

Applications and Challenges in Satisfiability Modulo Theories WING/ETAPS 2009–York

Leonardo de Moura Microsoft Research

Symbolic Reasoning

Verification/Analysis tools need some form of Symbolic Reasoning



Applications

Test case generation

Verifying Compilers

Predicate Abstraction

Invariant Generation

Type Checking

Model Based Testing

Research

Some Applications @ Microsoft



Test case generation



Is formula *F* satisfiable modulo theory *T* ?

SMT solvers have specialized algorithms for *T*



b + 2 = c and $f(read(write(a,b,3), c-2) \neq f(c-b+1))$



b + 2 = c and $f(read(write(a,b,3), c-2) \neq f(c-b+1))$

Arithmetic



b + 2 = c and $f(read(write(a,b,3), c-2) \neq f(c-b+1))$

Array Theory



b + 2 = c and $f(read(write(a,b,3), c-2) \neq f(c-b+1))$

Uninterpreted Functions



Theories

- A Theory is a set of sentences
- Alternative definition:
 A Theory is a class of structures



SMT@Microsoft: Solver

- Z3 is a new solver developed at Microsoft Research.
- Development/Research driven by internal customers.
- Free for academic research.
- Interfaces:



<u>http://research.microsoft.com/projects/z3</u>





For some theories, SMT can be reduced to SAT

Higher level of abstraction

 $bvmul_{32}(a,b) = bvmul_{32}(b,a)$



Ground formulas

For most SMT solvers: F is a set of ground formulas

Many Applications Bounded Model Checking Test-Case Generation



DPLL





DPLL

• Guessing (case-splitting) $p \mid p \lor q, \neg q \lor r$ $p, \neg q \mid p \lor q, \neg q \lor r$



DPLL

• Deducing $p \mid p \lor q, \neg p \lor s$ $p, s \mid p \lor q, \neg p \lor s$





Backtracking p, ¬s, q | $p \lor q, s \lor q, \neg p \lor \neg q$ p, s | $p \lor q, s \lor q, \neg p \lor \neg q$



Modern DPLL

- Efficient indexing (two-watch literal)
- Non-chronological backtracking (backjumping)
- Lemma learning



Solvers = DPLL + Decision Procedures

 Efficient decision procedures for conjunctions of ground atoms.

a=b, a<5 | ¬a=b ∨ f(a)=f(b), a < 5 ∨ a > 10

Efficient algorithms

Difference Logic	Belmann-Ford
Uninterpreted functions	Congruence closure
Linear arithmetic	Simplex



Theory Conflicts

a=b, a > 0, c > 0, a + c < 0 | F



Naïve recipe?

SMT Solver = DPLL + Decision Procedure

Standard question:

Why don't you use CPLEX for handling linear arithmetic?



Efficient SMT solvers

Decision Procedures must be: Incremental & Backtracking Theory Propagation

Research

Efficient SMT solvers

Decision Procedures must be: Incremental & Backtracking Theory Propagation Precise (theory) lemma learning a=b, a > 0, c > 0, a + c < 0 | F Learn clause: $\neg(a=b) \lor \neg(a>0) \lor \neg(c>0) \lor \neg(a+c<0)$ Imprecise! Precise clause: $\neg a > 0 \lor \neg c > 0 \lor \neg a + c < 0$

Applications and Challenges in Satisfiability Modulo Theories

Research



Verifying Compilers



pre/post conditions invariants and other annotations

Annotations: Example

class C {
 private int a, z;
 invariant z > 0

public void M()
 requires a != 0
 {
 z = 100/a;
}



Modeling execution traces



States and execution traces

- State
 - Cartesian product of variables
- Execution trace
 - Nonempty finite sequence of states
 - Infinite sequence of states
 - Nonempty finite sequence of states followed by special error state

(x: int, y: int, z: bool)

Command language



• x := 10

havoc x





• assert P P - P



Command language



• x := 10





● S;T



assume P

assert P

P -{ ● → ●



Command language



Reasoning about execution traces

- Hoare triple { P } S { Q } says that every terminating execution trace of S that starts in a state satisfying P
 - does not go wrong, and
 - terminates in a state satisfying Q

Reasoning about execution traces

- Hoare triple { P } S { Q } says that every terminating execution trace of S that starts in a state satisfying P
 - does not go wrong, and
 - terminates in a state satisfying Q
- Given S and Q, what is the weakest P' satisfying {P'} S {Q} ?
 - P' is called the *weakest precondition* of S with respect to Q, written wp(S, Q)
 - to check {P} S {Q}, check $P \Rightarrow P'$

Weakest preconditions

- wp(x := E, Q) =
- wp(havoc x, Q) =
- wp(assert P, Q) =
- wp(assume P, Q) =
- wp(S;T,Q) =

● wp(S□T,Q) =

Q[E/x] (∀x • Q) $P \wedge Q$ $P \Longrightarrow Q$ wp(S, wp(T,Q)) wp(S,Q) \land wp(T,Q)

