

An Introduction to TLA⁺


Stephan Merz

<https://members.loria.fr/Stephan.Merz/>

Inria Nancy – Grand Est & LORIA
Nancy, France



TLA⁺ Community Meeting @ ETAPS
Paris, April 2023




PhD 1972 (Brandeis University), Mathematics

- Mitre Corporation, 1962–65
- Marlboro College, 1965–69
- Massachusetts Computer Associates, 1970–77
- SRI International, 1977–85
- Digital Equipment Corporation/Compaq, 1985–2001
- Microsoft Research, since 2001

Pioneer of distributed algorithms **Turing Award 2013**

- Natl. Acad. of Engineering, Natl. Acad. of Sciences, American Acad. of Arts and Sciences
- PODC Influential Paper, ACM SIGOPS Hall of Fame (3x), J.C. Laprie Award (2x), LICS Award, John v. Neumann medal, E.W. Dijkstra Prize, NEC C&C Prize ...
- honorary doctorates: Rennes, Kiel, Lausanne, Lugano, Nancy, Brandeis

TLA⁺ specification language



- describe and verify distributed and concurrent systems
- based on mathematical set theory plus temporal logic TLA
- **TLA⁺ Video Course**
- book: Addison-Wesley, 2003 (free download for personal use)
- IDEs: TLA⁺ Toolbox, Visual Studio Code Extension

Some other publications

- Y. Yu, P. Manolios, L. Lamport: *Model checking TLA⁺ Specifications*. CHARME 1999, LNCS 1703.
- D. Cousineau et al.: *TLA⁺ Proofs*. Formal Methods (FM 2012), LNCS 7436.
- I. Konnov et al.: *TLA⁺ Model Checking Made Symbolic*. OOPSLA 2019.
- S. Merz: *The Specification Language TLA⁺*. Logics of Specification Languages, Springer 2008.

Objective of this presentation


- Explain basic concepts of TLA⁺
- TLA⁺ as a specification language
- Tool support for verification: model checking, proof, refinement
- Running example: distributed termination detection

Please interrupt for questions

Outline

- 1 Distributed Termination Detection
- 2 Checking Properties of the Specification
- 3 Safra's Algorithm for Termination Detection
- 4 Conclusion

Distributed Termination Detection




- Nodes perform some computation

- ▶ a node can be active (double circle) or inactive (simple circle)
- ▶ “master node” 0 wishes to detect when all nodes are inactive

- Relevant transitions

- ▶ active node finishes its computation and terminates
- ▶ master node detects termination

Distributed Termination Detection




- Nodes perform some computation

- ▶ a node can be active (double circle) or inactive (simple circle)
- ▶ “master node” 0 wishes to detect when all nodes are inactive

- Relevant transitions

- ▶ active node finishes its computation and terminates
- ▶ master node detects termination

Distributed Termination Detection



- Nodes perform some computation

- ▶ a node can be active (double circle) or inactive (simple circle)
- ▶ “master node” 0 wishes to detect when all nodes are inactive

- Relevant transitions

- ▶ active node finishes its computation and terminates
- ▶ master node detects termination
- ▶ active node sends a message to some node in the network
- ▶ node receives a message, waking up if inactive

Abstract Transition System for Describing the Problem

- State representation

- ▶ activation status per node
- ▶ number of pending messages
- ▶ termination detected?

$$\begin{aligned} \text{TypeOK} &\triangleq \\ &\wedge \text{active} \in [\text{Nodes} \rightarrow \text{BOOLEAN}] \\ &\wedge \text{pending} \in [\text{Nodes} \rightarrow \text{Nat}] \\ &\wedge \text{termDetect} \in \text{BOOLEAN} \\ \text{terminated} &\triangleq \forall n \in \text{Nodes} : \neg \text{active}[n] \wedge \text{pending}[n] = 0 \end{aligned}$$

Abstract Transition System for Describing the Problem

• State representation

- ▶ activation status per node
- ▶ number of pending messages
- ▶ termination detected?

$$\begin{aligned} \text{TypeOK} &\triangleq \\ &\wedge \text{active} \in [\text{Nodes} \rightarrow \text{BOOLEAN}] \\ &\wedge \text{pending} \in [\text{Nodes} \rightarrow \text{Nat}] \\ &\wedge \text{termDetect} \in \text{BOOLEAN} \\ \text{terminated} &\triangleq \forall n \in \text{Nodes} : \neg \text{active}[n] \wedge \text{pending}[n] = 0 \end{aligned}$$

