

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



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



Automated Software Analysis





Turing, 1936: “undecidable”

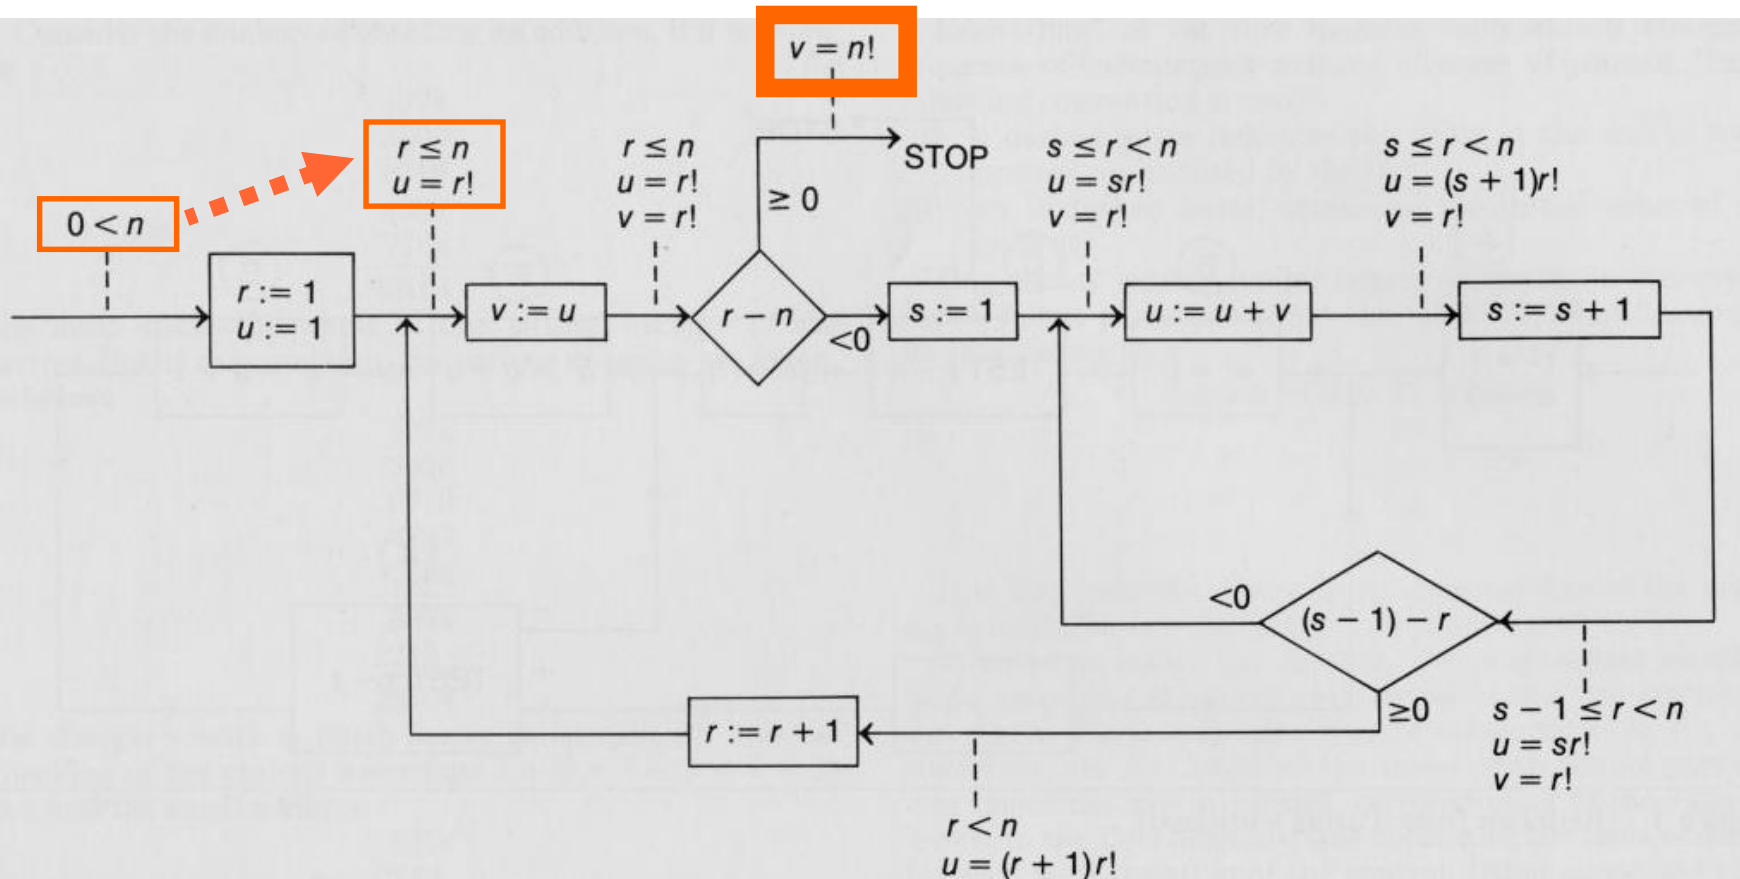


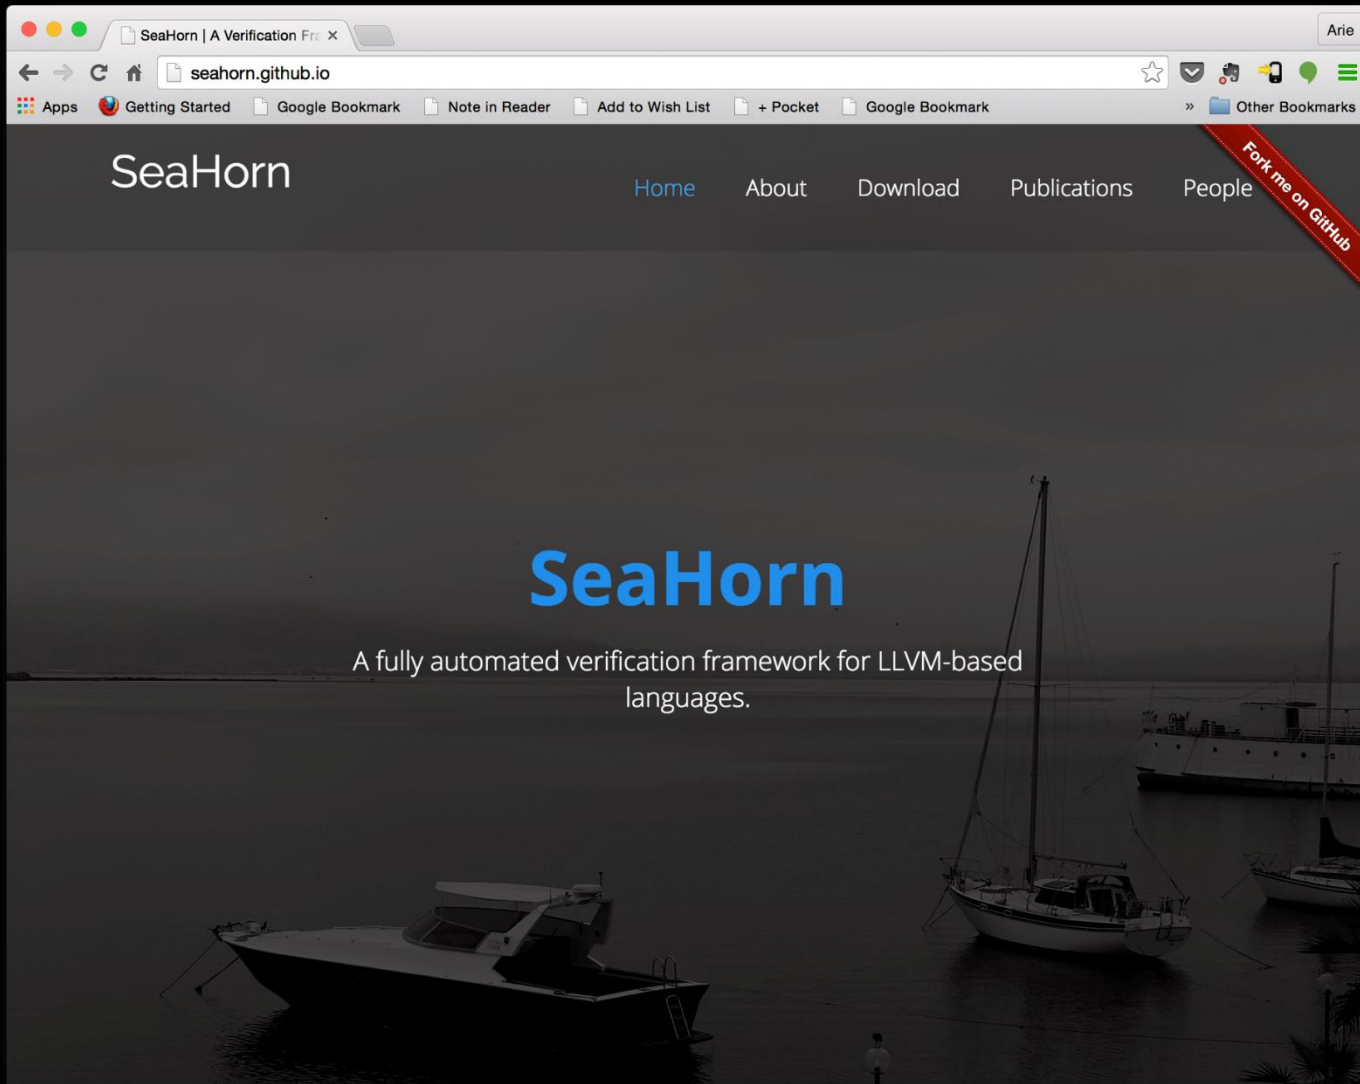
Turing, 1949

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.





<http://seahorn.github.io>



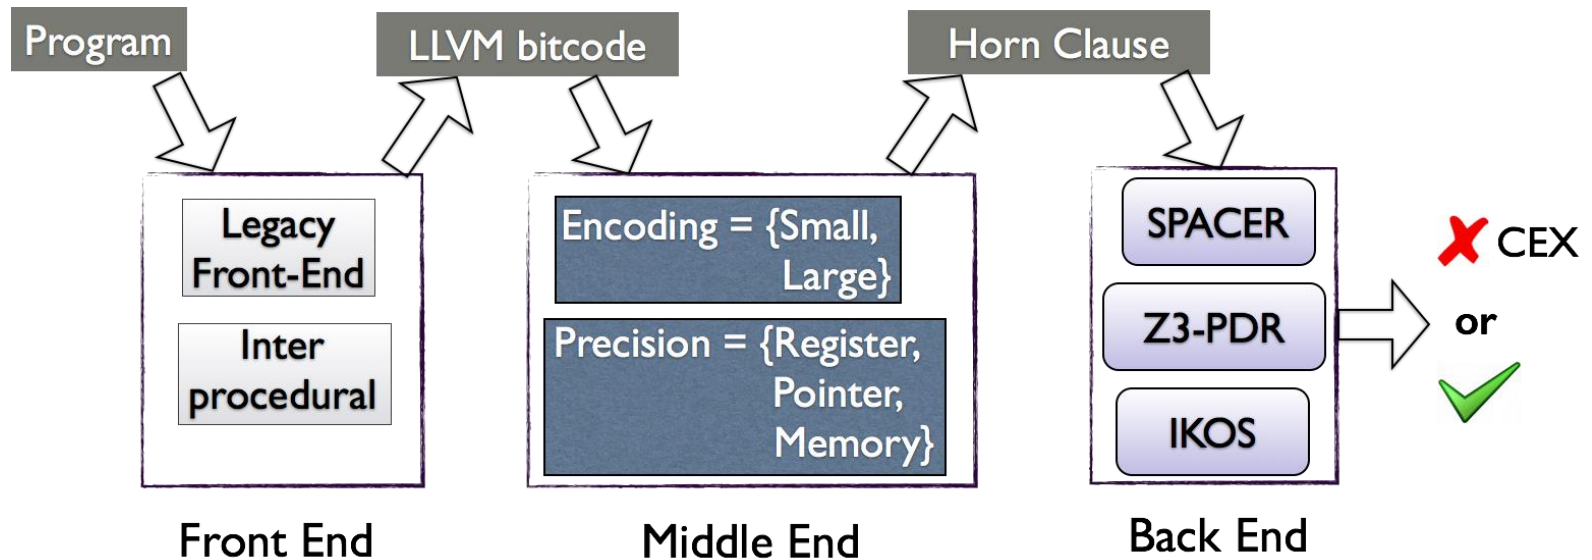
Software Engineering Institute

Carnegie Mellon University

Building Verifiers from Comp and SMT
Gurfinkel, 2015

© 2015 Carnegie Mellon University

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



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)



