Algorithmic Logic-Based Verification

Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213

Arie Gurfinkel

Software Engineering Institute Carnegie Mellon University

© 2016 Carnegie Mellon University

Copyright 2015 Carnegie Mellon University

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Department of Defense.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

This material has been approved for public release and unlimited distribution.

This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

DM-0002754



Software Engineering Institute | Carnegie Mellon University





SYMBOLIC REACHABILITY



Software Engineering Institute

Carnegie Mellon University

Symbolic Reachability Problem

P = (*V*, *Init*, *Tr*, *Bad*)

P is UNSAFE if and only if there exists a number *N* s.t.

$$Init(X_0) \land \left(\bigwedge_{i=0}^{N-1} Tr(X_i, X_{i+1})\right) \land Bad(X_N) \not\Rightarrow \bot$$

P is SAFE if and only if there exists a safe inductive invariant Inv s.t.

$$Init \Rightarrow Inv
 Inv(X) \land Tr(X, X') \Rightarrow Inv(X')
 Inv \Rightarrow \neg Bad
 Safe$$



Software Engineering Institute Carnegie Mellon University

Constrained Horn Clauses (CHC)

A Constrained Horn Clause (CHC) is a FOL formula of the forms

$$\forall V . (\phi \land p_1[X_1] \land ... \land p_n[X_n] \rightarrow p_{n+1}[X])$$

$$\forall V . (\phi \land p_1[X_1] \land ... \land p_n[X_n] \rightarrow false)$$

where

- $\bullet\,\phi$ is a constrained in a background theory A
 - -of combined theory of Linear Arithmetic, Arrays, Bit-Vectors, ...
- p_1, \ldots, p_{n+1} are n-ary predicates
- p_i[X] is an application of a predicate to first-order terms

Spacer: Solving SMT-constrained CHC

Spacer: a solver for SMT-constrained Horn Clauses

- stand-alone implementation in a fork of Z3
- <u>http://bitbucket.org/spacer/code</u>
- Support for Non-Linear CHC
 - model procedure summaries in inter-procedural verification conditions
 - model assume-guarantee reasoning
 - uses MBP to under-approximate models for finite unfoldings of predicates
 - uses MAX-SAT to decide on an unfolding strategy

Supported SMT-Theories

- Best-effort support for arbitrary SMT-theories
 - data-structures, bit-vectors, non-linear arithmetic
- Full support for Linear arithmetic (rational and integer)
- Quantifier-free theory of arrays
 - only quantifier free models with limited applications of array equality





IC3, PDR, and Friends (1)

IC3: A SAT-based Hardware Model Checker

- Incremental Construction of Inductive Clauses for Indubitable Correctness
- A. Bradley: SAT-Based Model Checking without Unrolling. VMCAI 2011

PDR: Explained and extended the implementation

- Property Directed Reachability
- N. Eén, A. Mishchenko, R. K. Brayton: Efficient implementation of property directed reachability. FMCAD 2011

PDR with Predicate Abstraction (easy extension of IC3/PDR to SMT)

- A. Cimatti, A. Griggio, S. Mover, St. Tonetta: IC3 Modulo Theories via Implicit Predicate Abstraction. TACAS 2014
- J. Birgmeier, A. Bradley, G. Weissenbacher: Counterexample to Induction-Guided Abstraction-Refinement (CTIGAR). CAV 2014

IC3, PDR, and Friends (2)

GPDR: Non-Linear CHC with Arithmetic constraints

- Generalized Property Directed Reachability
- K. Hoder and N. Bjørner: Generalized Property Directed Reachability. SAT 2012

SPACER: Non-Linear CHC with Arithmetic

- fixes an incompleteness issue in GPDR and extends it with under-approximate summaries
- A. Komuravelli, A. Gurfinkel, S. Chaki: SMT-Based Model Checking for Recursive Programs. CAV 2014

PolyPDR: Convex models for Linear CHC

- simulating Numeric Abstract Interpretation with PDR
- N. Bjørner and A. Gurfinkel: Property Directed Polyhedral Abstraction. VMCAI 2015

ArrayPDR: CHC with constraints over Airthmetic + Arrays

- Required to model heap manipulating programs
- A. Komuravelli, N. Bjørner, A. Gurfinkel, K. L. McMillan:Compositional Verification of Procedural Programs using Horn Clauses over Integers and Arrays. FMCAD 2015



