

# Let TLA<sup>+</sup> RiSE

Markus A. Kuppe

August 16, 2018

# Outline

## TLA+

### TLA+

Community

TLA Toolbox, Tools, and TLAPS

## Past, Present, Future Projects

### Scalability

Vertical Scaling

Horizontal Scaling

### Performance

### State Space Reduction

Randomization

Symmetry Reduction

## Conclusion & Outlook

- ▶ High-level specification language
  - ▶ Design above the code level
- ▶ Distributed and concurrent systems
- ▶ ⇒ Team wrote *two* RTOS [?]
  - ▶ (First version flew on the Rosetta spacecraft)
  - ▶ “*We witnessed first hand the brain washing done by years of C programming.*”
  - ▶ “*The TLA<sup>+</sup> abstraction helped a lot in coming to a much cleaner architecture.*”
  - ▶ “*One of the results was that the code size is about 10x less than the previous version.*”

- ▶ Untyped
- ▶ Zermelo Fraenkel set theory with Choice
- ▶ Linear-time framework: Temporal Logic of Actions (TLA)

MODULE *M*

VARIABLE *v*

*Init*  $\triangleq$  ... Defines initial states

*Next*  $\triangleq$   $v' = v + 1 \wedge \dots$  Constrains allowed transitions

*Spec*  $\triangleq$   $Init \wedge \square[Next]_v$  Defines system executions  
 $\wedge F$  and optionally weak or strong Fairness

*Safety*  $\triangleq$   $\square \dots$

*Liveness*  $\triangleq$   $\diamond \square \dots$

# PlusCal

- ▶ Imperative-style pseudo-code but precise
- ▶ Atomicity via labels
- ▶ Can embed TLA<sup>+</sup>
- ▶ Transpiles to TLA<sup>+</sup>
  - ▶ ⇒ Checkable with TLC
- ▶ “A gateway drug for programmers” (C. Newcombe)

---

```
--algorithm Euclid{
  variables x = M, y = N; {
    while (x ≠ y){
      if (x < y){y := y - x}
      else      {x := x - y}
    }
  } Sequential algorithm needs no labels
}
```

---

# Outline

## TLA+

TLA+

**Community**

TLA Toolbox, Tools, and TLAPS

## Past, Present, Future Projects

Scalability

Vertical Scaling

Horizontal Scaling

Performance

State Space Reduction

Randomization

Symmetry Reduction

## Conclusion & Outlook

## Some Adopters

- ▶ Microsoft
- ▶ Amazon
- ▶ Google
- ▶ Intel
- ▶ Oracle
- ▶ Huawei
- ▶ ARM
- ▶ Mongo
- ▶ Thales
- ▶ ...

## Community at large

- ▶ “tlaplus” Google Group with  $\sim 900$  members (almost self-sustaining)
- ▶ Microsoft internal “TLA Plus” group with  $\sim 150$  members
- ▶ GitHub ( $\sim 500$  stars) and  $\sim 10$  contributors
- ▶ Twitter, Reddit, Youtube, ...



- ▶ PlusCal to Go transpiler (Beschastnikh et. al.)
- ▶ Very early stages
  - ▶ Single PlusCal Process
- ▶ <https://github.com/UBC-NSS/pgo>

# BMCMT: Bounded Model Checking of TLA<sup>+</sup> with SMT

- ▶ Abstraction-based parameterized TLA<sup>+</sup> checker
  - ▶ Uses Z3
- ▶ Challenge: “Rich language. Specifications in TLA<sup>+</sup> are considerably more expressive than standard software: TLA<sup>+</sup> is untyped, it allows quantification over sets, comparison of cardinalities, and comparison and updates of the states of concurrent components.” [?]
- ▶ Recording of TLA<sup>+</sup> community event talk:
  - ▶ TLC faster for small models (especially when bound unknown)