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



SeaHorn Usage

> sea pf FILE.c

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

Additional options

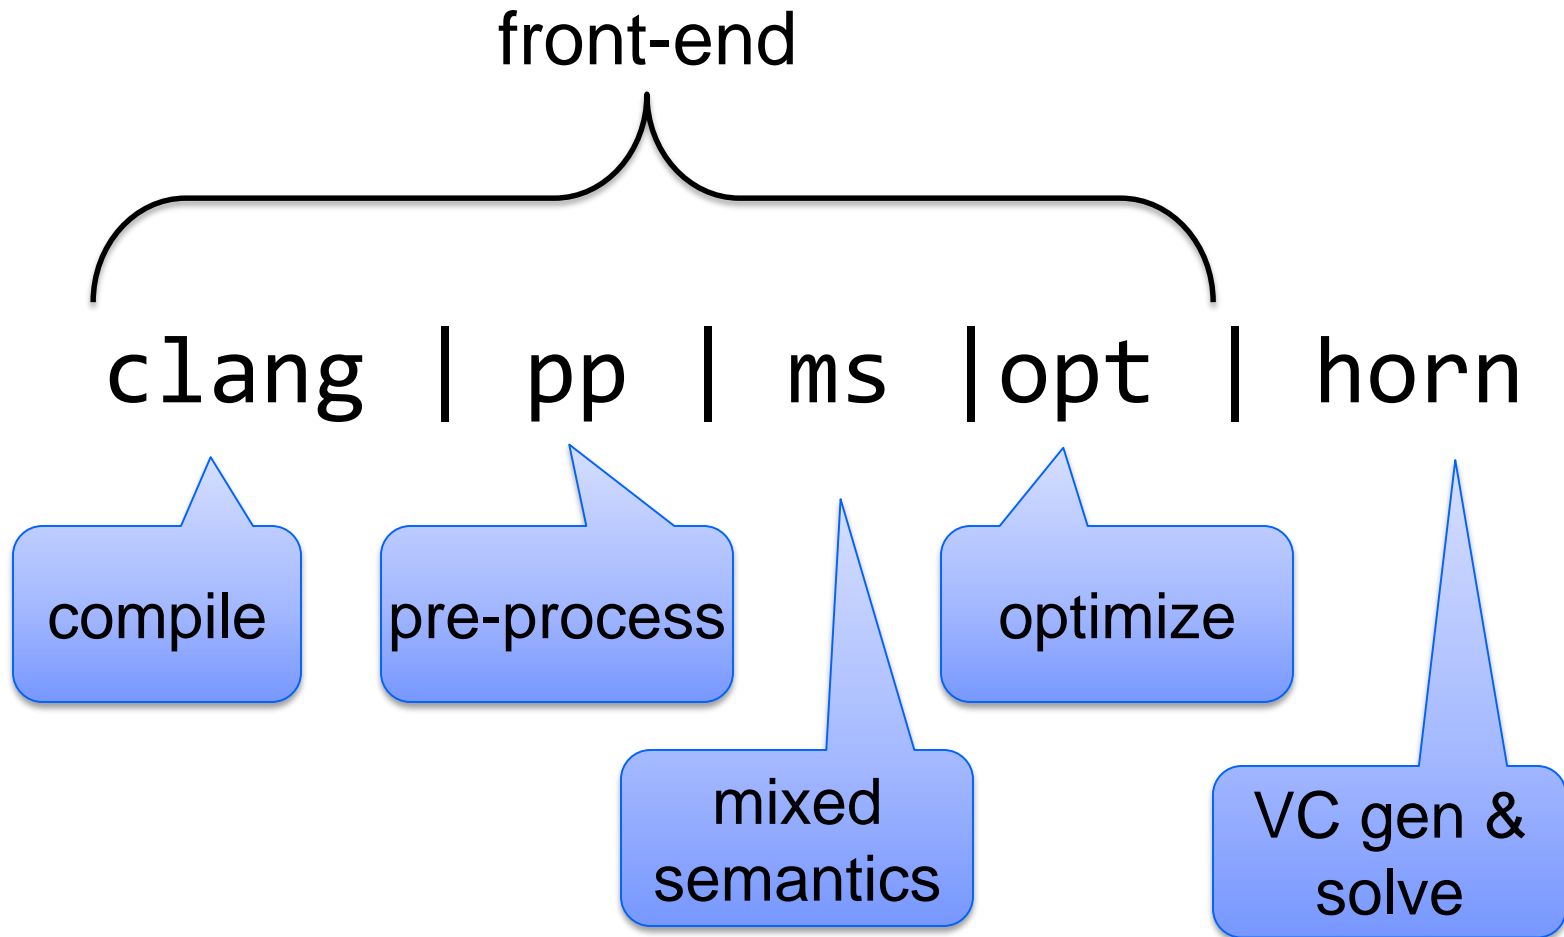
- `--cex=trace.xml` outputs a counter-example in SV-COMP'15 format
- `--track={reg,ptr,mem}` track registers, pointers, memory content
- `--step={large,small}` verification condition step-semantics
 - *small* == basic block, *large* == loop-free control flow block
- `--inline` inline all functions in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



Verification Pipeline



Constrained Horn Clauses

INTERMEDIATE REPRESENTATION



Constrained Horn Clauses (CHC)

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

$$\exists V . (\hat{A} \wedge p_1[X_1] \wedge \dots \wedge p_n[X_n] \rightarrow h[X]),$$

where

- A is a background theory (e.g., Linear Arithmetic, Arrays, Bit-Vectors, or combinations of the above)
- \hat{A} is a constrained in the background theory A
- p_1, \dots, p_n, h are n -ary predicates
- $p_i[X]$ is an application of a predicate to first-order terms

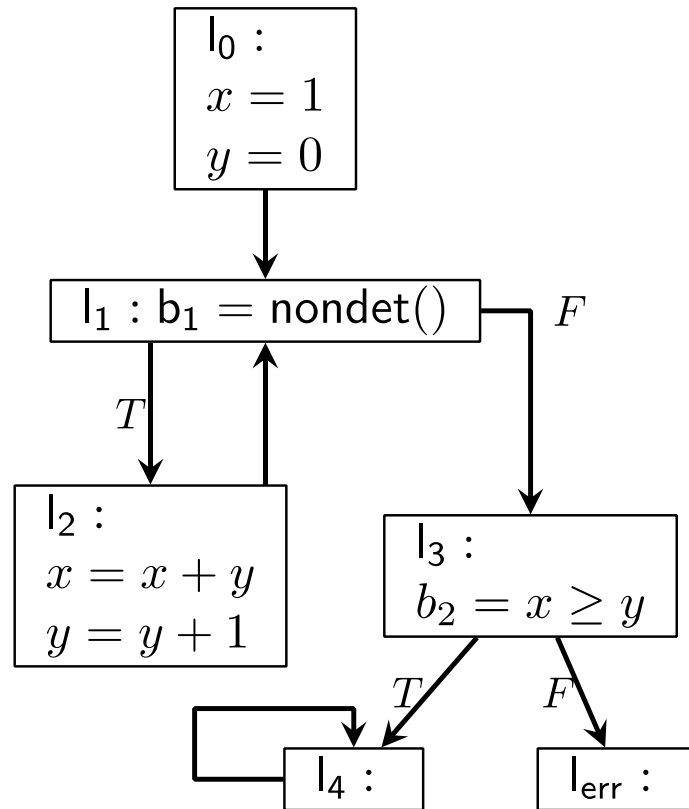


Example Horn Encoding

```

int x = 1;
int y = 0;
while (*) {
    x = x + y;
    y = y + 1;
}
assert(x ≥ y);

```



- ⟨1⟩ p_0 .
- ⟨2⟩ $p_1(x, y) \leftarrow p_0, x = 1, y = 0$.
- ⟨3⟩ $p_2(x, y) \leftarrow p_1(x, y)$.
- ⟨4⟩ $p_3(x, y) \leftarrow p_1(x, y)$.
- ⟨5⟩ $p_1(x', y') \leftarrow p_2(x, y), x' = x + y, y' = y + 1$.
- ⟨6⟩ $p_4 \leftarrow (x \geq y), p_3(x, y)$.
- ⟨7⟩ $p_{\text{err}} \leftarrow (x < y), p_3(x, y)$.
- ⟨8⟩ $p_4 \leftarrow p_4$.
- ⟨9⟩ $\perp \leftarrow p_{\text{err}}$.



CHC Terminology

head

body

constraint

Rule

$$h[X] \tilde{\wedge} p_1[X_1], \dots, p_n[X_n], \acute{A}.$$

Query

$$\text{false} \tilde{\wedge} p_1[X_1], \dots, p_n[X_n], \acute{A}.$$

Fact

$$h[X] \tilde{\wedge} \acute{A}.$$

Linear CHC

$$h[X] \tilde{\wedge} p[X_1], \acute{A}.$$

Non-Linear CHC

$$h[X] \tilde{\wedge} p_1[X_1], \dots, p_n[X_n], \acute{A}.$$

for $n > 1$



CHC Satisfiability

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

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

A model is **A-definable**, if each p_i is definable by a formula \tilde{A}_i in A



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

- 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



FROM PROGRAMS TO CLAUSES



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

$$\frac{Pre \wedge Inv \quad \{Inv \wedge C\} Body \{Inv\} \quad Inv \wedge \neg C \wedge Post}{\{Pre\} \mathbf{while} C \mathbf{do} Body \{Post\}}$$

Function Application

$$\frac{(Pre \wedge p=a) \wedge P \quad \{P\} Body_F \{Q\} \quad (Q \wedge p,r=a,b) \wedge Post}{\{Pre\} b = F(a) \{Post\}}$$