• Transitions

- ▶ termination of a node
- ▶ sending and receiving of messages
- ▶ termination detection

$$\begin{aligned} \text{RcvMsg}(i) &\triangleq \\ &\wedge \text{pending}[i] > 0 \\ &\wedge \text{active}' = [\text{active EXCEPT } ![i] = \text{TRUE}] \\ &\wedge \text{pending}' = [\text{pending EXCEPT } ![i] = @ - 1] \\ &\wedge \text{UNCHANGED } \text{termDetect} \end{aligned}$$

Abstract Transition System for Describing the Problem

• State representation

- ▶ activation status per node
- ▶ number of pending messages
- ▶ termination detected?

$$\begin{aligned} \text{TypeOK} &\triangleq \\ &\wedge \text{active} \in [\text{Nodes} \rightarrow \text{BOOLEAN}] \\ &\wedge \text{pending} \in [\text{Nodes} \rightarrow \text{Nat}] \\ &\wedge \text{termDetect} \in \text{BOOLEAN} \\ \text{terminated} &\triangleq \forall n \in \text{Nodes} : \neg \text{active}[n] \wedge \text{pending}[n] = 0 \end{aligned}$$

• Transitions

- ▶ termination of a node
- ▶ sending and receiving of messages
- ▶ termination detection

$$\begin{aligned} \text{RcvMsg}(i) &\triangleq \\ &\wedge \text{pending}[i] > 0 \\ &\wedge \text{active}' = [\text{active} \text{ EXCEPT } ![i] = \text{TRUE}] \\ &\wedge \text{pending}' = [\text{pending} \text{ EXCEPT } ![i] = @ - 1] \\ &\wedge \text{UNCHANGED } \text{termDetect} \end{aligned}$$

• Overall specification

$$\text{Spec} \triangleq \text{Init} \wedge \square [\text{Next}]_{\text{vars}} \wedge \text{WF}_{\text{vars}}(\text{DetectTermination})$$

Outline

- 1 Distributed Termination Detection
- 2 Checking Properties of the Specification**
- 3 Safra's Algorithm for Termination Detection
- 4 Conclusion

Expressing Correctness Properties

① Safety properties: “nothing bad ever happens”

- ▶ type correctness

$$Spec \Rightarrow \Box TypeOK$$

TypeOK is true throughout any execution of *Spec*

Expressing Correctness Properties

① Safety properties: “nothing bad ever happens”

- ▶ type correctness

$$Spec \Rightarrow \Box TypeOK$$

TypeOK is true throughout any execution of *Spec*

- ▶ safety of detection

$$Spec \Rightarrow \Box (termDetected \Rightarrow terminated)$$

formally again expressed as an invariant

Expressing Correctness Properties

① Safety properties: “nothing bad ever happens”

- ▶ type correctness

$$Spec \Rightarrow \Box TypeOK$$

TypeOK is true throughout any execution of *Spec*

- ▶ safety of detection

$$Spec \Rightarrow \Box (termDetected \Rightarrow terminated)$$

formally again expressed as an invariant

- ▶ quiescence of the system

$$Spec \Rightarrow \Box (terminated \Rightarrow \Box terminated)$$

Expressing Correctness Properties

1 Safety properties: “nothing bad ever happens”

- ▶ type correctness

$$Spec \Rightarrow \Box TypeOK$$

TypeOK is true throughout any execution of *Spec*

- ▶ safety of detection

$$Spec \Rightarrow \Box (termDetected \Rightarrow terminated)$$

formally again expressed as an invariant

- ▶ quiescence of the system

$$Spec \Rightarrow \Box (terminated \Rightarrow \Box terminated)$$

2 Liveness properties: “something good happens eventually”

- ▶ eventual detection of termination

$$Spec \Rightarrow \Box (terminated \Rightarrow \Diamond termDetected)$$

note: the system isn't guaranteed to terminate

Explicit-State Model Checking Using TLC

- Create a model: finite instance of TLA⁺ specification defined as a configuration
 - ▶ instantiate constant parameters, bound potentially infinite variable values
 - fix constants $N = 4$
 - add state constraint $\forall n \in Nodes : pending[n] \leq 3$
 - ▶ indicate formulas representing system specification and properties to be verified
 - ▶ TLC reports 4,097 distinct states (262,145 for $N = 6$)

Explicit-State Model Checking Using TLC

- Create a model: finite instance of TLA⁺ specification defined as a configuration

- ▶ instantiate constant parameters, bound potentially infinite variable values

fix constants $N = 4$

add state constraint $\forall n \in Nodes : pending[n] \leq 3$

- ▶ indicate formulas representing system specification and properties to be verified
- ▶ TLC reports 4,097 distinct states (262,145 for $N = 6$)