# Outline

## TLA<sup>+</sup>

TLA<sup>+</sup>

Community

TLA Toolbox, Tools, and TLAPS

## Past, Present, Future Projects

Scalability

Vertical Scaling

Horizontal Scaling

Performance

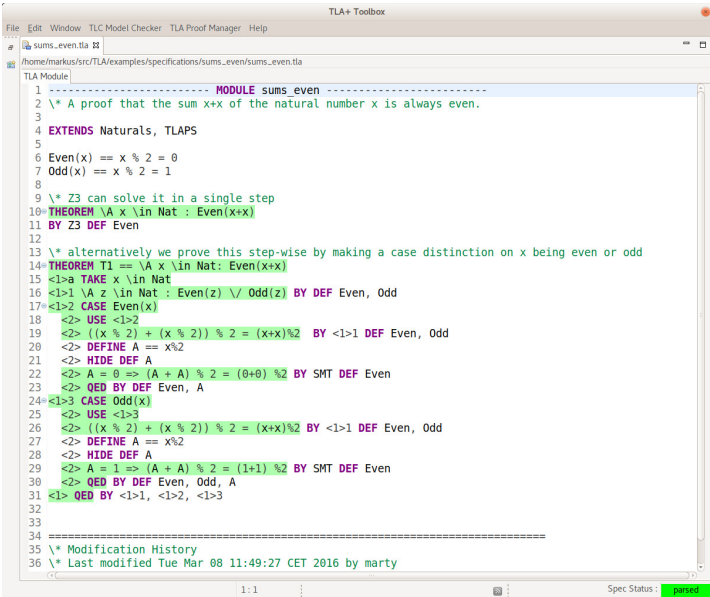
State Space Reduction

Randomization

Symmetry Reduction

## Conclusion & Outlook

# TLA Toolbox, Tools, and TLAPS



```
TLA+ Toolbox
File Edit Window TLC Model Checker TLA Proof Manager Help
sums_even.tla
/home/markus/src/TLA/examples/specifications/sums_even/sums_even.tla
TLA Module
1 ----- MODULE sums_even -----
2 \* A proof that the sum x+x of the natural number x is always even.
3
4 EXTENDS Naturals, TLAPS
5
6 Even(x) == x % 2 = 0
7 Odd(x) == x % 2 = 1
8
9 \* Z3 can solve it in a single step
10 THEOREM \A x \in Nat : Even(x+x)
11 BY Z3 DEF Even
12
13 \* alternatively we prove this step-wise by making a case distinction on x being even or odd
14 THEOREM T1 == \A x \in Nat: Even(x+x)
15 <1>a TAKE x \in Nat
16 <1>1 \A z \in Nat : Even(z) \V Odd(z) BY DEF Even, Odd
17 <1>2 CASE Even(x)
18 <2> USE <1>2
19 <2> ((x % 2) + (x % 2)) % 2 = (x+x)%2 BY <1>1 DEF Even, Odd
20 <2> DEFINE A == x%2
21 <2> HIDE DEF A
22 <2> A = 0 => (A + A) % 2 = (0+0) %2 BY SMT DEF Even
23 <2> QED BY DEF Even, A
24 <1>3 CASE Odd(x)
25 <2> USE <1>3
26 <2> ((x % 2) + (x % 2)) % 2 = (x+x)%2 BY <1>1 DEF Even, Odd
27 <2> DEFINE A == x%2
28 <2> HIDE DEF A
29 <2> A = 1 => (A + A) % 2 = (1+1) %2 BY SMT DEF Even
30 <2> QED BY DEF Even, Odd, A
31 <1> QED BY <1>1, <1>2, <1>3
32
33
34 =====
35 \* Modification History
36 \* Last modified Tue Mar 08 11:49:27 CET 2016 by marty
Spec Status : parsed
```