Recursion

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


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 , $Pre \Rightarrow wlp (P, Post)$



Horn Clauses by Weakest Liberal Precondition

Prog = **def** Main(x) { body_M }, ..., **def** P (x) { body_P }

wlp (x=E, Q) = **let** x=E **in** Q

wlp (**assert** (E) , Q) = E \wedge Q

wlp (**assume**(E), Q) = E \rightarrow Q

wlp (**while** E **do** S, Q) = I(w) \wedge

$\exists w . ((I(w) \wedge E) \rightarrow \text{wlp}(S, I(w))) \wedge ((I(w) \wedge :E) \rightarrow Q))$

wlp (y = P(E), Q) = p_{pre}(E) \wedge ($\exists r . p(E, r) \rightarrow Q[r/y]$)

ToHorn (**def** P(x) {S}) = wlp (x0=x ; **assume** (p_{pre}(x)); S, p(x0, ret))

ToHorn (Prog) = wlp (Main(), true) \wedge $\exists\{P \in \text{Prog}\} . \text{ToHorn}(P)$



Example of a WLP Horn Encoding

```
{Pre: y, 0}  
x0 = x;  
y0 = y;  
while y > 0 do  
  x = x+1;  
  y = y-1;  
{Post: x=x0+y0}
```

ToHorn



```
C1:  $I(x, y, x, y) \tilde{\wedge} y \geq 0.$   
C2:  $I(x+1, y-1, x_0, y_0) \tilde{\wedge} I(x, y, x_0, y_0), y > 0.$   
C3:  $\text{false} \tilde{\wedge} I(x, y, x_0, y_0), y \cdot 0, x \neq x_0 + y_0$ 
```

$\{y, 0\} P \{x = x_{\text{old}} + y_{\text{old}}\}$ is **true** iff the query C_3 is **satisfiable**



Dual WLP

Dual weakest liberal pre-condition

$$\mathbf{dual-wlp} (P, \text{Post}) = \mathbf{:wlp} (P, \text{:Post})$$

$s \models \mathbf{dual-wlp} (P, \text{Post})$ iff there exists an execution of P that starts in s and ends in Post

$\mathbf{dual-wlp} (P, \text{Post})$ is the weakest condition ensuring that an execution of P can reach a state in Post



Horn Clauses by Dual WLP

Assumptions

- each procedure is represented by a control flow graph
 - i.e., statements of the form $l_i: S ; \text{goto } l_j$, where S is loop-free
- program is unsafe iff the last statement of $\text{Main}()$ is reachable
 - i.e., no explicit assertions. All assertions are top-level.

For each procedure $P(x)$, create predicates

- $l(w)$ for each label, $p_{\text{en}}(x_0, x, w)$ for entry, $p_{\text{ex}}(x_0, r)$ for exit

The verification condition is a conjunction of clauses:

$$p_{\text{en}}(x_0, x) \tilde{\wedge} x_0 = x$$

$$l_i(x_0, w') \tilde{\wedge} l_j(x_0, w) \wedge \exists :wlp(S, :(w=w')), \text{ for each statement } l_i: S; \text{goto } l_j$$

$$p(x_0, r) \tilde{\wedge} p_{\text{ex}}(x_0, r)$$

$$\text{false} \tilde{\wedge} \text{Main}_{\text{ex}}(x, \text{ret})$$

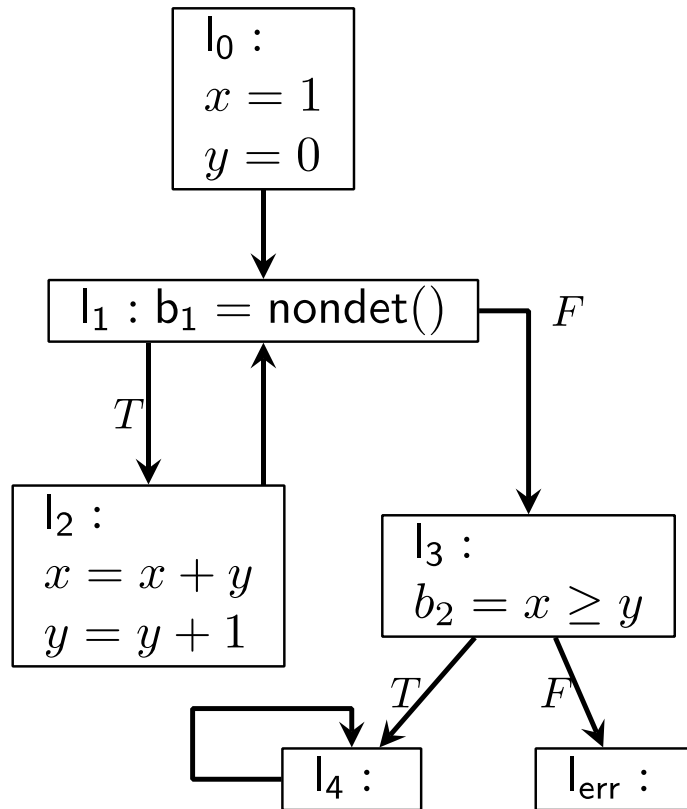


Example Horn Encoding

```

int x = 1;
int y = 0;
while (*) {
    x = x + y;
    y = y + 1;
}
assert(x ≥ y);

```



- ⟨1⟩ p_0 .
- ⟨2⟩ $p_1(x, y) \leftarrow p_0, x = 1, y = 0$.
- ⟨3⟩ $p_2(x, y) \leftarrow p_1(x, y)$.
- ⟨4⟩ $p_3(x, y) \leftarrow p_1(x, y)$.
- ⟨5⟩ $p_1(x', y') \leftarrow p_2(x, y), x' = x + y, y' = y + 1$.
- ⟨6⟩ $p_4 \leftarrow (x \geq y), p_3(x, y)$.
- ⟨7⟩ $p_{\text{err}} \leftarrow (x < y), p_3(x, y)$.
- ⟨8⟩ $p_4 \leftarrow p_4$.
- ⟨9⟩ $\perp \leftarrow p_{\text{err}}$.



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

$$x = \text{PHI} (v_0:\text{bb}_0, \dots, v_n:\text{bb}_n) \quad (\text{phi-assignment})$$

“x gets v_i if previously executed block was bb_i ”



Large Step Encoding: Single Static Assignment

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

```
0: goto 1  
1: x_0 = PHI(0:0, x_3:5);  
   y_0 = PHI(y:0, y_1:5);  
   if (x_0 < N) goto 2 else goto 6  
2: if (y_0 > 0) goto 3 else goto 4  
3: x_1 = x_0 + y_0; goto 5  
4: x_2 = x_0 - y_0; goto 5  
5: x_3 = PHI(x_1:3, x_2:4);  
   y_1 = -1 * y_0;  
   goto 1  
6:
```



Example: Single Static Assignment

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

```
0: goto 1  
1: x_0 = PHI(0:0, x_3:5);  
   y_0 = PHI(y:0, y_1:5);  
   if (x_0 < N) goto 2 else goto 6  
2: if (y_0 > 0) goto 3 else goto 4  
3: x_1 = x_0 + y_0; goto 5  
4: x_2 = x_0 - y_0; goto 5  
5: x_3 = PHI(x_1:3, x_2:4);  
   y_1 = -1 * y_0;  
   goto 1  
6:
```



Example: Large Step Encoding

```
0: goto 1
1: x_0 = PHI(0:0, x_3:5);
   y_0 = PHI(y:0, y_1:5);
   if (x_0 < N) goto 2 else goto 6
2: if (y_0 > 0) goto 3 else goto 4
3: x_1 = x_0 + y_0; goto 5
4: x_2 = x_0 - y_0; goto 5
5: x_3 = PHI(x_1:3, x_2:4);
   y_1 = -1 * y_0;
goto 1
6:
```