Structured if statement

if E then S else T end =

assume E; S

assume –E; T

Dijkstra's guarded command

if $E \rightarrow S \mid F \rightarrow T$ fi =

assert E ∨ F;
(
 assume E; S
 □
 assume F; T
While loop with loop invariant

while E invariant J where x denotes the assignment targets of S do end check that the loop invariant holds initially assert J; "fast forward" to an arbitrary
 iteration of the loop havoc x; assume J; assume E; S; assert J; assume false assume ¬E check that the loop invariant is maintained by the loop body



Verification conditions: Structure



Hypervisor: A Manhattan Project



- Meta OS: small layer of software between hardware and OS
- Mini: 60K lines of non-trivial concurrent systems C code
- Critical: must provide functional resource abstraction
- **Trusted**: a verification grand challenge

Hypervisor: Some Statistics

- VCs have several Mb
- Thousands of non ground clauses
- Developers are willing to wait at most 5 min per VC



Challenge: annotation burden

Partial solutions

- Automatic generation of: Loop Invariants
- Houdini-style automatic annotation generation



- Quantifiers, quantifiers, quantifiers, ...
- Modeling the runtime
 - ∀ h,o,f:
 IsHeap(h) ∧ o ≠ null ∧ read(h, o, alloc) = t
 ⇒
 read(h,o, f) = null ∨ read(h, read(h,o,f),alloc) = t



- Quantifiers, quantifiers, quantifiers, ...
- Modeling the runtime
- Frame axioms
 - ∀ o, f:
 - o ≠ null ∧ read(h₀, o, alloc) = t ⇒ read(h₁,o,f) = read(h₀,o,f) ∨ (o,f) ∈ M



- Quantifiers, quantifiers, quantifiers, ...
- Modeling the runtime
- Frame axioms
- User provided assertions
 - $\forall i,j: i \leq j \Rightarrow read(a,i) \leq read(b,j)$



- Quantifiers, quantifiers, quantifiers, ...
- Modeling the runtime
- Frame axioms
- User provided assertions
- Theories
 - ∀ x: p(x,x)
 - $\forall x,y,z: p(x,y), p(y,z) \Longrightarrow p(x,z)$
 - $\forall x,y: p(x,y), p(y,x) \Longrightarrow x = y$



- Quantifiers, quantifiers, quantifiers, ...
- Modeling the runtime
- Frame axioms
- User provided assertions
- Theories
- Solver must be fast in satisfiable instances.



We want to find bugs!





There is no sound and refutationally complete procedure for linear integer arithmetic + free function symbols



Many Approaches

Heuristic quantifier instantiation

Combining SMT with Saturation provers

Complete quantifier instantiation

Decidable fragments

Model based quantifier instantiation



E-matching & Quantifier instantiation

- SMT solvers use heuristic quantifier instantiation.
- E-matching (matching modulo equalities).

Example:





E-matching & Quantifier instantiation

- SMT solvers use heuristic quantifier instantiation.
- E-matching (matching modulo equalities).
- Example:



E-matching: why do we use it?

- Integrates smoothly with DPLL.
- Efficient for most VCs
- Decides useful theories:
 - Arrays
 - Partial orders

Θ.



Efficient E-matching

- E-matching is NP-Hard.
- In practice

Problem	Indexing Technique
Fast retrieval	E-matching code trees
Incremental E-Matching	Inverted path index



E-matching code trees



Similar triggers share several instructions.

Combine code sequences in a code tree Instructions:

- 1. init(f*,* 2)
- 2. check(r4, b, 3)
- 3. bind(r2, g, r5, 4)
- 4. compare(r1, r5, 5)
- 5. check(r6, a, 6)
- 6. bind(r3, h, r7, 7)
- 7. yield(r1, r7)



Applications and Challenges in Satisfiability Modulo Theories

Compiler

Challenge: modeling runtime

- Is the axiomatization of the runtime consistent?
- False implies everything
- E-matching doesn't work
 No ground terms to instantiate clauses
- Partial solution: SMT + Saturation Provers
- Found many bugs using this approach





Tight integration: DPLL + Saturation solver.







Inference rule:

$$\frac{C_1 \quad \dots \quad C_n}{C}$$

- DPLL(Γ) is parametric.
- Examples:
 - Resolution
 - Superposition calculus
 - ⊜..









Challenge: Robustness

Standard complain

"I made a small modification in my Spec, and Z3 is timingout"

- This also happens with SAT solvers (NP-complete)
- In our case, the problems are undecidable
- Partial solution: parallelization





- Joint work with Y. Hamadi (MSRC) and C. Wintersteiger
- Multi-core & Multi-node (HPC)
- Different strategies in parallel
- Collaborate exchanging lemmas



Microsoft[®]

Kecear

Challenge: Non-Linear Arithmetic

- Non-linear arithmetic is necessary for verifying embedded and hybrid systems
- Non-linear integer arithmetic is undecidable
- Many approaches for non linear real arithmetic
 - Cylindrical Algebraic Decomposition
 Doubly exponential procedure
 - Grobner Basis + "extensions"
 - Heuristics





Predicate Abstraction & Invariant Generation

Overview

- http://research.microsoft.com/slam/
- SLAM/SDV is a software model checker.
- Application domain: *device drivers*.
- Architecture:
 - c2bp C program → boolean program (*predicate abstraction*).
 bebop Model checker for boolean programs.
 newton Model refinement (check for path feasibility)
- SMT solvers are used to perform predicate abstraction and to check path feasibility.
- c2bp makes several calls to the SMT solver. The formulas are relatively small.