Figure: TLAPS

# TLA Toolbox, Tools, and TLAPS

The screenshot displays the TLA+ Toolbox interface. On the left, the TLA Module editor shows the specification for the Hanoi Tower problem, including constants for the number of disks (D) and towers (N), and invariants for the total sum of disks and the number of towers. The central pane shows the Model Overview tab with a warning detected. The right pane shows the Model\_1 error trace, indicating that the invariant 'NotSolved is violated'.

```
1 ..... MODULE Hanoi
2 EXTENDS Naturals, Bits, FiniteSets, TLC
3 .....
4 .....
5 (* TRUE iff i is a power of two
6 .....
7 PowerOfTwo(i) == i & (i-1) = 0
8 .....
9 .....
10 (* A set of all powers of two up to n
11 .....
12 SetOfPowerOfTwo(n) == {x \in 1..(2^n-1): PowerOfTwo(x)}
13 .....
14 .....
15 (* Copied from TLA+'s Bags standard library. The sum of
16 (* DOMAIN f.
17 .....
18 .....
19 Sum(f) == LET DSum[S \in SUBSET DOMAIN f] ==
20 LET elt == CHOOSE e \in S : TRUE
21 IN IF S = {} THEN 0
22 ELSE f[elt] + DSum[S \ {e}
23 IN DSum[DOMAIN f]
24 .....
25 (* D is number of disks and N number of towers
26 .....
27 CONSTANT D, N
28 .....
29 (* Towers of Hanoi with N towers
30 .....
31 .....
32 VARIABLES towers
33 vars == <<towers>>
34 .....
35 .....
36 (* The total sum of all towers must amount to the disk
37 .....
38 Inv == Sum(towers) = 2^D - 1
39 .....
40 (* Towers are naturals in the interval (0, 2^D *)
41 TypeOK == /\ \A i \in DOMAIN towers : /\ towers[i] \in
```

**Model Overview** 1 warning detected

**Model description**

**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.

- NotSolved
- Inv
- 

Properties

**What is the model?**

Specify the values of declared constants.

D <- 3

N <- 4

Advanced parts of the model: [Additional State constraints](#), [Action cor](#)

**How to run?**

TLC Parameters

Number of worker threads:

Fraction of physical memory allocated to:

Recover from checkpoint

Checkpoint ID:

Run in distributed mode

**Model\_1**

Invariant NotSolved is violated.

**Error-Trace Exploration**

Enter expressions to be evaluated at each state of the trace

Name	Value
<Initial predicate	State (num = 1)
<towers	<<7, 0, 0, 0>>
<Next line 89, State (num = 2)	
<towers	<<6, 1, 0, 0>>
<Next line 89, State (num = 3)	
<towers	<<4, 1, 2, 0>>
<Next line 89, State (num = 4)	
<towers	<<0, 1, 2, 4>>
<Next line 89, State (num = 5)	
<towers	<<0, 1, 0, 6>>
<Next line 89, State (num = 6)	
<towers	<<0, 0, 0, 7>>

Select line in Error Trace to show its value here.

Figure: Toolbox Model Checking

# TLC

- ▶ Explicit-state model checker for TLA<sup>+</sup>
- ▶ Disk-based (but you don't want it to go to disk)
- ▶ Handles a subclass of TLA<sup>+</sup> that seems to be useful in practice
  - ▶ E.g. no Temporal Existential Quantification, Composition of Actions, ...

- ▶ Safety checking corresponds to Breadth-First search over on-the-fly generated state graph
  - ▶ Fingerprints  $\sim 2^{64}$  (long)
- ▶ Liveness checking corresponds to Depth-First search over (partial) behavior graph [?]
  - ▶ Behavior graph is state graph x tableaux
  - ▶ Technically limited to  $\sim 2^{32}$  vertices