- Exploit the automation of TLC for gaining confidence in the specification

- ▶ check putative (non-)properties and make changes to specification
- ▶ e.g., allow inactive node to send messages

Using TLAPS to Prove Safety Properties

- TLAPS: proof assistant for verifying TLA⁺ specifications
 - ▶ proof effort is independent of the size of the instance
 - ▶ relies on user interaction to guide verification
 - ▶ uses automatic back-ends for discharging proof obligations

Using TLAPS to Prove Safety Properties

- TLAPS: proof assistant for verifying TLA⁺ specifications
 - ▶ proof effort is independent of the size of the instance
 - ▶ relies on user interaction to guide verification
 - ▶ uses automatic back-ends for discharging proof obligations
- TLAPS proof of type correctness

THEOREM *TypeCorrect* \triangleq *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 BY $\langle 1 \rangle 1, \langle 1 \rangle 2, PTL$ DEF *Spec*

- ▶ hierarchical proof language represents proof tree
- ▶ assertion follows from steps $\langle 1 \rangle 1$ and $\langle 1 \rangle 2$ by temporal logic
- ▶ prove non-temporal steps by expanding definitions and/or hierarchical subproofs

Proof of Main Safety Property

- Safety of termination detection is inductive relative to *TypeOK*

$Safe \stackrel{\Delta}{=} termDetect \Rightarrow terminated$

THEOREM $Safety \stackrel{\Delta}{=} Spec \Rightarrow \square Safe$

$\langle 1 \rangle 1. Init \Rightarrow Safe$

$\langle 1 \rangle 2. TypeOK \wedge Safe \wedge [Next]_{vars} \Rightarrow Safe'$

$\langle 1 \rangle 3. QED \quad \text{BY } \langle 1 \rangle 1, \langle 1 \rangle 2, \langle 1 \rangle 3, TypeCorrect, PTL \text{ DEF } Spec$

- ▶ use previously established theorem of type correctness
- ▶ proofs of steps $\langle 1 \rangle 1$ and $\langle 1 \rangle 2$ are similar as before

Proof of Main Safety Property

- Safety of termination detection is inductive relative to *TypeOK*

$Safe \stackrel{\Delta}{=} termDetect \Rightarrow terminated$

THEOREM $Safety \stackrel{\Delta}{=} Spec \Rightarrow \square Safe$

$\langle 1 \rangle 1. Init \Rightarrow Safe$

$\langle 1 \rangle 2. TypeOK \wedge Safe \wedge [Next]_{vars} \Rightarrow Safe'$

$\langle 1 \rangle 3. QED \quad \text{BY } \langle 1 \rangle 1, \langle 1 \rangle 2, \langle 1 \rangle 3, TypeCorrect, PTL \text{ DEF } Spec$

- ▶ use previously established theorem of type correctness
 - ▶ proofs of steps $\langle 1 \rangle 1$ and $\langle 1 \rangle 2$ are similar as before
- Proof of quiescence is similar
 - ▶ proofs of safety properties require minimal temporal logic
 - ▶ automation of TLA⁺ set theory is main concern

Proof of Main Safety Property

- Safety of termination detection is inductive relative to *TypeOK*

$Safe \stackrel{\Delta}{=} termDetect \Rightarrow terminated$

THEOREM $Safety \stackrel{\Delta}{=} Spec \Rightarrow \square Safe$

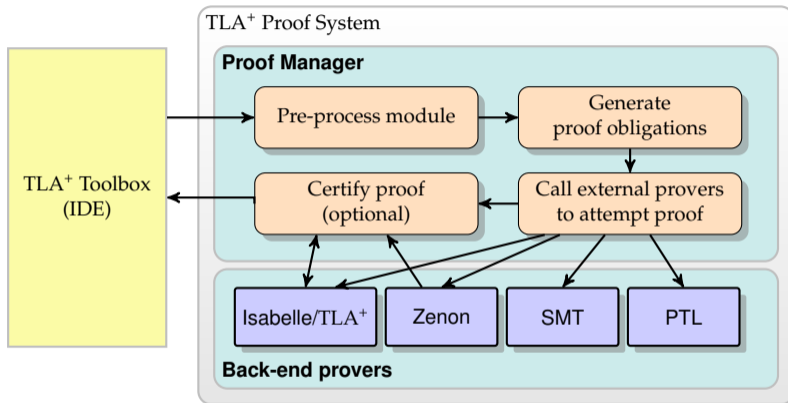
$\langle 1 \rangle 1. Init \Rightarrow Safe$

