Building Program Verifiers from Compilers and Theorem Provers

Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213

Arie Gurfinkel

based on joint work with Teme Kahsai, Jorge A. Navas, Anvesh Komuravelli, and Nikolaj Bjorner



Software Engineering Institute 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-0002433

Software Engineering Institute Carnegie Mellon University

Automated Software Analysis









Software Engineering Institute

Carnegie Mellon University



Turing, 1936: "undecidable"



Software Engineering Institute | Carnegie Mellon University

Turing, 1949

Alan M. Turing. "Checking a large routine", 1949

How can one check a routine in the sense of making sure that it is right?

programmer should make a number of definite assertions which can be checked individually, and from which the correctness of the whole programme easily follows.



Three-Layers of a Program Verifier

Compiler

- compiles surface syntax to some machine
- embodies syntax with semantics

Verification Condition Generator

- transforms a program and a property to verification condition in logic
- employs different abstractions, refinements, proof-search strategies, etc.

Automated Theorem Prover / Reasoning Engine

- discharges verification conditions
- general purpose constraint solver
- SAT, SMT, Abstract Interpreter, Temporal Logic Model Checker,...





tware Engineering Institute | Carnegie Mellon University



http://seahorn.github.io



Software Engineering Institute

Carnegie Mellon University

SeaHorn Philosophy

Build a state-of-the-art Software Model Checker

- useful to "average" users
 - user-friendly, efficient, trusted, certificate-producing, ...
- useful to researchers in verification
 - modular design, clean separation between syntax, semantics, and logic, ...

Stand on the shoulders of giants

- reuse techniques from compiler community to reduce verification effort
 - SSA, loop restructuring, induction variables, alias analysis, ...
 - static analysis and abstract interpretation
- reduce verification to logic
 - verification condition generation
 - Constrained Horn Clauses

Build reusable logic-based verification technology

• "SMT-LIB" for program verification



The Plan

This Lecture

- front-end: verification conditions, constrained horn clauses
- from verification problems to decision problems in logic

Next Lecture

- back-end: solving verification conditions
- IC3/PDR algorithms and extensions for software verification
- <u>http://arieg.bitbucket.org/pdf/gurfinkel_ssft15.pdf</u>

Labs

- we will play with tools
- get binary distribution from github.com/seahorn/seahorn
- get seahorn-tutorial repo from github.com/seahorn-tutorial
- pre-requisites: recent Linux or OSX + clang-3.6
- can use provided VM to satisfy the pre-requisites

SeaHorn Verification Framework



Distinguishing Features

- LLVM front-end(s)
- Constrained Horn Clauses to represent Verification Conditions
- Comparable to state-of-the-art tools at SV-COMP'15

Goals

- be a state-of-the-art Software Model Checker
- be a framework for experimenting and developing CHC-based verification

Software Engineering Institute Carnegie Mellon University

10

Related Tools

CPAChecker

- Custom front-end for C
- Abstract Interpretation-inspired verification engine
- Predicate abstraction, invariant generation, BMC, k-induction

SMACK / Corral

- LLVM-based front-end
- Reduces C verification to Boogie
- Corral / Q verification back-end based on Bounded Model Checking with SMT

UFO

- LLVM-based front-end (partially reused in SeaHorn)
- Combines Abstract Interpretation with Interpolation-Based Model Checking
- (no longer actively developed)

Software Engineering Institute Carnegie Mellon University

SeaHorn Usage

> sea pf FILE.c

Outputs sat for unsafe (has counterexample); unsat for safe Additional options

- --cex=trace.xml outputs a counter-example in SV-COMP'15 format
- --show-invars displays computed invariants
- --track={reg,ptr,mem} track registers, pointers, memory content
- --step={large,small} verification condition step-semantics
 - *small* == basic block, *large* == loop-free control flow block
- --inline inline all functions in the front-end passes

Additional commands

- sea smt generates CHC in extension of SMT-LIB2 format
- sea clp -- generates CHC in CLP format (under development)
- sea lfe-smt generates CHC in SMT-LIB2 format using legacy front-end