Example: Large Step Encoding

$$x_1 = x_0 + y_0$$

$$x_2 = x_0 - y_0$$

$$y_1 = -1 * y_0$$

```
1: x_0 = PHI(0:0, x_3:5);  
   y_0 = PHI(y:0, y_1:5);  
   if (x_0 < N) goto 2 else goto 6  
  
2: if (y_0 > 0) goto 3 else goto 4  
  
3: x_1 = x_0 + y_0 goto 5  
  
4: x_2 = x_0 - y_0 goto 5  
  
5: x_3 = PHI(x_1:3, x_2:4);  
   y_1 = -1 * y_0;  
   goto 1
```



Example: Large Step Encoding

$$x_1 = x_0 + y_0$$

$$x_2 = x_0 - y_0$$

$$y_1 = -1 * y_0$$

$$B_2 \rightarrow x_0 < N$$

$$B_3 \rightarrow B_2 \wedge y_0 > 0$$

$$B_4 \rightarrow B_2 \wedge y_0 \leq 0$$

$$B_5 \rightarrow (B_3 \wedge x_3 = x_1) \vee \\ (B_4 \wedge x_3 = x_2)$$

$$B_5 \wedge x'_0 = x_3 \wedge y'_0 = y_1$$

```
1: x_0 = PHI(0:0, x_3:5);
   y_0 = PHI(y:0, y_1:5);
   if (x_0 < N) goto 2 else goto 6
2: if (y_0 > 0) goto 3 else goto 4
3: x_1 = x_0 + y_0; goto 5
4: x_2 = x_0 - y_0; goto 5
5: x_3 = PHI(x_1:3, x_2:4);
   y_1 = -1 * y_0;
   goto 1
```

$p_1(x'_0, y'_0) \tilde{A} p_1(x_0, y_0), \hat{A}$.