Spacer In Pictures







Software Engineering Institute Ca

Carnegie Mellon University

Logic-based Algorithmic Verification



Software Engineering Institute

Carnegie Mellon University



http://seahorn.github.io



Software Engineering Institute Ca

Carnegie Mellon University

SeaHorn Usage

Example: in test.c, check that x is always greater than or equal to y test.c



SeaHorn Encoding of Verification Conditions

Software Engineering Institute

Carnegie Mellon University

PARAMETRIZED SYMBOLIC REACHABILITY

joint work with Sharon Shoham

Software Engineering Institute

Carnegie Mellon University

What we want to do ...

```
local
    pc: \{CHOOSE, TRY, WAIT, MOVE\};
    curr, next, desired : Location
def proc(i):
    do
        pc[i] = CHOOSE : desired[i] \leftarrow *; pc[i] \leftarrow TRY;
        pc[i] = \text{TRY} \land \forall j.i < j \Rightarrow curr[j] \neq desired[i] \land next[j] \neq desired[i]
              next[i] \leftarrow desired[i]; pc[i] \leftarrow WAIT;
        pc[i] = WAIT \land \forall j . j < i \Rightarrow next[i] \neq curr[j] \land next[i] \neq next[j]:
              pc[i] \leftarrow \text{MOVE};
        pc[i] = MOVE:
               curr[i] \leftarrow next[i]; pc[i] \leftarrow CHOOSE;
def init(i, j):
    pc[i] = CHOOSE \land curr[i] = next[i] \land (i \neq j \Rightarrow curr[i] \neq curr[j])
def bad(i, j):
    i \neq j \land curr[i] = curr[j]
```

Software Engineering Institute | Carnegie Mellon University

Parameterized Symbolic Reachability Problem

$$T = (\mathbf{v}, Init(N, \mathbf{v}), Tr(i, N, \mathbf{v}, \mathbf{v}'), Bad(N, \mathbf{v}))$$

- v is a set of state variables
 - each $v_k \in \mathbf{v}$ is a map $Nat \rightarrow Rat$
 - v is partitioned into Local(v) and Global(v)
- Init(N, v) and Bad(N, v) are initial and bad states, respectively
- *Tr*(*i*, N, v, v') is a transition relation, parameterized by a process identifier *i* and total number of processes N

All formulas are over the combined theories of arrays and LRA

Init(N,v) and Bad(N,v) contain at most 2 quantifiers

- Init(N,v) = $\forall x,y . \phi_{Init}(N, x, y, v)$, where ϕ_{Init} is quantifier free (QF)
- Bad(N,v) = $\forall x,y . \phi_{Bad}(N,x, y, v)$, where ϕ_{Bad} is QF

Tr contains at most 1 quantifier

• $Tr(i, N, v, v') = \forall j . \rho (i, j, N, v, v')$

A State of a Parameterized System

Global			
v ₀	V ₁	V_2	V 3

Software Engineering Institute Carnegie Mellon University

Parameterized Symbolic Reachability

T = (**v**, *Init*, *Tr*, *Bad*)

T is UNSAFE if and only if there exists a number *K* s.t. $Init(\boldsymbol{v}_0) \land (\bigwedge_{s \in [0,K)} Tr(i_s, N, \boldsymbol{v}_s, \boldsymbol{v}_{s+1})) \land Bad(\boldsymbol{v}_K) \not\Rightarrow \bot$

T is SAFE if and only if there exists a safe inductive invariant Inv s.t.

Software Engineering Institute Carnegie Mellon University

Parameterized vs Non-Parameterized Reachability

 $Init(\boldsymbol{v}) \Rightarrow Inv(\boldsymbol{v})$ $Inv(\boldsymbol{v}) \land Tr(i, N, \boldsymbol{v}, \boldsymbol{v}') \Rightarrow Inv(\boldsymbol{v}')$ $Inv(\boldsymbol{v}) \Rightarrow \neg Bad(\boldsymbol{v})$

VC(T)

Init, Bad, and Tr might contain quantifiers

- e.g., "ALL processes start in unique locations"
- e.g., "only make a step if ALL other processes are ok"
- e.g., "EXIST two distinct process in a critical section"

Inv cannot be assumed to be quantifier free

• QF Inv is either non-parametric or trivial

Decide existence of quantified solution for CHC

- stratify search by the number of quantifiers
- solutions with 1 quantifier, 2 quantifiers, 3 quantifiers, etc...

