

# Tree-Regular Analysis of Parallel Programs with Dynamic Thread Creation and Locks

Benedikt Nordhoff

Fachbereich Mathematik und Informatik

Arbeitsgruppe Softwareentwicklung und Verifikation

3. November 2012



- ▶ Static reachability analysis of recursive parallel programs.
  - ▶ Utilizing DPNs (this talk)
  - ▶ Implementation for real programming languages (Java)
  - ▶ Applications to information flow control
- ▶ Static data flow analyses for sequential programs with applications to information flow control.
  - ▶ Utilizing/combining PDGs, path conditions and abstract interpretation
- ▶ ...



- ▶ A DPN is a bunch of push down systems (PDS) which can dynamically *spawn* new PDS as a side effect of their transitions.
- ▶ The PDS in a Monitor-DPN may additionally synchronize via a finite set of reentrant locks which are bound to the stack.
- ▶ Allows to precisely model effects of:
  - ▶ Recursive procedures.
  - ▶ Dynamic thread creation.
  - ▶ Synchronization via a finite set of well nested locks.
- ▶ Allows for a finite abstraction of:
  - ▶ Method local state.
  - ▶ Thread local state.
- ▶ Abstracts from shared state.




- ▶ Each process in a Monitor-DPN has a control state and a stack which contains stack symbols and possibly locks.
- ▶ There are five kinds of transitions each depends on the control state of the process and the topmost stack symbol:
  - Base** Modifies the control state and top of stack.
  - Return** Modifies the control state and removes the top of stack with a possibly underlying lock.
  - Call** Modifies the control state, top of stack, and adds an additional stack symbol.
  - Spawn** Like base but create a new process with a given control state and stack (without locks).
  - Use** Like call but puts a lock under the new stack symbol. Can only be executed if the lock is not currently on the stack of any other process.




- ▶ Executions/traces interleave actions from different threads.
- ▶ Action Trees branch executions at spawns, this yields a tree with context free paths.
- ▶ Execution Trees additionally branch at procedure calls.
- ▶ The set of reachable execution trees is tree regular.


## Trace of two processes *a* and *b*




## Trace of two processes $a$ and $b$




## Action tree




## Trace of two proceses $a$ and $b$



## Action tree




## Execution tree








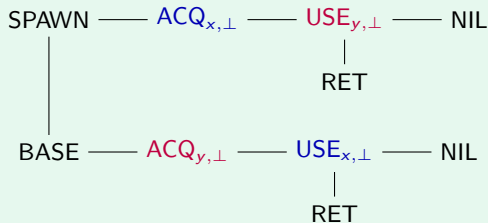
- ▶ Use tree-regularity to decide reachability of configurations with tree regular properties.
  1. Build a tree automaton accepting all reachable configurations.
  2. Build a tree automaton accepting configurations with property of interest.
  3. Check intersection for emptiness.
- ▶ Can also check for reachability from those reachable configurations.
  - ▶ Allows to check for arbitrary gen/kill properties e.g. def/use dependencies between two threads over a shared variable.









two final acquisitions of  $y$  in different threads







two final acquisitions of  $y$  in different threads



$y$  needs to be used after  
 $x$  has been finally acquired




two final acquisitions of  $y$  in different threads




$y$  needs to be used after  $x$  has been finally acquired

$x$  needs to be used after  $y$  has been finally acquired



two final acquisitions of  $y$  in different threads



$y$  needs to be used after  
 $x$  has been finally acquired

$x$  needs to be used after  
 $y$  has been finally acquired

- ▶ These properties are necessary, sufficient and tree-regular.

Lock sensitive schedulable execution trees.

Using a generalized version of Kahlon and Gupta's acquisition histories.

- ▶ State space:  $\{(A, U, G) \mid A \subseteq U \subseteq X, G \subseteq X \times X\}$
- ▶ Accepting states:  $\{(A, U, G) \in Q \mid G \text{ is acyclic}\}$

Interpretation: All operations only non reentrant.