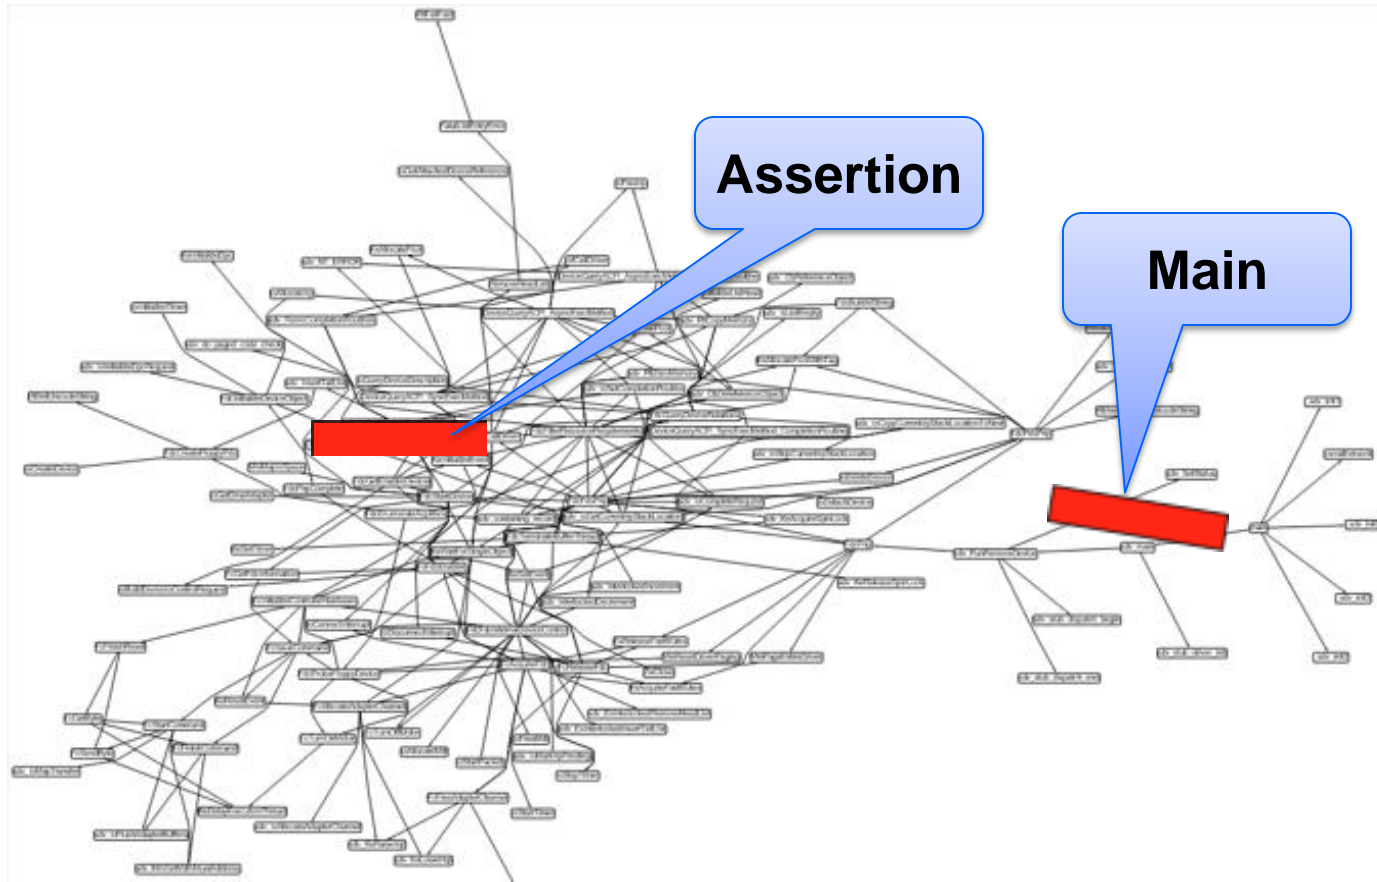


Mixed Semantics

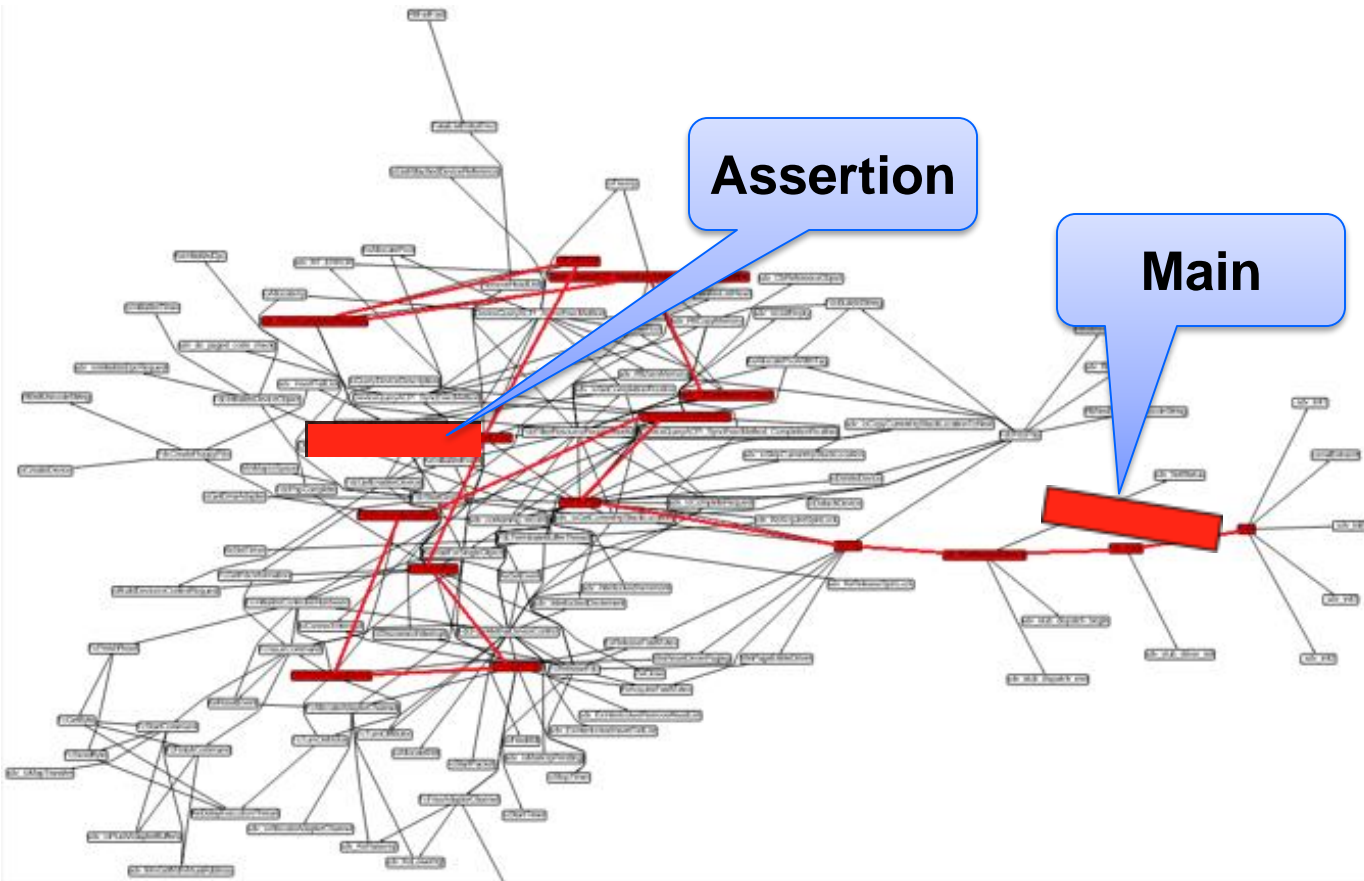
PROGRAM TRANSFORMATION



Deeply nested assertions



Deeply nested assertions



Counter-examples are long

Hard to determine (from main) what is relevant



Mixed Semantics

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]
 - $(\frac{3}{4}, \frac{3}{4}) \models f$ iff the execution of f on input state $\frac{3}{4}$ terminates and results in state $\frac{3}{4}$