# Outline

TLA<sup>+</sup>

TLA<sup>+</sup>

Community

TLA Toolbox, Tools, and TLAPS

Past, Present, Future Projects

Scalability

Vertical Scaling

Horizontal Scaling

Performance

State Space Reduction

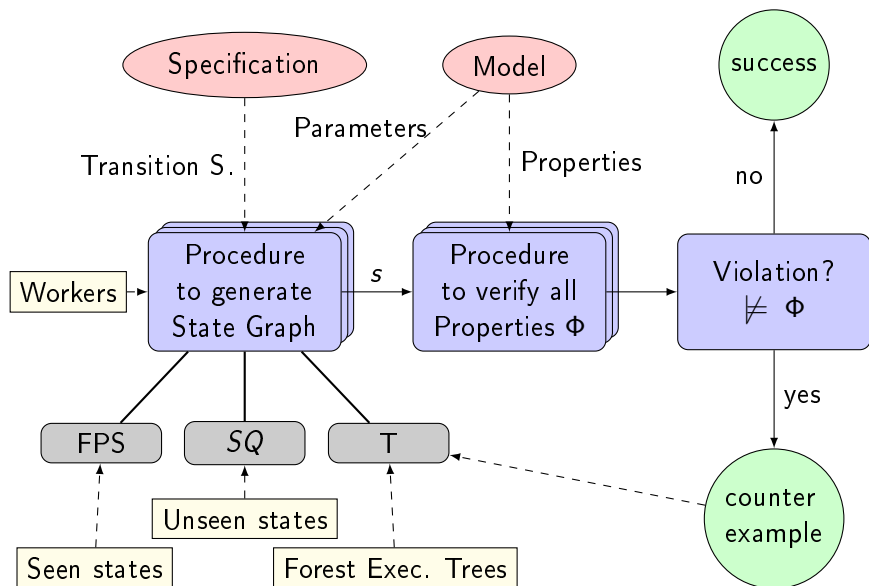
Randomization

Symmetry Reduction

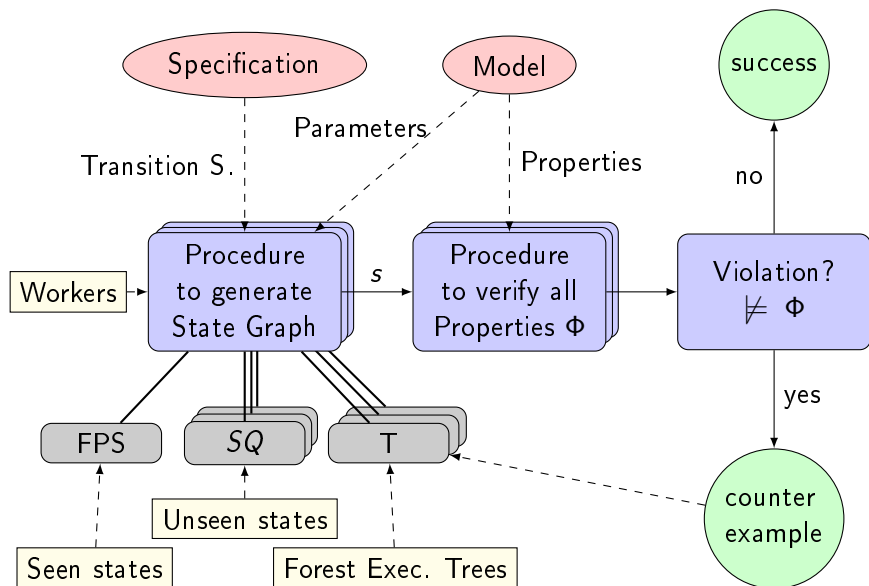
Conclusion & Outlook



# Schematic TLC




# Schematic TLC



# Lock-Striping FPS

1. A global lock for *FPS* to guard **concurrent** find-or-puts does not scale due to *lock contention*
2.  $\Rightarrow$  Partition *FPS* and use one lock per partition

“[...] *lock striping* seems much more promising because the size of the stripe set can be increased as processor counts increase.”



# Lock-Striping FPS

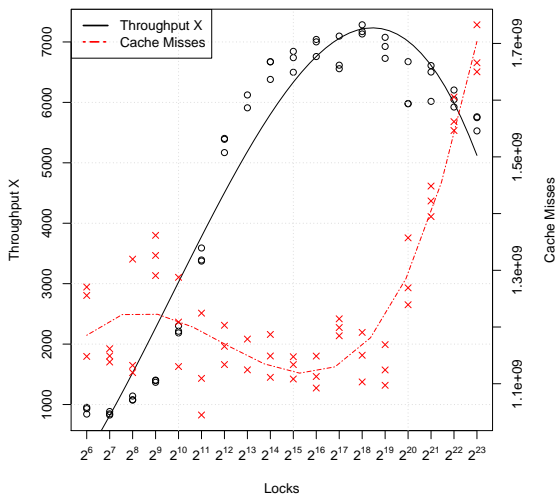


Figure: Lock striping exhibits lock coherence


# Lock-Free & Shared-Nothing

Minimize worker contention via:

- ▶ Lock-Free (CAS) Partitioned/Sharded FPS
  - ▶ Parallel adaptive sort & parallel eviction to disk
  - ▶ Raw memory to avoid GC
- ▶ Shared-Nothing Trace T per worker
- ▶ (Shared-Nothing State Queue SQ)
  - ▶ Overly optimistic assumptions about average shape/properties of state graphs!