- $A$  Locks finally acquired within the tree.
- $U$  Locks used or finally acquired within the tree.
- $G$  Acquisition graph,  $x \rightarrow x' \in G \Leftrightarrow x'$  is used or finally acquired *after*  $x$  has been finally acquired. (Order constrain)



Transitions:

$$\text{NIL} \rightarrow (\emptyset, \emptyset, \emptyset)$$

$$\text{RET} \rightarrow (\emptyset, \emptyset, \emptyset)$$

$$\text{BASE } \alpha \rightarrow \alpha$$

$$\text{NCALL } \alpha \rightarrow \alpha$$

$$\text{ACQ}_{x,\top} \alpha \rightarrow \alpha$$

$$f(A, U, G) (A', U', G') \rightarrow (A \cup A', U \cup U', G \cup G') \quad A \cap A' = \emptyset \\ f \in \{\text{RCALL}, \text{USE}_{x,\top}, \text{SPAWN}\}$$

$$\text{USE}_{x,\perp} (A, U, G) (A', U', G') \rightarrow (A \cup A', U \cup U' \cup \{x\}, G \cup G') \quad A \cap A' = \emptyset$$

$$\text{ACQ}_{x,\perp} (A, U, G) \rightarrow (A \cup \{x\}, U \cup \{x\}, G \cup \{(x, u) \mid u \in U\}) \quad x \notin U$$

The product automaton of these three automata accepts all lock sensitive execution trees of the DPN.

Let  $R, W$  be two sets of stack symbols. E.g. reads and writes of some variable.

- ▶ State space:  $2^{\{r,w\}}$
- ▶ Accepting states:  $\{\{r, w\}\}$

$$\text{NIL}_\gamma \rightarrow \{r\} \quad \gamma \in R \qquad \text{NIL}_\gamma \rightarrow \{w\} \quad \gamma \in W$$

$$\text{NIL}_\gamma \rightarrow \emptyset \quad \gamma \notin (R \cup W) \quad \text{RET} \rightarrow \emptyset$$

$$\{\text{ACQ, BASE, NCALL}\} \alpha \rightarrow \alpha$$

$$\{\text{RCALL, SPAWN, USE}\} \alpha \beta \rightarrow \alpha \cup \beta$$

- ▶ This tree automaton accepts all trees in which both sets are reached simultaneously. E.g. there exists a datarace.
- ▶  $\mathcal{L}(\mathcal{T}_M) \cap \mathcal{L}(\mathcal{T}_{ah}) \cap \mathcal{L}(\mathcal{T}_{CFL}^{R,W}) = \emptyset$  iff there exists no conflict.



- ▶ Have characterized (in some sense)  $post_{\mathcal{M}}^*({\text{NIL}}_{\gamma_0}^{p_0})$
- ▶ For tree-regular  $A \subseteq post_{\mathcal{M}}^*({\text{NIL}}_{\gamma_0}^{p_0})$  can characterize  $post_{\hat{\mathcal{M}}}^*(A)$ 
  - ▶ A tree transducer *marks* an intermediate configuration from  $A$  in the execution trees.
  - ▶ Release structures ensure the locks held at the intermediate configuration can be released before they are needed.
  - ▶  $\hat{\mathcal{M}}$  can be a restriction of the DPN  $\mathcal{M}$ .

- ▶ Have characterized (in some sense)  $post_{\mathcal{M}}^*({NIL}_{\gamma_0}^{p_0})$
- ▶ For tree-regular  $A \subseteq post_{\mathcal{M}}^*({NIL}_{\gamma_0}^{p_0})$  can characterize  $post_{\hat{\mathcal{M}}}^*(A)$ 
  - ▶ A tree transducer *marks* an intermediate configuration from  $A$  in the execution trees.
  - ▶ Release structures ensure the locks held at the intermediate configuration can be released before they are needed.
  - ▶  $\hat{\mathcal{M}}$  can be a restriction of the DPN  $\mathcal{M}$ .
- ▶ To calculate def/use chains over a shared variable let
  - $W$  be the set of execution trees which reached a configuration writing the variable at a given point (possibly in some context).
  - $\hat{\mathcal{M}}$  the DPN without killing transitions on the variable.
  - $R$  the set of execution trees which reached a configuration reading the variable at a given point (possibly in some context).

Check:  $post_{\hat{\mathcal{M}}}^*(post_{\mathcal{M}}^*({NIL}_{\gamma_0}^{p_0}) \cap W) \cap R \stackrel{?}{=} \emptyset$




- ▶ Use unary CUT-nodes to mark an arbitrary intermediate configuration in execution trees.
- ▶ Can use a tree transducer to obtain the intermediate configuration.
- ▶ A tree transducer is a tree automata with output.
- ▶ The inverse image of regular sets under tree transducers is regular.
  - ▶ Obtains trees where the marked intermediate configuration is accepted by a given tree automata.
- ▶ Check for lock sensitive schedulability using acquisition and release structures.



