## Parallel Bug-finding in **Concurrent Programs via Reduced Interleaving Instances**\*

#### Truc L. Nguyen<sup>1</sup>, **Peter Schrammel**<sup>2</sup>, Salvatore La Torre<sup>4</sup>, Gennaro Parlato<sup>1</sup> Bernd Fischer<sup>3</sup>

<sup>1</sup> University of Southampton, United Kingdom

- <sup>2</sup> University of Sussex, United Kingdom
- <sup>3</sup> Stellenbosch University, South Africa
- <sup>4</sup> Universita degli Studi di Salerno, Italy



\* published at ASE'17





DI SALERNO



UNIVERSITÀ DEGLI STUDI UNIVERSITEIT•STELLENBOSCH•UNIVER jou kennisvennoot · your knowledge partner

## Concurrency makes bug finding harder.

State space explosion (i.e., large number of interleavings):



Problem: modern hardware means concurrency is everywhere  $\Rightarrow$  software is increasingly concurrent

## **Concurrency makes bug finding** *easier***.**



Concurrent hardware allows us to **run many (smaller)** analysis tasks in parallel:



## **Concurrency makes bug finding** *easier***.**



Concurrent hardware allows us to **run many (smaller)** analysis tasks in parallel:



How can we partition a task into independent smaller tasks?

## Strategy competition vs. task competition US

#### Strategy competition: run different settings on same task

(first counterexample "wins" and aborts other tasks)

#### **Swarm Verification Techniques**

Gerard J. Holzmann, Rajeev Joshi, and Alex Groce

Abs the l bene tin p (anticipate the appearance of systems with large numbers of CPU cores, but without matching increases in clock speeds... describe a model checking strategy that leverages this trend"

systems. For the near-term future, we can enterpate the appearance of systems with large numbers of CPU cores, but without matching increases in clockspeeds. We will describe a model checking strategy that can allow us to leverage this trend, and that allows us to tackle significantly larger problem sizes than before.

Index Terms—software engineering tools and techniques, logic model checking, distributed algorithms, software verification.

#### This talk:

Task competition: run same prover (setting) on different tasks  $\Rightarrow$  How can we partition a task into independent smaller tasks?

#### **Compositional Safety Refutation Techniques**

Kumar Madhukar<sup>1,3</sup>, Peter Schrammel<sup>2</sup>, and Mandayam Srivas<sup>3</sup>

<sup>1</sup> TCS Research, Pune, India

<sup>2</sup> University of Sussex, School of Engineering and Informatics, Brighton, UK <sup>3</sup> Chennai Mathematical Institute, Chennai, India

**Abstract.** One of the most successful techniques for refuting safety properties is to find counterexamples by bounded model checking. However, for large programs, bounded model checking instances often exceed the limits of resources available. Generating such counterexamples in a modular way could speed up refutation, but it is challenging because of the inherently non-compositional nature of these counterexamples. We start from the morelithic sofety worifection problem and present a start







#### **Compositional Safety Refutation Techniques**

Kumar Madhukar<sup>1,3</sup>, Peter Schrammel<sup>2</sup>, and Mandayam Srivas<sup>3</sup>

<sup>1</sup> TCS Research, Pune, India
 <sup>2</sup> University of Sussex, School of Engineering and Informatics, Brighton, UK
 <sup>3</sup> Chennai Mathematical Institute, Chennai, India



Horizontal/hierarchical decomposition, i.e. component-wise:

- E.g. Rely-guarantee, etc
- Easy for verification, difficult for refutation





#### **Compositional Safety Refutation Techniques**

Kumar Madhukar<sup>1,3</sup>, Peter Schrammel<sup>2</sup>, and Mandayam Srivas<sup>3</sup>

<sup>1</sup> TCS Research, Pune, India
 <sup>2</sup> University of Sussex, School of Engineering and Informatics, Brighton, UK
 <sup>3</sup> Chennai Mathematical Institute, Chennai, India



Horizontal/hierarchical decomposition, i.e. component-wise:

- E.g. Rely-guarantee, etc
- Easy for verification, difficult for refutation





#### **Compositional Safety Refutation Techniques**

Kumar Madhukar<sup>1,3</sup>, Peter Schrammel<sup>2</sup>, and Mandayam Srivas<sup>3</sup>

<sup>1</sup> TCS Research, Pune, India
 <sup>2</sup> University of Sussex, School of Engineering and Informatics, Brighton, UK
 <sup>3</sup> Chennai Mathematical Institute, Chennai, India



Horizontal/hierarchical decomposition, i.e. component-wise:

- E.g. Rely-guarantee, etc
- Easy for verification, difficult for refutation

Vertical decomposition, i.e. partition of execution traces:

- E.g. testing, path-wise symbolic execution, etc
- Easy for refutation, difficult for verification





#### **Compositional Safety Refutation Techniques**

Kumar Madhukar<sup>1,3</sup>, Peter Schrammel<sup>2</sup>, and Mandayam Srivas<sup>3</sup>

<sup>1</sup> TCS Research, Pune, India
 <sup>2</sup> University of Sussex, School of Engineering and Informatics, Brighton, UK
 <sup>3</sup> Chennai Mathematical Institute, Chennai, India



Horizontal/hierarchical decomposition, i.e. component-wise:

- E.g. Rely-guarantee, etc
- Easy for verification, difficult for refutation

Vertical decomposition, i.e. partition of execution traces:

- E.g. testing, path-wise symbolic execution, etc
- Easy for refutation, difficult for verification



#### **Reduced interleaving instances**

UNIVERSITY OF SUSSEX

#### Our goal:

**Split** set of **interleavings**  $I_k(P)$  into **subsets** that can be **analyzed symbolically** and **independently**.



#### **Our solution:**

**Derive** program variants  $P_{\vartheta}$  that allow context switches only in subsets of statements (tiles) s.t.  $I_k(P) = \bigcup_{\vartheta} I_k(P_{\vartheta})$ .





#### **Assumption: bounded concurrent programs**



- finite #threads, fixed (but arbitrary) schedule
  - captures all bounded round-robin computations for given bound
- bugs manifest within very few rounds [Musuvathi, Qadeer, PLDI'07]



#### Assumption: bounded concurrent programs



- finite #stmts
- control can only go forward
  - simplifies analysis and tiling



#### Tiles:



- tile: (contiguous) subset of visible statements
  - other tile types possible: random subsets, data-flow driven, ...
- **tiling**: partition of program into tiles
- uniform window tiling: all tiles have same size
  - number of visible statements



#### Tile selection:



- z-selection: subset of z tiles for each thread
  - context switches are only allowed from selected tiles
  - ⇒ context switches can only go into other selected tiles (or first thread statement)
- each z-selection specifies a reduced interleaving instance



#### **Completeness of selections:**



 each interleaving with *k* context switches can be covered by a [*k*/2]-selection ∂ ∈ Θ<sub>P</sub>

– each thread can only switch out at most  $\lfloor k/2 \rfloor$  times

⇒ set of all [*k*/2]-selections together covers all interleavings with *k* context switches:  $I_k(P) = \bigcup_{\vartheta} I_k(P_{\vartheta})$ 



#### **Completeness of selections:**



number of selections grows exponentially
 ⇒ sampling



# VERISMART (Verification Smart)

http://users.ecs.soton.ac.uk/gp4/cseq/

## VERISMART implements swarm verification US by task competition for multi-threaded C

Target:

- C programs with "rare" concurrency bugs, i.e.,
  - "large" number of interleavings
  - "few" interleavings lead to a bug
- automatic bug-finding (bounded analysis, not complete)
- reachability
  - assertion failure
  - out-of-bound array, division-by-zero, ...