$\langle 1 \rangle 2. TypeOK \wedge Safe \wedge [Next]_{vars} \Rightarrow Safe'$

$\langle 1 \rangle 3. QED \quad \text{BY } \langle 1 \rangle 1, \langle 1 \rangle 2, \langle 1 \rangle 3, TypeCorrect, PTL \text{ DEF } Spec$

- ▶ use previously established theorem of type correctness
- ▶ proofs of steps $\langle 1 \rangle 1$ and $\langle 1 \rangle 2$ are similar as before
- Proof of quiescence is similar
 - ▶ proofs of safety properties require minimal temporal logic
 - ▶ automation of TLA⁺ set theory is main concern
- Liveness proofs require establishing enabledness predicate
 - ▶ supported in development version of TLAPS

TLAPS Architecture




- Isabelle/TLA⁺: faithful encoding of TLA⁺ in Isabelle's meta-logic
- PTL: decision procedure for propositional temporal logic

Outline

- 1 Distributed Termination Detection
- 2 Checking Properties of the Specification
- 3 Safra's Algorithm for Termination Detection**
- 4 Conclusion


Overall Idea of Safra's algorithm (EWD 998, 1986)

- Token circulating on the ring



Overall Idea of Safra's algorithm (EWD 998, 1986)


- Token circulating on the ring



- ▶ nodes remember difference between numbers of messages sent and received
- ▶ token accumulates sum of differences
- ▶ receiving node becomes “stained”, passing token collects “stain”

Overall Idea of Safra's algorithm (EWD 998, 1986)

- Token circulating on the ring



- ▶ nodes remember difference between numbers of messages sent and received
- ▶ token accumulates sum of differences
- ▶ receiving node becomes “stained”, passing token collects “stain”

- Condition for detecting termination

- ▶ sum of counters at master node and token is zero
- ▶ master node is inactive and clean, and it holds a clean token

Analyzing the Algorithm Using TLC

- Similar correctness properties as for the abstract state machine
 - ▶ type correctness, safety, quiescence, liveness
- Explicit-state model checking with TLC

# nodes	small bounds		modest bounds	
	# states	time	# states	time
3	0.23 M	0:00:09	1.5 M	0:00:21
4	18.7 M	0:03:54	248 M	0:33:53
5	1150 M	2:05:00	–	–

bounds on counters:

- ▶ small: all counters ≤ 2
 - ▶ modest: nodes ≤ 3 , token ≤ 9
- used 32 cores for 5 nodes

Analyzing the Algorithm Using TLC

- Similar correctness properties as for the abstract state machine
 - ▶ type correctness, safety, quiescence, liveness
- Explicit-state model checking with TLC

# nodes	small bounds		modest bounds	
	# states	time	# states	time
3	0.23 M	0:00:09	1.5 M	0:00:21
4	18.7 M	0:03:54	248 M	0:33:53
5	1150 M	2:05:00	–	–

bounds on counters:

- ▶ small: all counters ≤ 2
 - ▶ modest: nodes ≤ 3 , token ≤ 9
- used 32 cores for 5 nodes

- Does this give you enough confidence?
 - ▶ model checking suffers from state space explosion
 - ▶ one modification was incorrect for $N = 4$, but correct for $N = 3$
 - ▶ TLC supports random exploration, finds seeded bugs in majority of runs

An Inductive Invariant for Safra's Algorithm

- Inductive invariant adapted from Dijkstra

$$\begin{aligned} \text{Sum}(f, S) &\triangleq \text{FoldFunctionOnSet}(+, 0, f, S) \\ \text{Inv} &\triangleq \wedge \text{Sum}(\text{pending}, \text{Node}) = \text{Sum}(\text{counter}, \text{Node}) \\ &\wedge \vee \wedge \forall i \in \text{token.pos} + 1 .. N - 1 : \text{active}[i] = \text{FALSE} \\ &\quad \wedge \text{token.q} = \text{Sum}(\text{counter}, (\text{token.pos} + 1) .. (N - 1)) \\ &\quad \vee \text{Sum}(\text{counter}, 0 .. \text{token.pos}) + \text{token.q} > 0 \\ &\quad \vee \exists i \in 0 .. \text{token.pos} : \text{color}[i] = \text{"orange"} \\ &\quad \vee \text{token.color} = \text{"orange"} \end{aligned}$$