12

Verification Pipeline



Software Engineering Institute Carnegie Mellon University

Constrained Horn Clauses INTERMEDIATE REPRESENTATION



Software Engineering Institute

Carnegie Mellon University

Constrained Horn Clauses (CHC)

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

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

where

- A is a background theory (e.g., Linear Arithmetic, Arrays, Bit-Vectors, or combinations of the above)
- ϕ is a constrained in the background theory A
- p₁, ..., p_n, h are n-ary predicates
- p_i[X] is an application of a predicate to first-order terms

15

Example Horn Encoding



Software Engineering Institute | Carnegie Mellon University



Software Engineering Institute Carnegie Mellon University

CHC Satisfiability

A **model** of a set of clauses Π is an interpretation of each predicate p_i that makes all clauses in Π valid

A set of clauses is **satisfiable** if it has a model, and is unsatisfiable otherwise

A model is **A-definable**, it each p_i is definable by a formula ψ_i in A



Software Engineering Institute Ca

Carnegie Mellon University

Relationship between CHC and Verification

A program satisfies a property iff corresponding CHCs are satisfiable

satisfiability-preserving transformations == safety preserving

Models for CHC correspond to verification certificates

inductive invariants and procedure summaries

Unsatisfiability (or derivation of FALSE) corresponds to counterexample

• the resolution derivation (a path or a tree) is the counterexample

CAVEAT: In SeaHorn the terminology is reversed

tware Engineering Institute

- SAT means there exists a counterexample a BMC at some depth is SAT
- UNSAT means the program is safe BMC at all depths are UNSAT

Carnegie Mellon University Gurfi

19

FROM PROGRAMS TO **CLAUSES**



Software Engineering Institute Carnegie Mellon University

Hoare Triples

A Hoare triple {Pre} P {Post} is valid iff every terminating execution of P that starts in a state that satisfies *Pre* ends in a state that satisfies *Post*

Inductive Loop Invariant

Function Application

 $\begin{array}{ll} (\text{Pre} \land \text{p=a}) \Rightarrow \text{P} & \{\text{P}\} \ \text{Body}_{\text{F}}\{\text{Q}\} & (\text{Q} \land \text{p,r=a,b}) \Rightarrow \text{Post} \\ \\ \{\text{Pre}\} \ \text{b} = \text{F}(a) \ \{\text{Post}\} & \end{array}$

Recursion

 $\{Pre\} b = F(a) \{Post\} \vdash \{Pre\} Body_F \{Post\}$

 $\{Pre\} b = F(a) \{Post\}$



Weakest Liberal Pre-Condition

Validity of Hoare triples is reduced to FOL validity by applying a predicate transformer

Dijkstra's weakest liberal pre-condition calculus [Dijkstra'75]

wlp (P, Post)

weakest pre-condition ensuring that executing P ends in Post

{Pre} P {Post} is valid

 \Leftrightarrow Pre \Rightarrow wlp (P, Post)



Software Engineering Institute Carnegie Mellon University

A Simple Programming Language

```
Prog ::= def Main(x) { body<sub>M</sub> }, ..., def P (x) { body<sub>P</sub> }
```

```
body ::= stmt (; stmt)*
```

```
stmt ::= x = E | assert (E) | assume (E) |
          while E do S | y = P(E) |
          L:stmt | goto L
                                      (optional)
```

```
:= expression over program variables
Ε
```



Software Engineering Institute Carnegie Mellon University

Horn Clauses by Weakest Liberal Precondition

```
Prog ::= def Main(x) { body<sub>M</sub> }, ..., def P (x) { body<sub>P</sub> }
```

```
\begin{split} & \text{wlp} \; (\text{x=E, Q}) = \text{let x=E in Q} \\ & \text{wlp} \; (\text{assert}(E) \;, \; Q) = E \land Q \\ & \text{wlp} \; (\text{assume}(E), \; Q) = E \rightarrow Q \\ & \text{wlp} \; (\text{while E do } S, \; Q) = I(w) \land \\ & \forall w \; . \; ((I(w) \land E) \rightarrow \text{wlp} \; (S, \; I(w))) \land ((I(w) \land \neg E) \rightarrow Q)) \\ & \text{wlp} \; (y \; = \; P(E), \; Q) = p_{\text{pre}}(E) \land (\forall \; r. \; p(E, \; r) \rightarrow Q[r/y]) \end{split}
```

