- 3.3 Expressing the Problem in PlusCal

4 Euclid's Algorithm

- 4.1 The Greatest Common Divisor
 - 4.1.1 Divisors
 - 4.1.2 `CHOOSE` and the Maximum of a Set
 - 4.1.3 The GCD Operator

Specifying Systems