tware Engineering Institute Carnegie Mellon University

ONE QUANTIFIER TWO QUANTIFIER

Software Engineering Institute

Carnegie Mellon University

One Quantifier (Solution)

$$Init(i, i, \mathbf{v}) \implies Inv_{1}(i, \mathbf{v})$$

$$Inv_{1}(i, \mathbf{v}) \wedge Tr(i, \mathbf{v}, \mathbf{v}') \implies Inv_{1}(i, \mathbf{v}')$$

$$j \neq i \wedge Inv_{1}(i, \mathbf{v}) \wedge Inv_{1}(j, \mathbf{v}) \wedge Tr(j, \mathbf{v}, \mathbf{v}') \implies Inv_{1}(i, \mathbf{v}')$$

$$Inv_{1}(i, \mathbf{v}) \wedge Inv_{1}(j, \mathbf{v}) \implies \neg Bad(i, j, \mathbf{v})$$

$$VC_{1}(T)$$

Claim

- If $VC_1(T)$ is QF-SAT then VC(T) is SAT
- If *Tr* does not contain functions that range over PIDs, then VC₁(T) is QF-SAT only if VC(T) admits a solution definable by a *simple* single quantifier formula
 - simple == quantified id variables do not appear as arguments to functions

 $VC_1(T)$ is essentially Owicki-Gries for 2 processes *i* and *j* If there are no global variables then (3) is unnecessary

• VC₁(T) is linear

How do we get it

- 1. Restrict Inv to a fixed number of quantifiers
 - e.g., replace Inv(N, v) with $\forall k.Inv_1(k, N, v)$

2. Case split consecution Horn clause based on the process that makes the move

- w+1 cases for w-quantifiers
 - one for each quantified id variable
 - one for interference by "other" process (only for global variables)
- 3. Instantiate the universal quantifier in \forall k.lnv₁(k, N, v)
 - use symmetry to reduce the space of instantiations
- 4. Other instantiations might be needed for quantifiers if
 - id variables appear as arguments to functions

How do we get it

$$Inv(\boldsymbol{v}) \wedge Tr(j, \boldsymbol{v}, \boldsymbol{v}') \implies Inv(\boldsymbol{v}')$$

$$(\forall k \cdot Inv_1(k, \boldsymbol{v})) \wedge Tr(j, \boldsymbol{v}, \boldsymbol{v}') \implies Inv_1(i, \boldsymbol{v}')$$

$$(\forall k \cdot Inv_1(k, \boldsymbol{v})) \wedge Tr(i, \boldsymbol{v}, \boldsymbol{v}') \implies Inv_1(i, \boldsymbol{v}')$$

$$(\forall k \cdot Inv_1(k, \boldsymbol{v})) \wedge j \neq i \wedge Tr(j, \boldsymbol{v}, \boldsymbol{v}') \implies Inv_1(i, \boldsymbol{v}')$$

$$Inv_1(i, \boldsymbol{v}) \wedge Tr(i, \boldsymbol{v}, \boldsymbol{v}') \implies Inv_1(i, \boldsymbol{v}')$$

$$Inv_1(i, \boldsymbol{v}) \wedge Inv_1(j, \boldsymbol{v}) \wedge j \neq i \wedge Tr(j, \boldsymbol{v}, \boldsymbol{v}') \implies Inv_1(i, \boldsymbol{v}')$$

Software Engineering Institute Carnegie Mellon University

Two Quantifier Solution

$$Init(i, j, \boldsymbol{v}) \wedge Init(j, i, \boldsymbol{v}) \wedge Init(i, i, \boldsymbol{v}) \wedge Init(j, j, \boldsymbol{v}) \Rightarrow I_2(i, j, \boldsymbol{v})$$
$$I_2(i, j, \boldsymbol{v}) \wedge Tr(i, \boldsymbol{v}, \boldsymbol{v}') \Rightarrow I_2(i, j, \boldsymbol{v}')$$
$$I_2(i, j, \boldsymbol{v}) \wedge Tr(j, \boldsymbol{v}, \boldsymbol{v}') \Rightarrow I_2(i, j, \boldsymbol{v}')$$
$$I_2(i, j, \boldsymbol{v}) \wedge I_2(j, z, \boldsymbol{v}) \wedge Tr(z, \boldsymbol{v}, \boldsymbol{v}') \wedge z \neq i \wedge z \neq j \Rightarrow I_2(i, j, \boldsymbol{v}')$$
$$I_2(i, j, \boldsymbol{v}) \Rightarrow \neg Bad(i, j, \boldsymbol{v})$$