An Inductive Invariant for Safra's Algorithm

- Inductive invariant adapted from Dijkstra

$$\begin{aligned} \text{Sum}(f, S) &\triangleq \text{FoldFunctionOnSet}(+, 0, f, S) \\ \text{Inv} &\triangleq \wedge \text{Sum}(\text{pending}, \text{Node}) = \text{Sum}(\text{counter}, \text{Node}) \\ &\quad \wedge \vee \wedge \forall i \in \text{token.pos} + 1 .. N - 1 : \text{active}[i] = \text{FALSE} \\ &\quad \quad \wedge \text{token.q} = \text{Sum}(\text{counter}, (\text{token.pos} + 1) .. (N - 1)) \\ &\quad \quad \vee \text{Sum}(\text{counter}, 0 .. \text{token.pos}) + \text{token.q} > 0 \\ &\quad \quad \vee \exists i \in 0 .. \text{token.pos} : \text{color}[i] = \text{"orange"} \\ &\quad \quad \vee \text{token.color} = \text{"orange"} \end{aligned}$$

- Verification using TLAPS

- ▶ first prove type correctness invariant
- ▶ prove $\text{Spec} \Rightarrow \Box \text{Inv}$, based on type correctness
- ▶ also prove that Inv implies main safety property
- ▶ proofs require auxiliary facts about Sum

Correctness by Refinement

- Specifications and properties are TLA⁺ formulas

▶ THEOREM $Spec \Rightarrow Prop$ every run of $Spec$ satisfies property $Prop$

Correctness by Refinement

- Specifications and properties are TLA⁺ formulas

- ▶ THEOREM $Spec \Rightarrow Prop$

every run of *Spec* satisfies property *Prop*

- ▶ THEOREM $Impl \Rightarrow Spec$

every run of *Impl* corresponds to a run of *Spec*

Correctness by Refinement

- Specifications and properties are TLA⁺ formulas

- ▶ THEOREM $Spec \Rightarrow Prop$ every run of *Spec* satisfies property *Prop*
- ▶ THEOREM $Impl \Rightarrow Spec$ every run of *Impl* corresponds to a run of *Spec*
- ▶ stuttering invariance of TLA⁺ formulas is important here:
allow for low-level steps in *Impl* that are invisible to *Spec*

Correctness by Refinement

- Specifications and properties are TLA⁺ formulas

- ▶ THEOREM $Spec \Rightarrow Prop$ every run of $Spec$ satisfies property $Prop$
- ▶ THEOREM $Impl \Rightarrow Spec$ every run of $Impl$ corresponds to a run of $Spec$
- ▶ stuttering invariance of TLA⁺ formulas is important here:
allow for low-level steps in $Impl$ that are invisible to $Spec$

- Use existing tools for verifying refinement

$TD \stackrel{\Delta}{=} \text{INSTANCE } TerminationDetection$ THEOREM $Spec \Rightarrow TD!Spec$

Correctness by Refinement

- Specifications and properties are TLA⁺ formulas

- ▶ THEOREM $Spec \Rightarrow Prop$ every run of *Spec* satisfies property *Prop*
- ▶ THEOREM $Impl \Rightarrow Spec$ every run of *Impl* corresponds to a run of *Spec*
- ▶ stuttering invariance of TLA⁺ formulas is important here:
allow for low-level steps in *Impl* that are invisible to *Spec*

- Use existing tools for verifying refinement

$TD \stackrel{\Delta}{=} \text{INSTANCE } TerminationDetection$ THEOREM $Spec \Rightarrow TD!Spec$

- ▶ TLC verifies refinement just like it checks correctness properties
- ▶ refinement proof checked by TLAPS, based on previous inductive invariant

Outline

- 1 Distributed Termination Detection
- 2 Checking Properties of the Specification
- 3 Safra's Algorithm for Termination Detection
- 4 Conclusion**

Summing Up

- **TLA⁺: mathematical language for specifying systems**
 - ▶ highly expressive and flexible language encourages abstract descriptions
 - ▶ state machine specifications represent system behavior
 - ▶ no distinction between systems and properties
 - ▶ refinement (and composition) reflected in logic
- **Tool support**
 - ▶ IDEs: TLA⁺ Toolbox / VS Code Extension
 - ▶ TLC: push-button verification, support for random exploration
 - ▶ Apalache: bounded symbolic model checking (see separate tutorial)
 - ▶ TLAPS: interactive proof platform, automatic proof back-ends
 - ▶ PlusCal: front-end for generating TLA⁺ from “pseudo code” language
- **More information**

<http://lamport.azurewebsites.net/tla/tla.html>

Google discussion group