# Lock-Free & Shared-Nothing

Dataset: 2017-02-21\_x32 & 2017-02-22\_x32



- ▶ TLC beats scalability of SPIN
  - ▶ (Bakery spec)
- ▶ SPIN outperforms throughput of TLC

# Liveness Checking

- ▶ Check Liveness: Find and check lassos for fulfilling cycles
- ▶ Strongly Connected Components with ?
  - ▶ DFS, linear time, implemented iteratively
- ▶ Liveness checking runs periodically (stops safety)




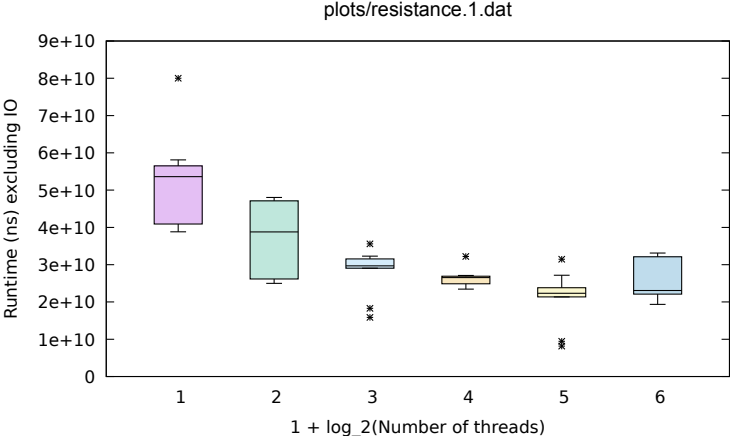
Figure: CPU usage with periodic liveness checking (32 core machine)

# Concurrent SCC

- ▶ R.E. Tarjan drafted a concurrent algorithm for us
  - ▶ Scalability of prototype not promising, abandoned idea for lower hanging fruits
- ▶ “Multi-Core on-the-Fly SCC Decomposition” [?]
  - ▶ GSoC student Parv Mor implemented prototype this summer
  - ▶ Results look more promising
- ▶ Contention/Coherence Union find data-structure?!
- ▶ “A Randomized Concurrent Algorithm for Disjoint Set Union” [?]