- ▶ A tree-automata checks that the cut-nodes actually mark an intermediate configuration by ensuring that:
  - ▶ Every process spawned before the intermediate configuration contains exactly one cut-node.
  - ▶ No process spawned after the intermediate configuration contains a cut-node.




- ▶ A tree-automata checks that the cut-nodes actually mark an intermediate configuration by ensuring that:
  - ▶ Every process spawned before the intermediate configuration contains exactly one cut-node.
  - ▶ No process spawned after the intermediate configuration contains a cut-node.
- ▶ Still need to check for lock-sensitive schedulability.




*x* is held throughout the second phase by the spawned thread but needed by the first.







$x$  is held throughout the second phase by the spawned thread but needed by the first.



Deadlock each thread needs the a lock held by the other thread to release the lock it holds.



$x$  is held throughout the second phase by the spawned thread but needed by the first.



Deadlock each thread needs the a lock held by the other thread to release the lock it holds.

- ▶ Combined with the previous acquisition histories these properties are necessary, sufficient and tree-regular.

Ensures no lock finally acquired at the cut is used afterwards.

State space:  $\mathbb{B} \times 2^X \times 2^X$

Accepting states:  $\{(C, A, U) \in \mathbb{B} \times 2^X \times 2^X \mid A \cap U = \emptyset\}$

	RCALL $(\perp, A, U) (\perp, A', U') \rightarrow (\perp, A \cup A', U \cup U')$
	RCALL $(\top, A, U) (\perp, \_, U') \rightarrow (\top, A, U \cup U')$
NIL $\rightarrow (\perp, \emptyset, \emptyset)$	RCALL $(\perp, A, \_) (\top, A', U') \rightarrow (\top, A \cup A', U')$
RET $\rightarrow (\perp, \emptyset, \emptyset)$	USE <sub>x,⊥</sub> $(\perp, A, U) (\perp, A', U') \rightarrow (\perp, A \cup A', U \cup U' \cup \{x\})$
BASE $p \rightarrow p$	USE <sub>x,⊥</sub> $(\top, A, U) (\perp, \_, U') \rightarrow (\top, A, U \cup U')$
NCALL $p \rightarrow p$	USE <sub>x,⊥</sub> $(\perp, A, \_) (\top, A', U') \rightarrow (\top, A \cup A', U')$
ACQ <sub>x,⊥</sub> $p \rightarrow p$	ACQ <sub>x,⊥</sub> $(\perp, A, U) \rightarrow (\perp, A, U \cup \{x\})$
CUT $(\perp, A, U) \rightarrow (\top, A, U)$	ACQ <sub>x,⊥</sub> $(\top, A, U) \rightarrow (\top, A \cup \{x\}, U)$
	SPAWN $(\_, A, U) (C, A', U') \rightarrow (C, A \cup A', U \cup U')$

(Reentrant use nodes are handled like returning calls.)

Ensures that all locks used at the cut can be released.

State space  $\mathbb{B} \times 2^X \times 2^{X \times X}$

Accepting states:  $\mathbb{B} \times 2^X \times \{G \in 2^{X \times X} \mid G \text{ is acyclic}\}$

NIL  $\rightarrow (\perp, \emptyset, \emptyset)$

RET  $\rightarrow (\perp, \emptyset, \emptyset)$

BASE  $p \rightarrow p$

NCALL  $p \rightarrow p$

ACQ<sub>x,-</sub>  $p \rightarrow p$

CUT  $(\perp, U, G) \rightarrow (\top, U, G)$

RCALL  $(C, U_c, G_c) (\perp, U_r, G_r) \rightarrow (C, U_c \cup U_r, G_c \cup G_r)$

RCALL  $(\perp, \_, G_c) (\top, U_r, G_r) \rightarrow (\top, U_r, G_c \cup G_r)$

USE<sub>x,⊤</sub>  $(C, U_c, G_c) (\perp, U_r, G_r) \rightarrow (C, U_c \cup U_r, G_c \cup G_r)$

USE<sub>x,-</sub>  $(\perp, \_, G_c) (\top, U_r, G_r) \rightarrow (\top, U_r, G_c \cup G_r)$

USE<sub>x,⊥</sub>  $(\perp, U_c, G_c) (\perp, U_r, G_r) \rightarrow (\perp, U_c \cup U_r \cup \{x\}, G_c \cup G_r)$

USE<sub>x,⊥</sub>  $(\top, U_c, G_c) (\perp, U_r, G_r) \rightarrow (\top, U_c \cup U_r, G_c \cup G_r \cup \{(u, x) \mid u \in U_c\})$