ToHorn (def P(x) {S}) = wlp (x0=x;assume($p_{pre}(x)$); S, p(x0, ret)) ToHorn (Prog) = wlp (Main(), true) $\land \forall \{P \in Prog\}$. ToHorn (P)

Software Engineering Institute Carnegie Mellon University

Example of a WLP Horn Encoding



C1:
$$I(x,y,x,y) \leftarrow y \ge 0$$
.
C2: $I(x+1,y-1,x_o,y_o) \leftarrow I(x,y,x_o,y_o), y \ge 0$.
C3: false $\leftarrow I(x,y,x_o,y_o), y \le 0, x \ne x_o + y_o$

 $\{y \ge 0\} P \{x = x_{old} + y_{old}\}$ is **true** iff the query C₃ is **satisfiable**

Software Engineering Institute Carnegie Mellon University

Dual WLP

Dual weakest liberal pre-condition

dual-wlp (P, Post) = \neg wlp (P, \neg Post)

s ⊨ **dual-wlp** (P, Post) iff there exists an execution of P that starts in s and ends in Post

dual-wlp (P, Post) is the weakest condition ensuring that an execution of P can reach a state in Post



Software Engineering Institute Carneg

Carnegie Mellon University

Control Flow Graph

A CFG is a graph of basic blocks

• edges represent different control flow

A CFG corresponds to a program syntax

• where statements are restricted to the form

L_i:S ; goto L_j

and S is control-free (i.e., assignments and procedure calls)





Software Engineering Institute | C

Carnegie Mellon University

Horn Clauses by Dual WLP

Assumptions

• each procedure is represent by a control flow graph

- i.e., statements of the form $I_i:S$; goto I_i , where S is loop-free
- program is unsafe iff the last statement of Main() is reachable
 - i.e., no explicit assertions. All assertions are top-level.

For each procedure P(x), create predicates

• I(w) for each label, $p_{en}(x_0,x,w)$ for entry, $p_{ex}(x_0,r)$ for exit

The verification condition is a conjunction of clauses:

 $\begin{array}{l} p_{en}(x_{0},x) \leftarrow x_{0}^{=}x \\ I_{i}(x_{0}^{},w') \leftarrow I_{j}(x_{0}^{},w) \wedge \neg wlp \; (S, \neg(w=w')), \, \text{for each statement } I_{i} : S; \, \text{goto } I_{j} \\ p \; (x_{0}^{},r) \leftarrow p_{ex}(x_{0}^{},r) \\ \text{false} \leftarrow \text{Main}_{ex}(x, \, \text{ret}) \end{array}$

Carnegie Mellon University

28

Example Horn Encoding



Software Engineering Institute | Carnegie Mellon University

From CFG to Cut Point Graph

A *Cut Point Graph* hides (summarizes) fragments of a control flow graph by (summary) edges

Vertices (called, *cut points*) correspond to *some* basic blocks

An edge between cut-points c and d summarizes all finite (loop-free) executions from c to d that do not pass through any other cut-points



oftware Engineering Institute | Carne

Cut Point Graph Example





Software Engineering Institute | Car

Carnegie Mellon University

From CFG to Cut Point Graph

A *Cut Point Graph* hides (summarizes) fragments of a control flow graph by (summary) edges

Cut Point Graph preserves reachability of (not-summarized) control location.

Summarizing loops is undecidable! (Halting program)

A *cutset summary* summarizes all location except for a *cycle cutset* of a CFG. Computing minimal cutset summary is NP-hard (minimal feedback vertex set).

A reasonable compromise is to summarize everything but heads of loops. (Polynomial-time computable).