Claim

- If $VC_2(T)$ is QF-SAT then VC(T) is SAT
- If *Tr* does not contain functions that range over PIDs, then VC₂(T) is QF-SAT only if VC(T) admits a solution definable by a *simple* two quantifier formula
- At least 2 quantifiers are "needed" for systems with global guards

Extends to K-quantifiers

Software Engineering Institute Carnegie Mellon University

Putting it all together

Solve for Inductive Invariant k := 1;while true do $Inv_k(i_1,\ldots,i_k,\boldsymbol{v}) := \text{Solve}(U^k(VC^{\omega}(T)));$ if $Inv_k(i_1,\ldots,i_k,\boldsymbol{v}) \neq null$ then return "inductive invariant found: $\forall i_1,\ldots,i_k$. $Inv(i_1,\ldots,i_k,\boldsymbol{v})$ " $res := ModelCheck(T_k)$; Look for bugs if res = cex then **return** "counterexample found for k processes" k := k + 1

Finite vs Infinite Number of Processes

$$\begin{array}{ll} \operatorname{def \ proc}(i):\\ \operatorname{do}\\ & b[i]=0 \ : \ b[i]:=1 \ ;\\ & b[i]=1 \ : \ b[i]:=0 \ ;\\ & (\forall j\neq i \ . \ b[j]\neq b[i]) \ : \ pc[i]:=E \ ;\\ \end{array}$$

$$\begin{array}{ll} \operatorname{def \ init}(i,j): \ pc[i]=I \land b[i]=0 \ ;\\ \operatorname{def \ bad}(i,j): \ pc[i]=E \ ; \end{array}$$

Tr does not depend on N (number of processes) Safe for infinitely many processes

$$Inv \equiv (\forall i . b[i] \in [0, 1] \land pc[i] = I) \land (\forall i, j, k . distinct(i, j, k) \Rightarrow \neg distinct(b[i], b[j], b[k]))$$

Cex for N = 2

Software Engineering Institute | Carnegie Mellon University

Evaluation and Implementation

Python-based Implementation

- Simple language for specifying concurrent protocols
- Local and Universally guarded transitions
- Constraints over arrays and integer arithmetic
- Reduce to CHC using the rules and solve using Spacer

Evaluated on Simple/Tricky Well-Know Protocols

- Dining philosophers, bakery1, bakery2, collision avoidance, TICKET
- Models are pretty close to an implementation
 - limit abstraction in modeling, try to make verification hard
- Safe inductive invariants computed within seconds

ftware Engineering Institute Carnegie Mellon University

Related Work

Kedar Namjoshi et al.

- Local Proofs for Global Safety Properties, and many other papers
- systematic derivation of proof rules for *concurrent* systems
- finite state and fixed number of processes

Andrey Rybalchenko et al.

- Compositional Verification of Multi-Threaded Programs, and others
- compositional proof rules for concurrent systems are CHC
- infinite state and fixed number of processes

Lenore Zuck et al.

- Invisible Invariants
- finite state and parametric number of processes
- finite model theorem for special classes of parametric systems

Nikolaj Bjørner, Kenneth L. McMillan, and Andrey Rybalchenko

• On Solving Universally Quantified Horn Clauses. SAS 2013:

Conclusion

Parameterized Verification == Quantified solutions for CHC

Quantifier instantiation to *systematically* derive proof rules for verification of safety properties of parameterized systems

• Parameterized systems definable with SMT-LIB syntax

Lazy vs Eager Quantifier Instantiation

- eager instantiation in this talk
- would be good to extend to lazy / dynamic / model-based instantiation

Connections with other work in parameterized verification

- complete instantiation = decidability ?
- relative completeness

ftware Engineering Institute Carnegie Mellon University

?

?

?

?

Software Engineering Institute

Carnegie Mellon University

Algorithmic Logic-Based Verification Gurfinkel, 2016 © 2016 Carnegie Mellon University ?

Contact Information

Arie Gurfinkel, Ph. D. Principle Researcher CSC/SSD Telephone: +1 412-268-5800 Email: info@sei.cmu.edu

Web

www.sei.cmu.edu/contact.cfm

U.S. Mail

Software Engineering Institute Customer Relations 4500 Fifth Avenue Pittsburgh, PA 15213-2612 USA

Customer Relations

Email: info@sei.cmu.edu Telephone: +1 412-268-5800 SEI Phone: +1 412-268-5800 SEI Fax: +1 412-268-6257