- 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

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

$K \models EF (pc=L)$, $K^m \models EF (pc=L)$ and $K \models EG (pc \neq L)$, $K^m \models EG (pc \neq L)$



```

def main()
1: int x = nd();
2: x = x+1;
3: while(x>=0)
4:   x=f(x);
5:   if(x<0)
6:     Error;
7:
8: END;

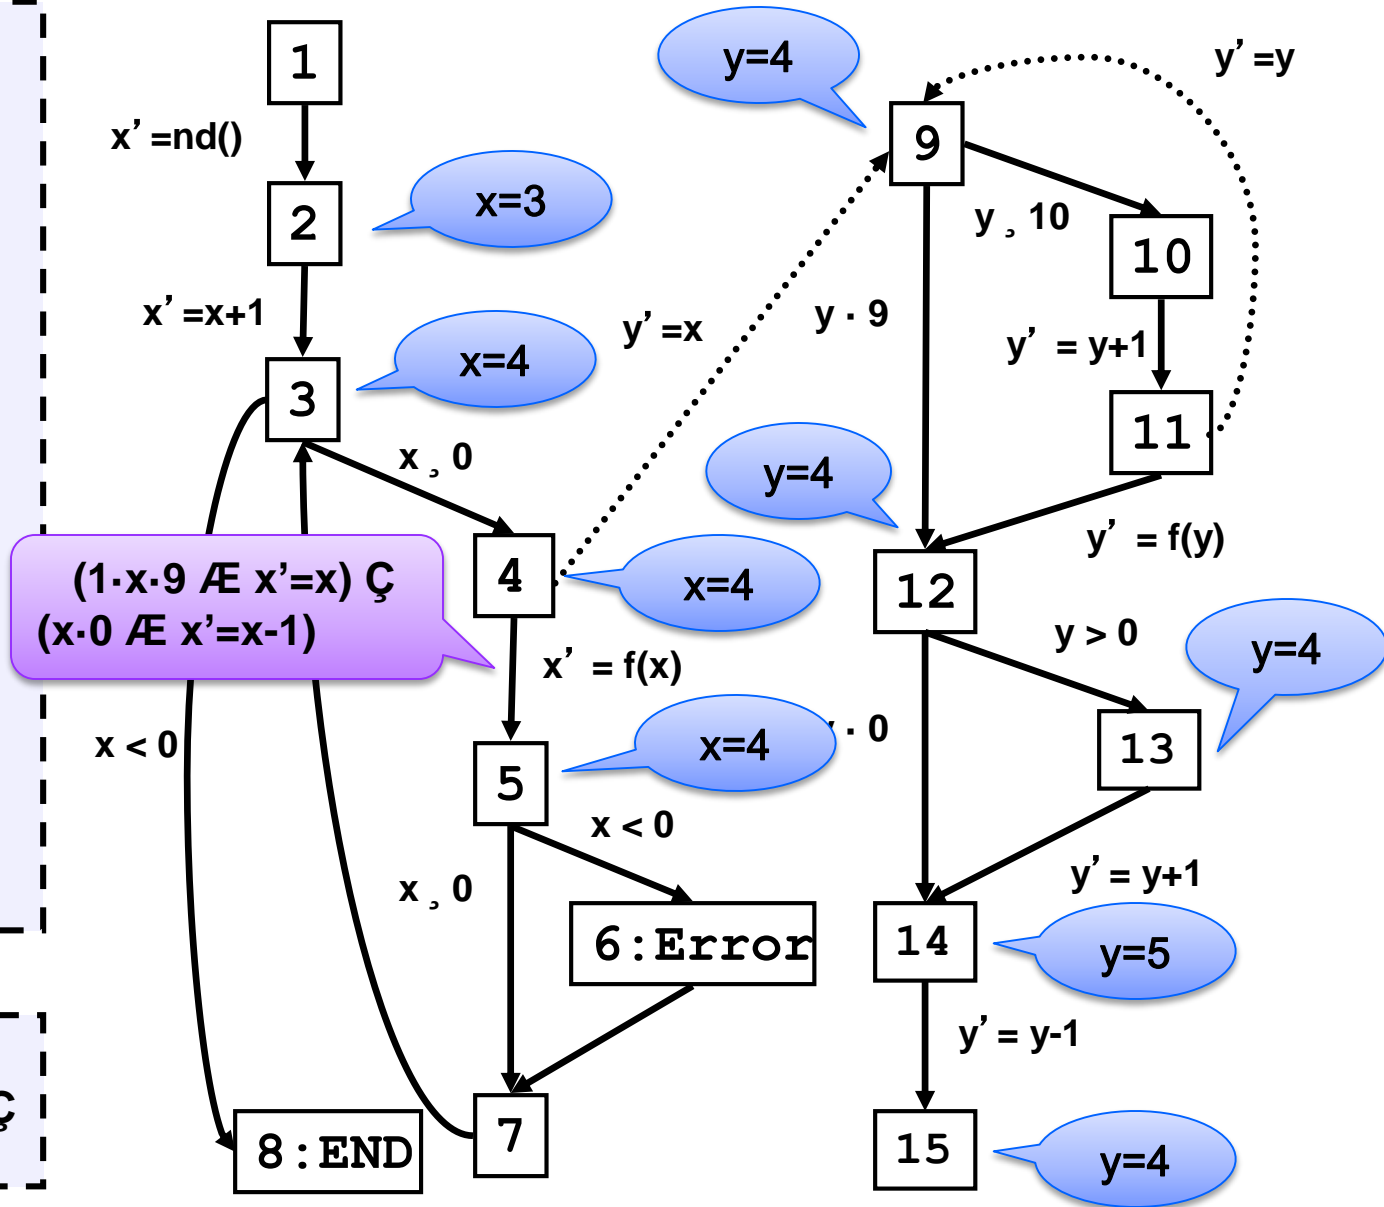
def f(int y): ret y
9: if(y,10){
10:   y=y+1;
11:   y=f(y);
12: else if(y>0)
13:   y=y+1;
14: y=y-1
15:

```

```

Summary of f(y)
(1..y-9 Æ y'=y) Ç
(y-0 Æ y'=y-1)

```



Mixed Semantics as Program Transformation

```
main ()
  p1 (); p1 ();
  assert (c1);
p1 ()
  p2 ();
  assert (c2);
p2 ()
  assert (c3);
```



Mixed Semantics

```
mainnew ()
  if (*) goto p1entry;
  else p1new ();
  if (*) goto p1entry;
  else p1new ();
  if ( $\neg$ c1) goto error;
  assume (false);
p1entry :
  if (*) goto p2entry;
  else p2new ();
  if ( $\neg$ c2) goto error;
p2entry :
  if ( $\neg$ c3) goto error;
  assume (false);
error : assert (false);
p1new ()
  p2new ();
  assume (c2);
p2new ()
  assume (c3);
```



Implementing Mixed Semantics in LLVM

Something about how this can be implemented as a simple transformation in LLVM

in the Lab, show how to do this transformation by hand by modifying the bitcode and using opt to execute the optimization



SOLVING CHC WITH SMT



Programs, Cexs, Invariants

A program $P = (V, \text{Init}, \frac{1}{2}, \text{Bad})$

- Notation: $F(X) = \exists u . (X \in \frac{1}{2}) \subseteq \text{Init}$

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

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

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

$$\left. \begin{array}{l} \text{Init}(u) \Rightarrow \text{Inv}(u) \\ \text{Inv}(u) \wedge \rho(u, v) \Rightarrow \text{Inv}(v) \end{array} \right\} \begin{array}{l} \text{Inductive} \\ \text{Safe} \end{array}$$
$$\text{Inv}(u) \Rightarrow \neg \text{Bad}(u)$$




IC3/PDR Algorithm Overview

```
Input: Transition system  $T = (Init, Tr, Bad)$ 
1  $F_0 \leftarrow Init ; N \leftarrow 0$ 
2 repeat
3    $G \leftarrow \text{PDRMKSAFE}([F_0, \dots, F_N], Bad)$ 
4   if  $G = []$  then return UNSAFE;
5    $\forall 0 \leq i \leq N \cdot F_i \leftarrow G[i]$ 
6    $F_0, \dots, F_N \leftarrow \text{PDRPUSH}([F_0, \dots, F_N])$ 
   //  $F_0, \dots, F_N$  is a safe  $\delta$ -trace
7   if  $\exists 0 \leq i \leq N \cdot F_i = \emptyset$  then return SAFE;
8    $N \leftarrow N + 1 ; F_N \leftarrow \emptyset$ 
9 until  $\infty$ ;
```

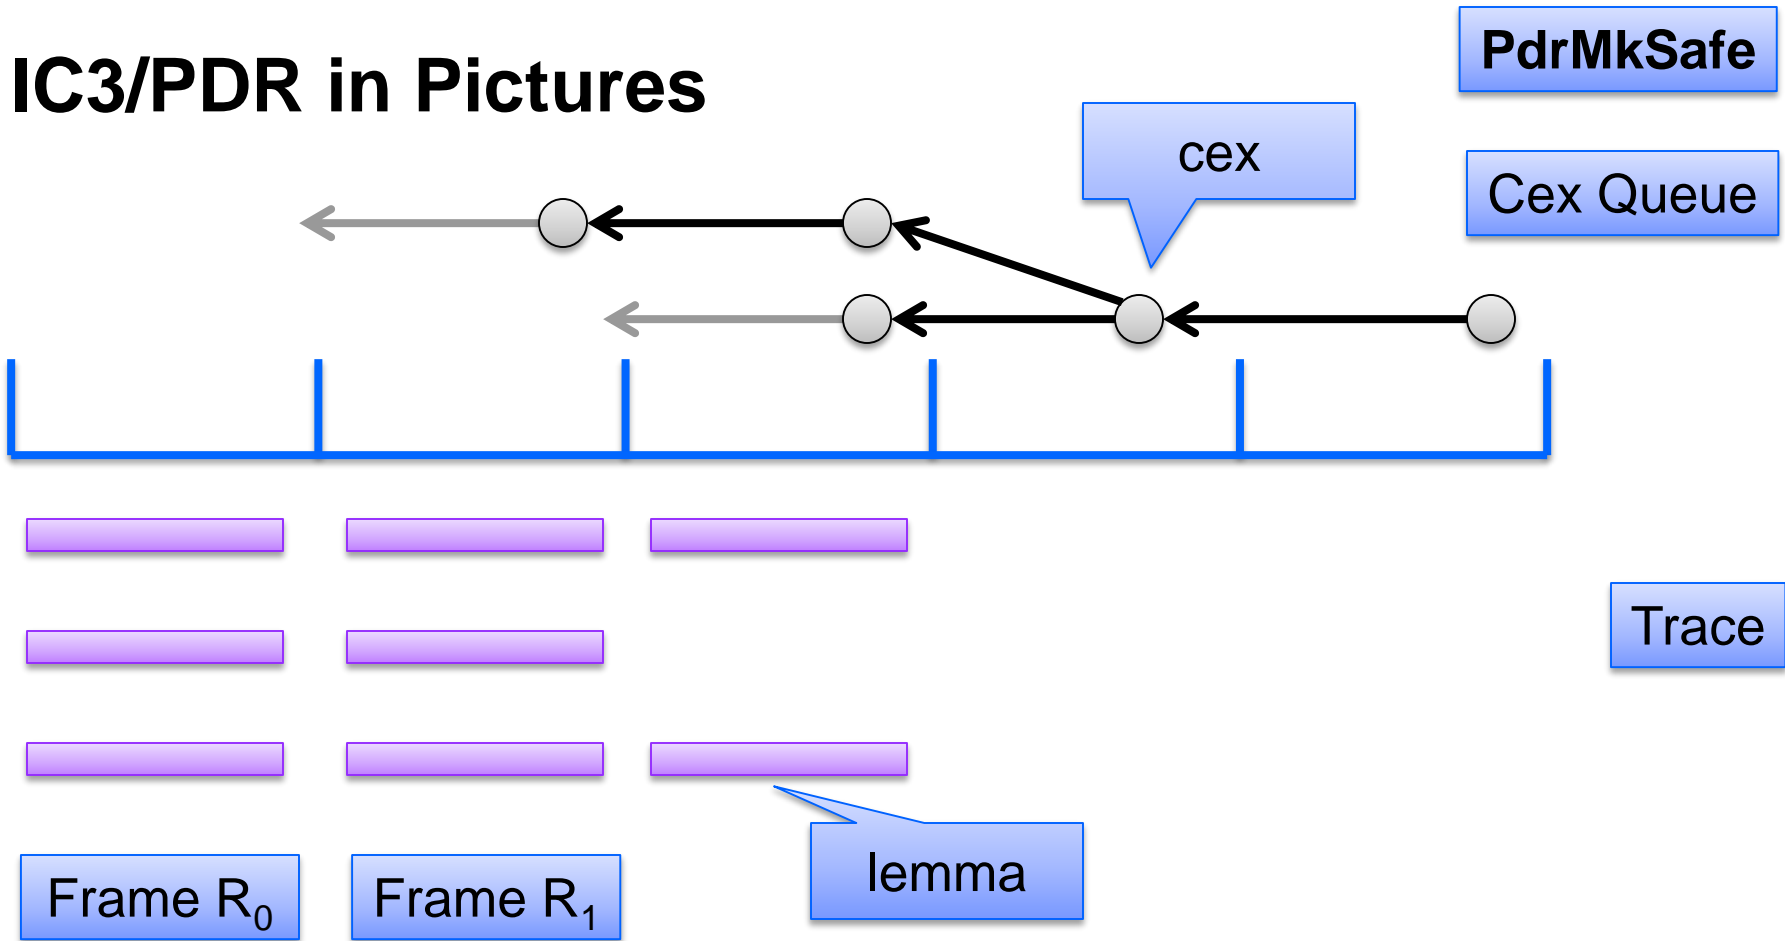
Aaron R. Bradley: SAT-Based Model Checking without Unrolling. VMCAI 2011: 70-87