32

Single Static Assignment

SSA == every value has a unique assignment (a *definition*) A procedure is in SSA form if every variable has exactly one definition

SSA form is used by many compilers

- explicit def-use chains
- simplifies optimizations and improves analyses

PHI-function are necessary to maintain unique definitions in branching control flow

 $x = PHI (v_0:bb_0, ..., v_n:bb_n))$

(phi-assignment)

"x gets v_i if previously executed block was bb_i"



ftware Engineering Institute | Carnegie Mellon University

Single Static Assignment: An Example

int x, y, n; x = 0;while (x < N) { if (y > 0)x = x + y;else X = X - Y;y = -1 * y;

Software Engineering Institute Carnegie Mellon University

Large Step Encoding

Problem: Generate a compact verification condition for a loop-free block of code

Software Engineering Institute

Carnegie Mellon University

Large Step Encoding: Extract all Actions

$$x_1 = x_0 + y_0$$

 $x_2 = x_0 - y_0$
 $y_1 = -1 * y_0$

Software Engineering Institute Carnegie Mellon University

Example: Encode Control Flow

 $X_1 = X_0 + Y_0$ $\mathbf{x}_2 = \mathbf{x}_0 - \mathbf{y}_0$ $y_1 = -1 * y_0$ $B_2 \rightarrow x_{\rho} < N$ $B_3 \rightarrow B_2 \wedge y_0 > 0$ $B_4 \rightarrow B_2 \wedge y_0 \leq 0$ $B_5 \rightarrow (B_3 \land X_3 = X_1) \lor$ $(B_4 \wedge X_3 = X_2)$ $B_5 \wedge x'_{\theta} = x_3 \wedge y'_{\theta} = y_1$ $p_1(x'_0, y'_0) \leftarrow p_1(x_0, y_0), \phi.$

Software Engineering Institute Carnegie Mellon University

Mixed Semantics PROGRAM TRANSFORMATION



Software Engineering Institute Ca

Carnegie Mellon University

Deeply nested assertions





Software Engineering Institute | Carne

Carnegie Mellon University

Deeply nested assertions



Counter-examples are long Hard to determine (from main) what is relevant

Software Engineering Institute Carnegie Mellon University

Mixed Semantics

[GWC'08,LQ'14]

Stack-free program semantics combining:

- operational (or small-step) semantics
 - i.e., usual execution semantics
- natural (or big-step) semantics: function summary [Sharir-Pnueli 81]
 - (σ , σ `) \in ||f|| iff the execution of f on input state σ terminates and results in state σ '
- some execution steps are big, some are small
- Non-deterministic executions of function calls
 - update top activation record using function summary, or
 - enter function body, forgetting history records (i.e., no return!)

Preserves reachability and non-termination

<u>Theorem:</u> Let K be the operational semantics, K^m the stack-free semantics, and L a program location. Then,

 $\mathsf{K}\vDash\mathsf{EF}(\mathsf{pc=L})\Leftrightarrow\mathsf{K}^{\mathsf{m}}\vDash\mathsf{EF}(\mathsf{pc=L})\quad\text{ and }\quad\mathsf{K}\vDash\mathsf{EG}(\mathsf{pc\neq L})\Leftrightarrow\mathsf{K}^{\mathsf{m}}\vDash\mathsf{EG}(\mathsf{pc\neq L})$





Software Engineering Institute

Carnegie Mellon University

Mixed Semantics as Program Transformation





Software Engineering Institute Carnegie Mellon University

Mixed Semantics: Summary

Every procedure is inlined at most once

- in the worst case, doubles the size of the program
- can be restricted to only inline functions that directly or indirectly call error() function

Easy to implement at compiler level

- create "failing" and "passing" versions of each function
- reduce "passing" functions to returning paths
- in main(), introduce new basic block bb.F for every failing function F(), and call failing.F in bb.F
- inline all failing calls
- replace every call to F to non-deterministic jump to bb.F or call to passing F
- Increases context-sensitivity of context-insensitive analyses
 - context of failing paths is explicit in main (because of inlining)
 - enables / improves many traditional analyses