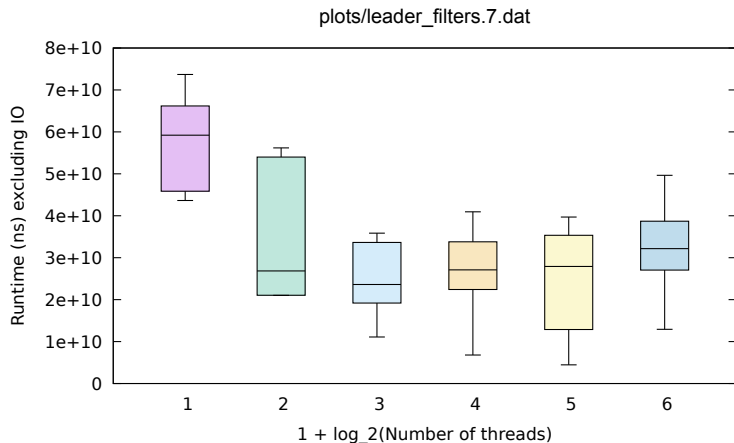


# Concurrent SCC



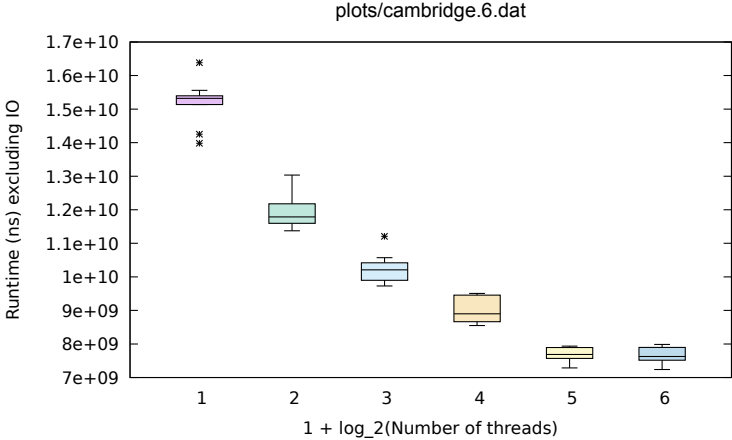
Vertices: 13.8M, Arcs: 32.1M, SCCs: 3 (max 13.8M), Diameter: 60391

# Concurrent SCC



Vertices: 26.3M, Arcs: 91.7M, SCCs: 26.3M (max 1), Diameter: 71


# Concurrent SCC



Vertices: 3.4M, Arcs: 9.5M, SCCs: 8413 (max 3.3M), Diameter: 418

# Distributed TLC

- ▶ Executes TLC on network of machines
- ▶ Distributed Fingerprint Set (DHT)
  - ▶ Nearby memory faster than (local) disks
- ▶ Limitations
  - ▶ Master is bottleneck & SPOF
    - ▶ Checkpointing
  - ▶ No liveness checking
  - ▶ Difficult to setup



# Distributed TLC

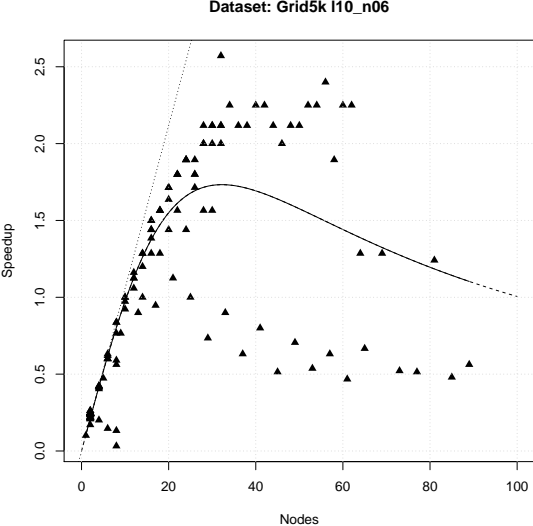


Figure: Scalability distributed TLC:  $Cost/State = 2^{10}$

# Distributed TLC

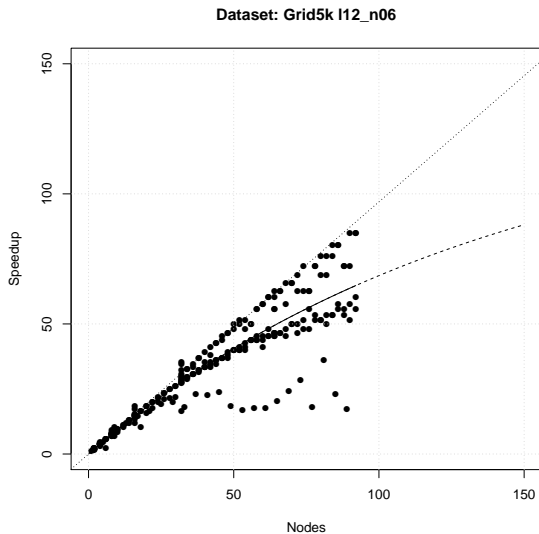


Figure: Scalability distributed TLC:  $Cost/State = 2^{12}$

# Distributed TLC




Figure: Scalability distributed TLC:  $Cost/State = 2^{14}$

# Outline

TLA<sup>+</sup>

TLA<sup>+</sup>

Community

TLA Toolbox, Tools, and TLAPS

## Past, Present, Future Projects

Scalability

Vertical Scaling

Horizontal Scaling

Performance

State Space Reduction

Randomization

Symmetry Reduction

## Conclusion & Outlook



# Cloud TLC

- ▶ Push-Button model checking in the cloud
  - ▶ Support for Azure and AWS
    - ▶ Just compute APIs for portability reasons
    - ▶ Hide away idiosyncrasies of TLC and cloud platform
- ▶ Support for single node TLC and Distributed TLC
- ▶ Can be started from Toolbox and CLI within seconds
  - ▶ Cold-start in the range of minutes
- ▶ Easily check several models concurrently
  - ▶ Instance count is elastic with regards to resource demand
- ▶ Instances dispose automatically after inactivity