SPAWN  $(\perp, \_, \emptyset) (\perp, U_r, \emptyset) \rightarrow (\perp, U_r, \emptyset)$

SPAWN  $(\top, \_, G_s) (B, U_r, G_r) \rightarrow (B, U_r, G_s \cup G_r)$

- ▶ To check for reachability of  $B$  from  $A$ ,  
where  $A$  is lock-sensitively reachable.

$$(\mathcal{L}(\mathcal{T}_M) \cap \mathcal{L}(\mathcal{T}_{rs}) \cap \mathcal{L}(\mathcal{T}_{cwf}) \cap \mathcal{T}_{ct}^{-1}(A))|_{\text{CUT}} \cap \mathcal{L}(\mathcal{T}_{ah}) \cap B \stackrel{?}{=} \emptyset$$

- ▶ To check for reachability of  $B$  from  $A$ ,  
where  $A$  is lock-sensitively reachable.

$$(\mathcal{L}(\mathcal{T}_M) \cap \mathcal{L}(\mathcal{T}_{rs}) \cap \mathcal{L}(\mathcal{T}_{cwf}) \cap \mathcal{T}_{ct}^{-1}(A)) \mid_{\text{CUT}} \cap \mathcal{L}(\mathcal{T}_{ah}) \cap B \stackrel{?}{=} \emptyset$$

- ▶ Execution trees of the DPN  $\mathcal{M}$
- ▶ Locks held throughout and release structures
- ▶ Cut well-formed execution trees
- ▶ Inverse image of  $A$  under cut transducer
- ▶ Transducer *removing* cuts
- ▶ Acquisition histories



- ▶ To check for reachability of  $B$  from  $A$ ,  
where  $A$  is lock-sensitively reachable.

$$(\mathcal{L}(\mathcal{T}_M) \cap \mathcal{L}(\mathcal{T}_{rs}) \cap \mathcal{L}(\mathcal{T}_{cwf}) \cap \mathcal{T}_{ct}^{-1}(A)) \mid_{\text{CUT}} \cap \mathcal{L}(\mathcal{T}_{ah}) \cap B \stackrel{?}{=} \emptyset$$

- ▶ Execution trees of the DPN  $\mathcal{M}$
- ▶ Locks held throughout and release structures
- ▶ Cut well-formed execution trees
- ▶ Inverse image of  $A$  under cut transducer
- ▶ Transducer *removing* cuts
- ▶ Acquisition histories
- ▶ Before projecting out the cut-nodes, one can also check that some transitions don't occur after/before the cut to *switch* the DPN.



- ▶ Utilizing the T.J. Watson Libraries for Analysis (WALA) from IBM.
- ▶ Using XSB a tabulating Prolog system to evaluate tree automata.
- ▶ Some automata are extremely nondeterministic when evaluated top down and other when evaluated bottom up.
  - ▶ Emptiness check for product automata explores some automata bottom up and other top down.
  - ▶ Avoids exploration of unfeasible states.
- ▶ Witnesses can be generated.
- ▶ Integrated with *Joana* an information flow control tool for Java.
  - ▶ Currently checking individual def/use dependencies between threads for feasibility.
  - ▶ Want to check additionally:
    - ▶ Interferences in given contexts during the slicing phase.
    - ▶ Multiple interferences on critical paths.



Thanks for listening!



Ahmed Bouajjani, Markus Müller-Olm, and Tayssir Touili.  
Regular symbolic analysis of dynamic networks of pushdown systems.  
In *CONCUR 2005 – Concurrency Theory, 2005*.



Peter Lammich, Markus Müller-Olm, and Alexander Wenner.  
Predecessor sets of dynamic pushdown networks with tree-regular constraints.  
In *Computer Aided Verification, 2009*.



Thomas Gawlitza, Peter Lammich, Markus Müller-Olm, Helmut Seidl, and Alexander Wenner.  
Join-lock-sensitive forward reachability analysis for concurrent programs with dynamic process creation.  
In *Verification, Model Checking, and Abstract Interpretation, 2011*.



Vineet Kahlon, Franjo Ivančić, and Aarti Gupta.  
Reasoning about threads communicating via locks.  
In *In CAV. Springer, 2005*.



Benedikt Nordhoff, Peter Lammich, and Markus Müller-Olm.  
Iterable forward reachability analysis of Monitor-DPNs.  
Submitted for publication, 2012.