SOLVING CHC WITH SMT



Software Engineering Institute Carnegie Mellon University

Programs, Cexs, Invariants

A program $P = (V, Init, \rho, Bad)$

• Notation: $\mathcal{F}(X) = \exists \boldsymbol{u} . (X \land \rho) \lor \text{Init}$

Software Engineering Institute

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

$$Init(v_0) \land \left(\bigwedge_{i=0}^{N-1} \rho(v_i, v_{i+1})\right) \land Bad(v_N) \not\Rightarrow \bot$$

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

$$Init(u) \Rightarrow Inv(u)
 Inv(u) \land \rho(u, v) \Rightarrow Inv(v)$$

$$Inv(u) \Rightarrow \neg Bad(u)$$
Inv(u) Safe



Carnegie Mellon University Gur

IC3/PDR Algorithm Overview

bounded safety

Input: Safety problem $\langle Init(X), Tr(X, X'), Bad(X, J) \rangle$

 $F_0 \leftarrow Init; N \leftarrow 0$ repeat

 $\mathbf{G} \leftarrow \text{PdrMkSafe}([F_0, \dots, F_N], Bad)$

if $\mathbf{G} = []$ then return *Reachable*; $\forall 0 \leq i \leq N \cdot F_i \leftarrow \mathbf{G}[i]$

$$F_0, \ldots, F_N \leftarrow \text{PdrPush}([F_0, \ldots, F_N])$$

if $\exists 0 \leq i < N \cdot F_i = F_{i+1}$ **then return** Unreg hable; $N \leftarrow N + 1$; $F_N \leftarrow \emptyset$ strengthen result

until ∞ ;

Software Engineering Institute Carnegie Mellon University

47

IC3/PDR in Pictures





Software Engineering Institute Ca

Carnegie Mellon University

48

PdrMkSafe





Software Engineering Institute Ca

Carnegie Mellon University

IC3/PDR in Pictures





Software Engineering Institute | Ca

Carnegie Mellon University

Building Verifiers from Comp and SMT Gurfinkel, 2015 © 2015 Carnegie Mellon University

PdrPush

IC3/PDR in Pictures





Software Engineering Institute | Carn

Carnegie Mellon University

Building Verifiers from Comp and SMT Gurfinkel, 2015 © 2015 Carnegie Mellon University

PdrPush

Spacer: Solving CHC in Z3

Spacer: 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





RESULTS



Software Engineering Institute | Carnegie Mellon University

SV-COMP 2015

4th Competition on Software Verification held (here!) at TACAS 2015 Goals

- Provide a snapshot of the state-of-the-art in software verification to the community.
- Increase the visibility and credits that tool developers receive.
- Establish a set of benchmarks for software verification in the community.

Participants:

 Over 22 participants, including most popular Software Model Checkers and Bounded Model Checkers

Benchmarks:

- C programs with error location (programs include pointers, structures, etc.)
- Over 6,000 files, each 2K 100K LOC
- Linux Device Drivers, Product Lines, Regressions/Tricky examples
- http://sv-comp.sosy-lab.org/2015/benchmarks.php

54

Results for DeviceDriver category



Software Engineering Institute Carnegie Mellon University

Detecting Buffer Overflow in Auto-pilot software

Show absence of Buffer Overflows in

paparazzi and mnav autopilots



Automatically instrument buffer accesses with runtime checks

Use SeaHorn to validate that run-time checks never fail

- somewhat slower than pure abstract interpretation
- much more precise!



Carnegie Mellon University

Conclusion

SeaHorn (http://seahorn.github.io)

- a state-of-the-art Software Model Checker
- LLVM-based front-end
- CHC-based verification engine
- a framework for research in logic-based verification



The future

- making SeaHorn useful to users of verification technology
 - counterexamples, build integration, property specification, proofs, etc.
- targeting many existing CHC engines
 - specialize encoding and transformations to specific engines
 - communicate results between engines
- richer properties
 - termination, liveness, synthesis



Software Engineering Institute Carnegie Mellon University

57

Contact Information

Arie Gurfinkel, Ph. D. Sr. 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



Software Engineering Institute | Carnegi

Carnegie Mellon University ^G