

# Efficient Sampling of SAT solutions for Testing (ICSE '18)

Rafael Dutra, Kevin Laeufer, Jonathan Bachrach and Koushik Sen (EECS Department, UC Berkeley)

# Background

- In software testing, generating a lot random solutions to the constraints is a important problem.
  - Conventional symbolic execution and dynamic symbolic execution uses SMT solver to generate ONE solution for the path constraint.
    - Not very scalable due to path explosion
  - Generating multiple solutions to constraint can test multiple paths having the same path prefix

# Background: SAT problem

- SAT problem: determining if there exists an assignment (of variables) which satisfies a boolean formula. (First problem proven to be NP-complete)
- How to solve?
  - DPLL algorithms (introduced in 60s, still the basis for modern solvers)
    - CNF form (  $(a \vee \neg b) \wedge (\neg a \vee b)$  ) One solution: [1,1]
    - Backtracking: Assign true/false for one variable, and then solve the sub-problems (branching/splitting).
    - Pruning: Unit propagation/Pure literal elimination
    - Heuristics: Which variable to try first? (e.g. the variable that has the most occurrences )
  - Non-DPLL algorithms
    - Stochastic Local Search (WalkSAT)
      - Pick a random assignment, then try to flip one variable.

# Background: Translate SMT problem to SAT problem

- A SMT problem asks to decide if a logic-formula (background theories expressed in first-order logic) can be satisfied.
- Eager approach - encoding and translating (bit-blasting)
  - Example  $(x \neq 0) \wedge (y \mid 2 = z)$ 
    - Let  $x=[b1,b2]$ ,  $y=[b3,b4]$ ,  $z=[b5,b6]$
    - $b1 \vee b2, (b3 = 1) \wedge (b4 = b6)$
    - CNF form:  $(b1 \vee b2) \wedge (b3) \wedge (b4 \vee \neg b6) \wedge (\neg b4 \vee b6)$
- Lazy approach
  - First asks SAT for an assignment and then checks for consistency.
  - Example:  $(x > 0 \vee y = 100) \wedge (x < 3 \vee y = 200)$ 
    - SAT solver assigns [False, False] to  $(x > 0)$ ,  $(x < 3)$ . Inconsistency!
    - [True, True], [True, False], [False, True] all are satisfiable assignments.

# Problem: how to get multiple assignments quickly and uniformly

- Uniformity:
  - Given the set of all satisfiable assignments  $R$ , the solutions should be uniformly sampled from  $R$ .
- Benefits of uniformity:
  - Ensure the diversity the inputs, exploring more program states.

# Related works (baselines):

- Based on universal hashing (e.g. UniGen)
  - Idea: select a set of universal hashing functions to uniformly partition the solution sparse and then plug the hash function (XOR of boolean variables) to the constraints ( e.g. original constraints  $\wedge$  hash function )
  - Strong uniformity guarantee.
  - Bad performance (for each sampling, needs to make a call to SAT)
- SearchTreeSampler
  - Also uses SAT solver as a black-box
  - Maintain a tree of pseudo-solutions. Level  $i$  in the tree stores partial-solutions with the first  $i$  boolean variables assigned.
  - Recursively build the tree
    - Sample a pseudo-solutions uniformly from level  $i$ , and then call SAT to enumerate all satisfiable pseudo-solutions in level  $i+1$  ( e.g. original constraints  $\wedge$  pseudo-solution of  $i$   $\wedge$  new probing bits in level  $i+1$ )
- Others: treat SAT solver as white boxes

# Quick Sampler Algorithm

- Overall idea:
  - Sample a random solution
  - Explore the neighbors (delta-mutations)
  - Combining two mutations
- Intuition:
  - $\delta_a$  and  $\delta_b$  consist of a minimal set of bits which can be flipped while still preserving the satisfiability of the formula.
  - So the bits in  $\delta_a$  are likely to be closely related to each other by some clauses in the formula.
  - It is likely that those clauses would still be satisfied in  $\sigma \oplus (\delta_a \vee \delta_b)$ , where we flip all the bits from  $\delta_a$  in addition to the bits from  $\delta_b$ .

$$\begin{aligned}\sigma &: 0 1 0 0 0 1 0 1 1 0 1 1 \\ \delta_a &: \textcolor{red}{1} 0 0 0 \textcolor{red}{1} \textcolor{red}{1} 0 0 0 0 0 0 \\ \sigma_a = \sigma \oplus \delta_a &: \textcolor{blue}{1} 1 0 0 \textcolor{blue}{1} 0 0 1 1 0 1 1 \\ \delta_b &: 0 \textcolor{red}{1} 0 0 0 \textcolor{red}{1} \textcolor{red}{1} 0 \textcolor{red}{1} 0 0 0 \\ \sigma_b = \sigma \oplus \delta_b &: 0 \textcolor{blue}{0} 0 0 0 0 \textcolor{blue}{1} 1 0 0 1 1 \\ (\delta_a \vee \delta_b) &: \textcolor{red}{1} 1 0 0 \textcolor{red}{1} \textcolor{red}{1} \textcolor{red}{1} 0 \textcolor{red}{1} 0 0 0 \\ \tilde{\sigma} = \sigma \oplus (\delta_a \vee \delta_b) &: \textcolor{blue}{1} 0 0 0 \textcolor{blue}{1} 0 \textcolor{blue}{1} 1 0 0 1 1\end{aligned}$$