# Outline

TLA<sup>+</sup>

TLA<sup>+</sup>

Community

TLA Toolbox, Tools, and TLAPS

## Past, Present, Future Projects

Scalability

Vertical Scaling

Horizontal Scaling

**Performance**

State Space Reduction

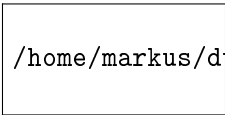
Randomization

Symmetry Reduction

Conclusion & Outlook

## Performance Next-State

- ▶ No intermediate language, no compiler, just AST interpreter
- ▶ Simple left-to-right evaluation of expressions
  - ▶ Recursion but no tail call optimization in Java
- ▶  $\Rightarrow$  Evaluation of next-state at least two orders magnitude slower compared to SPIN



```
/home/markus/dump/syncthing/work/LabInt
```

**Figure:** Throughput (ops/s) normal evaluation (red) vs. module overwrite (blue)

(see online)

# No Partial Evaluation

MODULE *Frob*

VARIABLES  $x, y$

$Init \triangleq x = 0 \wedge y = 0$

$expensiveOp(n) \triangleq \text{CHOOSE } e \in \text{SUBSET } (1 .. n) : \text{TRUE}$

$NextOuch \triangleq \wedge x' \in 1 .. 100$

$\wedge y' = expensiveOp(23)$

$NextYeah \triangleq \wedge y' = expensiveOp(23)$

$\wedge x' \in 1 .. 100$

# TLA<sup>+</sup> Compiler

- ▶ “The Truffle language development framework allows running programming languages efficiently on GraalVM.”<sup>1</sup>
- ▶ “The guest language developer gets a high-performance language implementation, but does not need to be a compiler expert.” [?]
  - ▶ Speedup of evaluation at runtime over special-purpose compilers:
    - ▶ Ruby 3.8x
    - ▶ R 5x
- ▶ Translate AST emitted by SANY to Truffle AST
- ▶ ⇒ Partial Evaluation for TLA<sup>+</sup>

---

<sup>1</sup>GraalVM is a just-in-time compiler for OpenJDK.

# Outline

TLA<sup>+</sup>

TLA<sup>+</sup>

Community

TLA Toolbox, Tools, and TLAPS

## Past, Present, Future Projects

Scalability

Vertical Scaling

Horizontal Scaling

Performance

**State Space Reduction**

Randomization

Symmetry Reduction

Conclusion & Outlook

# Find Inductive Invariant candidates with TLC

Goal: Proof invariance of  $I$  with TLAPS

Find inductive invariant  $Inv$  that satisfies:

1.  $Init \Rightarrow Inv$ , which means that  $Inv$  is true in all initial states.
2.  $Inv \Rightarrow I$ , which means that  $I$  is true in every state on which  $Inv$  is true.
3.  $Inv \wedge Next \Rightarrow Inv'$ , which means if  $Inv$  is true on any state  $s$ , then it is true on any state reachable from  $s$  by a  $Next$  step.

3.1 Let TLC check:

$$CheckInductiveSpec \triangleq Inv \wedge \Box[Next]_{vars}$$

[?]

## Find Inductive Invariant candidates with TLC

MODULE *Foo*

EXTENDS *Naturals*

VARIABLE *x*

$TypeOK \triangleq x \in \text{SUBSET } (1 .. 500)$  or  $x \in \text{Nat}, \dots$

$H \triangleq \dots$  "interesting part"

$Inv \triangleq TypeOK \wedge H$

$CheckInductiveSpec \triangleq Inv \wedge \square[Next]_v$  Make *Inv* the initial predicate.