Predicate Abstraction: c2bp

- **Given** a C program *P* and $F = \{p_1, \dots, p_n\}$.
- Produce a Boolean program B(P, F)
 - Same control flow structure as P.
 - Boolean variables $\{b_1, \dots, b_n\}$ to match $\{p_1, \dots, p_n\}$.
 - Properties true in *B*(*P*, *F*) are true in *P*.
- Each p_i is a pure Boolean expression.
- Each p_i represents set of states for which p_i is true.
- Performs modular abstraction.



Abstracting Expressions via F

Implies_F (e)

- Best Boolean function over F that implies e.
- ImpliedBy_F (e)
 - Best Boolean function over F that is implied by e.
 - ImpliedBy_F (e) = not Implies_F (not e)



Implies_F(e) and ImpliedBy_F(e)



Computing Implies_F(e)

- minterm $m = l_1 \land ... \land l_n$, where $l_i = p_i$, or $l_i = not p_i$.
- Implies_F (e): disjunction of all minterms that imply e.
- Naive approach
 - Generate all 2ⁿ possible minterms.
 - For each minterm m, use SMT solver to check validity of m ⇒ e.
- Many possible optimizations



Computing Implies_F(e)

- F = { x < y, x = 2}
- *e*:y>1
- Minterms over F
 - !x<y, !x=2 implies y>1
 - x<y, !x=2 implies y>1 🚫
 - !x<y, x=2 implies y>1
 - x<y, x=2 implies y>1

 $Implies_{f}(y>1) = b_{1}^{2}y \wedge b_{2}^{2}=2$



Challenge: Rich API

All-SAT

- Better (more precise) Predicate Abstraction
- Unsatisfiable cores
 - Why the abstract path is not feasible?
 - Fast Predicate Abstraction



Unsatisfiable cores

- Let S be an unsatisfiable set of formulas.
- $S' \subseteq S$ is an unsatisfiable core of S if:
 - S' is also unsatisfiable, and
 - There is not $S'' \subset S'$ that is also unsatisfiable.
- Computing Implies_F(e) with $F = \{p_1, p_2, p_3, p_4\}$
 - Assume $p_1, p_2, p_3, p_4 \Rightarrow e$ is valid
 - That is $p_1, p_2, p_3, p_4, \neg e$ is unsat
 - Now assume $p_1, p_3, \neg e$ is the unsatisfiable core
 - Then it is unnecessary to check:

•
$$p_1, \neg p_2, p_3, p_4 \Rightarrow e$$

•
$$p_1, \neg p_2, p_3, \neg p_4 \Rightarrow e$$

• $p_1, p_2, p_3, \neg p_4 \Rightarrow e$

Research

Loop invariants



How to find loop invariant /?

Template based approach

- I is a Boolean combination of $F = \{p_1, \dots, p_n\}$
- Unknown invariant on the LHS constraints how weak / can be

 $I(\mathbf{x}) \wedge \neg c(\mathbf{x}) \Rightarrow Post(\mathbf{x})$ $I(\mathbf{x}) \Rightarrow \neg c(\mathbf{x}) \Rightarrow Post(\mathbf{x})$

 Unknown invariant on the RHS constraints how strong / can be

 $\Theta(\mathbf{x}) \Longrightarrow \mathbf{I}(\mathbf{x})$

 More details: Constraint-based Invariant Inference over Predicate Abstraction, S. Gulwani et al, VMCAI 2009





Bit-precise test case generation
Test case generation

unsigned GCD(x, y) {	$(y_0 > 0)$ and	$x_0 = 2$
requires($y > 0$);	$(m_0 = x_0 \% y_0)$ and	$v_{0} = 4$
while (true) {	not $(m_0 = 0)$ and Solver	$m_0 = 2$
unsigned m = x % y;	$(x_1 = y_0)$ and	x = 4
if (m == 0) return y;	(y = m) and	χ_1
x = y;		y ₁ = 2
y = m;	$(m_1 = x_1 \% y_1)$ and	$m_{1}^{2} = 0$
}	$(m_1 = 0)$	



Applications and Challenges in Satisfiability Modulo Theories

Challenge: Machine Arithmetic

- Most solvers use bit-blasting
- **bvmul**₃₂(a,b) is converted into a multiplier circuit
- Solvers may run out of memory
- "Smart" algorithms are usually less efficient than bit-blasting

Challenge: Floating Point Arithmetic

- I'm unaware of any SMT solver for floating point arithmetic
- Approximate using Reals
 Unsound!
 Incomplete!



Applications and Challenges in Satisfiability Modulo Theories

Conclusion

- Logic as a platform
- Most verification/analysis tools need symbolic reasoning
- SMT is a hot area
- Many applications & challenges
- http://research.microsoft.com/projects/z3





Applications and Challenges in Satisfiability Modulo Theories