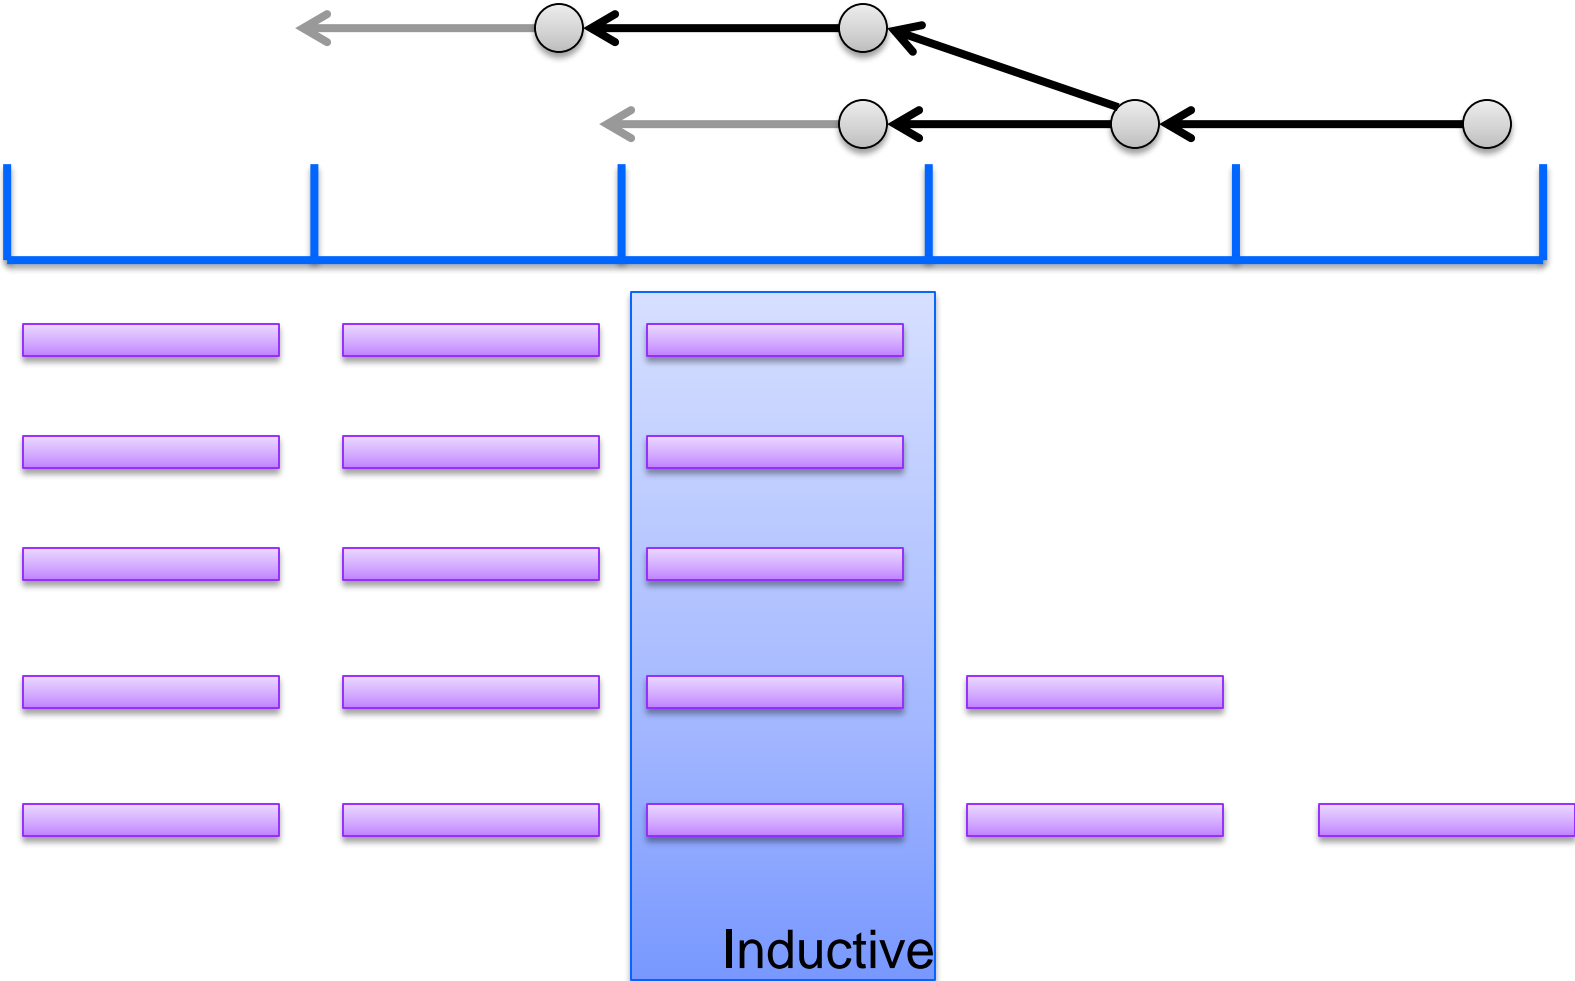
IC3/PDR in Pictures





IC3/PDR in Pictures



IC3/PDR in Pictures




IC3/PDR in Pictures




PDR Invariants

$R_i \rightarrow : \text{Bad}$ $\text{Init} \rightarrow R_i$

$R_i \rightarrow R_{i+1}$ $R_i \not\rightarrow \frac{1}{2} \rightarrow R_{i+1}$



Spacer: Solving CHC in Z3



Spacer: solver for SMT-constrained Horn Clauses

- stand-alone implementation in a fork of Z3
- <http://bitbucket.org/spacer/code>

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



SV-COMP 2015

<http://sv-comp.sosy-lab.org/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>



Results for DeviceDriver category



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



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

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