Approach:

- source-to-source translation to generate instances (for tiling)
  - instances are bounded concurrent programs
- use cluster to run Lazy-CSeq over instances [Inverso et al., CAV'14]

### VERISMART architecture



- Inline/unwind module:
  - concurrent program  $\rightarrow$  bounded concurrent program
- Numerical labels module:
  - inject numerical labels at each visible statement
- Instrument module:
  - instrument the code with guarded commands (yield) that can enable/disable context switch points at numerical labels
- Split module:
  - generate variants with configuration from tiling and #tiles
  - randomize number of generated variants when #variants is large

#### Why does this work?

UNIVERSITY OF SUSSEX

Remember:

Each  $P_{\vartheta}$  allows only a (small) subset of P's interleavings

We assume bugs are rare,

- so for most  $\vartheta$ ,  $P_{\vartheta}$  does not exhibit the bug...
- ... and the analysis will run out of time
- but if  $P_{\vartheta}$  does exhibit the bug...
- ... the analysis will find it quick(er)

Hence,

- overall CPU time consumption goes (way) up...
- ... but with enough cores CPU time is free and...
- mean wall clock time to find failure goes down



# **Experimental Evaluation** on lock-free data structures

## eliminationstack



- C implementation of Hendler et al.'s EliminationStack (lock-free data structure) [Hendler, Shavit, Yerushalmi, SPAA'04]
- analyzed under SC
- annotated with several assertions to check linearizability [Bouajjani, Emmi, Enea, Hamza, POPL'15]
- ABA problem: requires 7 threads for exposure
- Lazy-CSeq can find bug in ~13h and 4GB
  - #unwind=1, #rounds=2, #threads=8, size=52 visible stmts
- all other tools fail (afaik)

#### safestack



- small implementation of lock-free stack (<100 loc)
- bug requires context bound of 5

$\leftarrow$ $\rightarrow$ C $\bullet$ https://social.msdn.mic	r <b>osoft.com</b> /Forums/en-US/	91c1971c-519f-	4ad2-816d-149e6b2fd916/bi	ug-with-a-co	ntext-swit 🛠 🤇	) 🖪 🚺
Microsoft   Developer Network				Sign in	MSDN subscriptions	Get tools
Downloads 🗸 Programs	~ Community ~	Documenta	ation 🗸			٩
Ask a question Quick access 🔻	Search related threads	Se	arch forum questions			Q
Asked by:	General discussion Hi guys,	icy Heisenbugs				
Dmitry Vyukov Joined Sep 2007 Dmitry Vyukov' 2 Show activity Top related threads Context and Context bound objects	This may be of interest to you. Here and in research papers you postulate that most bugs can be found with just a few context switches. Here is a nice counter-example with the most large context bound I ever saw. The code is very simple, it's no more than a try to avoid ABA in a trivial lock-free stack (the code was actually posted in a discussion forum by a guy who seems was read to put it into production). The bug can be detected with at least 3 threads and a context switch bound of 5 (!) Did you ever see such a bug? Can you spot it? ;) I think it's of interest at least as a reference example. AFAIR, the most context bound you reporte in papers is 3.					an be trivial is ready <mark>itch</mark> eported
Context switch	Here is original	source:				

#### safestack



- small implementation of lock-free stack (<100 loc)
- bug requires context bound of 5
- analyzed for SC, PSO, TSO
- Lazy-CSeq can find bug in ~7h and 6.5GB
  - #unwind=3, #rounds=4, #threads=4, size=152 visible stmts
- all other tools fail (afaik)

### eliminationstack: Results



- Lazy-CSeq: 46764 sec, 4.2 GB
- CBMC (sequential): 80.8 sec, 0.7 GB
  - average over 3000 interleavings, bug not found



## eliminationstack: Expected bug finding time



Number of cores

## safestack (SC): Results



- Lazy-CSeq: 24139 sec, 6.6 GB
- CBMC (sequential): 55.4 sec, 0.7 GB

- average over 3000 interleavings, bug not found



 $\Rightarrow$  similar picture, but less advantage for VERISMART

# safestack (SC): Expected bug finding time







# safestack (PSO): Expected bug finding time



OF SUSSEX



#### Conclusions



- first task-competitive swarm verification approach
- exploits availability of many cores to reduce mean wall clock time to find failure
  - allows us to handle very hard problems
  - high speed-ups already for 5-50 cores
- reduced interleaving instances boost bug-finding capabilities

#### **Future Work**

- other backends (testing)
- other tiling styles
- fast over-approximations to filter out safe instances