**Figure 1: Combining two mutations.**

# Implementation


- Use SAT solver as an oracle, to answer MAX-SAT queries
- MAX-SAT query
  - Given a set of hard constraints and a set of soft constraints, can you satisfy all the hard constraints and maximum possible number of soft constraints.
- How to find a random solution?
  - Randomly assign boolean variables.
  - Then call MAX-SAT(hard, soft),
    - where hard is the original constraints,
    - soft is that the **assignment for each variable = randomly assigned one**
- How to find a delta?
  - For one solution, flip one bit of it
  - Then call MAX-SAT(hard, soft)
    - where hard is the original constraints ^ flipped bit must be flipped
    - Soft is that **assignment for each variable = original one**

# Evaluation


- Correctness: 75%
- Performance

Benchmark	<i>S</i>	Vars	Clauses	Solutions	QUICKSAMPLER					SEARCHTREESAMPLER		UNIGEN2		
					<i>n</i>	Calls	Samples	Valid	<i>t<sub>q</sub></i> (μs)	<i>t<sub>q</sub><sup>*</sup></i> (μs)	Samples	<i>t<sub>s</sub></i> / <i>t<sub>q</sub></i>	Samples	<i>t<sub>u</sub></i> / <i>t<sub>q</sub></i>
blasted_case47	28	118	328	262144	244	6616	10010929	0.564	7.5	26	11694350	41.3	3932170	426
blasted_case110	17	287	1263	16384	1387	22208	10001202	0.822	28.3	29	8502350	14.9	245762	34

- Uniqueness



# Evaluation - Uniformity



**Figure 6: blasted\_case47 histogram**

# More related works: Sampling using SMT solver as an oracle


- **PANGOLIN: Incremental Hybrid Fuzzing with Polyhedral Path Abstraction (S&P 2020)**
- Treat SMT solver as an oracle, determining the path abstraction (range)
- Sampling the range using Dikin walk algorithm

$$x \leq 2 \wedge y \leq 5 \wedge x^2 - 5x + 4 \leq y \quad (1)$$

$$\begin{cases} 0 \leq x \leq 2 \\ 0 \leq y \leq 5 \\ 4 \leq 5x + y \leq 15 \end{cases}$$

**Algorithm 1** Polyhedral path abstraction inference.


```
1: procedure INFERENCE( $pc \triangleq \sigma_1 \wedge \sigma_2 \dots \wedge \sigma_n$ )
2:    $pc$ , path constraint.  $\hat{pc}$ , polyhedral path abstraction.
3:
4:    $\hat{pc} \leftarrow \text{true}$ 
5:   for all input variable  $v_i$  in  $pc$  do
6:      $min \leftarrow SMToptMin(v_i, pc)$ 
7:      $max \leftarrow SMToptMax(v_i, pc)$ 
8:      $\hat{pc} \leftarrow \hat{pc} \wedge min \leq v_i \leq max$ 
9:   end for
10:  for all atomic predicate  $\sigma_i$  in  $pc$  do
11:    if  $\sigma_i$  contains linear expression  $\iota_i$  then
12:       $min \leftarrow SMToptMin(\iota_i, pc)$ 
13:       $max \leftarrow SMToptMax(\iota_i, pc)$ 
14:       $\hat{pc} \leftarrow \hat{pc} \wedge min \leq \iota_i \leq max$ 
15:    end if
16:  end for
17:
18:  return  $\hat{pc}$ 
19: end procedure
```



# How does SMT-opt work?

Symbolic Optimization with SMT  
Solvers (POPL '14)


$$\varphi \equiv 0 \leq x \leq 3 \wedge 0 \leq z \leq 2 \wedge (2y \leq -x + 4 \vee 4y = 3x + 3),$$



**Figure 2.** Illustration of SYMBA on a 3-D example.

# More related works: leverage extra information to speedup SMT/SAT solving

- Range
  - Pangolin (SP20) - add the range constraints to the formula
  - Trident (ISSTA20)
    - Assigning boolean variables before search
      - E.g if we know,  $x < C$  (we reduce 32 bits to  $\log C$  bits in the vector)
- Variable dependency (add another heuristic)



**Figure 3: Data-dependence graph for the constraint  $\phi \equiv x = u + w \wedge y = 2 * u - w \wedge x + 2 * y < 10 \wedge u * w < 60$ .**