# New Standard Module Randomization

## MODULE *Randomization*

$\text{RandomSubset}(k, S)$  equals a randomly chosen subset of  $S$  containing  $k$  elements, where  $0 < k < \text{Cardinality}(S)$ .

$\text{RandomSubset}(k, S) \triangleq \text{CHOOSE } T \in \text{SUBSET } S : \text{Cardinality}(T) = k$

$\text{RandomSetOfSubsets}(k, n, S)$  equals a pseudo-randomly chosen set of subsets of  $S$  – that is, a randomly chosen subset of  $\text{SUBSET } S$ . Thus, each element  $T$  of this set is a subset of  $S$ . Each such  $T$  is chosen so that each element of  $S$  has a probability  $n / \text{Cardinality}(S)$  of being in  $T$ . Thus, the average number of elements in each chosen subset  $T$  is  $n$ . The set  $\text{RandomSetOfSubsets}(k, n, S)$  is obtained by making  $k$  such choices of  $T$ . Because this can produce duplicate choices, the number of elements  $T$  in this set may be less than  $k$ .

$\text{RandomSetOfSubsets}(k, n, S) \triangleq$   
 $\text{CHOOSE } T \in \text{SUBSET SUBSET } S : \text{Cardinality}(T) \leq n$

# New Standard Module Randomization

MODULE *Foo*

EXTENDS *Integers, Randomization*

VARIABLE *x*

$TypeOK \triangleq x \in RandomSubset(4711, SUBSET (1 .. 500))$

$H \triangleq \dots$

$Inv \triangleq TypeOK \wedge H$

$CheckInductiveSpec \triangleq Inv \wedge \square[\dots] \dots$

# Symmetry Reduction

- ▶ Chooses a representative of equivalence classes (orbit) of states
- ▶ Constructive Orbit Problem - in general - is NP-hard [see ?]

MODULE *Symmetry*

EXTENDS *FiniteSets*

VARIABLE  $x$

CONSTANT  $S$

ASSUME ( $\text{Cardinality}(S) \geq 9$ )

$\text{Spec} \triangleq (x \in S) \wedge \square [x' \in S]_{\langle x \rangle}$  Without symmetry: 9 states, without: 1

- ▶ For each state enumerate  $|vars| * |A|! * |B|!$  where  $A$  and  $B$  are two symmetry sets
- ▶ Not supported by liveness checking (TLC prints warning)

## Liveness under Symmetry

- ▶ TLA<sup>+</sup> actions (labeled arcs) hard to account for in quotient graph
  - ▶ Approach resulted in incompleteness of liveness checking
- ▶ ⇒ Abandoned idea
- ▶ Maybe: Use quotient graph to find SCCs, re-generate actual SCC for all elements of symmetry set
  - ▶ Inefficient if SCCs are large (which they tend to be)

## Partial Order Reduction for TLA<sup>+</sup>?

- ▶ (Static) POR - similar to SPIN's implementation - explored by S. Merz
  - ▶ ⇒ Didn't work too well
  - ▶ SPIN fine-grained atomicity similar to programming language
  - ▶ TLA<sup>+</sup> due to abstractions coarse-grained atomicity
    - ▶ Not looked at PlusCal (fine-grained atomicity)
- ▶ Dynamic POR might be different (open question)

## Conclusion & Outlook

- ▶ Continue to focus on scalability of parallel and distributed TLC
  - ▶ Concurrent SCC search with lock-free union find
  - ▶ Scalability StateQueue
- ▶ “TLA+ compiler” to speed-up evaluation of next-state relation
- ▶ Shift Toolbox maintenance to community
  
- ▶ Machine Learning combined with Cloud TLC
  - ▶ Optimize scalability and performance => Less manual tuning of TLC
  - ▶ Predict size of state graph/time to check => User defines “when”
  
- ▶ Start new with TLC-Next instead of continue with existing TLC
  - ▶ Feature cost and technical debt to drag on
- ▶ OTS data-structures not (yet?) ready for multicore “revolution”
- ▶ Scalability & Performance too much of an art

Q&A

Q&A

# Bibliography I