

Symbolic Execution and Advanced Test Coverage Criteria

Nikolai Kosmatov

joint work with Sébastien Bardin, Omar Chebaro, Mickaël
Delahaye. . .

CEA, LIST, Software Security Lab
Paris-Saclay, France

USE 2015, Oslo, June 23, 2015

Dynamic Symbolic Execution [dart,cute,pathcrawler,exe,sage,pex,klee,...]

- ✓ very powerful approach to white-box test generation
- ✓ many tools and many successful case-studies since mid 2000's
- ✓ arguably one of the most wide-spread use of formal methods in “common software” [SAGE at Microsoft]

Dynamic Symbolic Execution [dart,cute,pathcrawler,exe,sage,pex,klee,...]

- ✓ very powerful approach to white-box test generation
- ✓ many tools and many successful case-studies since mid 2000's
- ✓ arguably one of the most wide-spread use of formal methods in “common software” [SAGE at Microsoft]

Symbolic Execution [King 70's]

- consider a program P on input v , and a given path σ
- a **path predicate** φ_σ for σ is a formula s.t. for any input v
 v satisfies $\varphi_\sigma \Leftrightarrow P(v)$ follows σ
- old idea, recently renewed interest [requires powerful solvers]

Dynamic Symbolic Execution [dart,cute,pathcrawler,exe,sage,pex,klee,...]

- ✓ very powerful approach to white-box test generation
- ✓ many tools and many successful case-studies since mid 2000's
- ✓ arguably one of the most wide-spread use of formal methods in “common software” [SAGE at Microsoft]

Symbolic Execution [King 70's]

- consider a program P on input v , and a given path σ
- a **path predicate** φ_σ for σ is a formula s.t. for any input v
 v satisfies $\varphi_\sigma \Leftrightarrow P(v)$ follows σ
- old idea, recently renewed interest [requires powerful solvers]

Dynamic Symbolic Execution [Korel+, Williams+, Godefroid+]

- interleaves dynamic and symbolic executions
- drives the search towards feasible paths for free
- gives hints for relevant under-approximations

Dynamic Symbolic Execution (2)

input : a program P

output : a test suite TS covering all feasible paths of $Paths^{\leq k}(P)$

- pick an uncovered path $\sigma \in Paths^{\leq k}(P)$
- is the path predicate φ_{σ} satisfiable?
- if SAT(s) then add a new pair $\langle s, \sigma \rangle$ into TS
- loop until no more paths to cover

[smt solver]


Dynamic Symbolic Execution (2)

input : a program P

output : a test suite TS covering all feasible paths of $Paths^{\leq k}(P)$

- pick an uncovered path $\sigma \in Paths^{\leq k}(P)$
- is the path predicate φ_{σ} satisfiable?
- if SAT(s) then add a new pair $\langle s, \sigma \rangle$ into TS
- loop until no more paths to cover

[smt solver]




Dynamic Symbolic Execution (2)

input : a program P

output : a test suite TS covering all feasible paths of $Paths^{\leq k}(P)$

- pick an uncovered path $\sigma \in Paths^{\leq k}(P)$
- is the path predicate φ_{σ} satisfiable?
- if SAT(s) then add a new pair $\langle s, \sigma \rangle$ into TS
- loop until no more paths to cover

[smt solver]




Dynamic Symbolic Execution (2)

input : a program P

output : a test suite TS covering all feasible paths of $Paths^{\leq k}(P)$

- pick an uncovered path $\sigma \in Paths^{\leq k}(P)$
- is the path predicate φ_{σ} satisfiable?
- if SAT(s) then add a new pair $\langle s, \sigma \rangle$ into TS
- loop until no more paths to cover

[smt solver]




Dynamic Symbolic Execution (2)

input : a program P

output : a test suite TS covering all feasible paths of $Paths^{\leq k}(P)$

- pick an uncovered path $\sigma \in Paths^{\leq k}(P)$
- is the path predicate φ_{σ} satisfiable?
- if SAT(s) then add a new pair $\langle s, \sigma \rangle$ into TS
- loop until no more paths to cover

[smt solver]




Dynamic Symbolic Execution (2)

input : a program P

output : a test suite TS covering all feasible paths of $Paths^{\leq k}(P)$

- pick an uncovered path $\sigma \in Paths^{\leq k}(P)$
- is the path predicate φ_{σ} satisfiable?
- if SAT(s) then add a new pair $\langle s, \sigma \rangle$ into TS
- loop until no more paths to cover

[smt solver]




Dynamic Symbolic Execution (2)

input : a program P

output : a test suite TS covering all feasible paths of $Paths^{\leq k}(P)$

- pick an uncovered path $\sigma \in Paths^{\leq k}(P)$
- is the path predicate φ_σ satisfiable?
- if SAT(s) then add a new pair $\langle s, \sigma \rangle$ into TS
- loop until no more paths to cover

[smt solver]



Dynamic Symbolic Execution

- ✓ very powerful approach to white-box test generation
- ✓ arguably one of the most wide-spread use of formal methods in “common software”

Dynamic Symbolic Execution

- ✓ very powerful approach to white-box test generation
 - ✓ arguably one of the most wide-spread use of formal methods in “common software”
 - ✗ lack of support for many coverage criteria
-

Dynamic Symbolic Execution

- ✓ very powerful approach to white-box test generation
- ✓ arguably one of the most wide-spread use of formal methods in “common software”
- ✗ lack of support for many coverage criteria

Challenge : extend DSE to a large class of coverage criteria

- well-known problem
- recent efforts in this direction through instrumentation
[Active Testing, Mutation DSE, Augmented DSE]
- limitations :
 - ▶ exponential explosion of the search space [APEX : 272x avg]
 - ▶ very implementation-centric mechanisms
 - ▶ unclear expressiveness

Labels : a well-defined specification mechanism for coverage criteria

- ▶ based on predicates, can easily encode a large class of criteria
- ▶ w.r.t related work : semantic view, more formal treatment

DSE* : an efficient integration of labels into DSE

- ▶ no exponential blowup of the search space
- ▶ can be added to DSE in a black-box manner

Implem. in `PATHCRAWLER`

- ▶ huge savings compared to existing approaches
- ▶ handles labels with a very low overhead (2x average, up to 7x)

[Bardin et al., ICST 2014, TAP 2014, ICST 2015]

- Introduction
- Simulation of coverage criteria by labels
- Efficient DSE for labels
- Experiments
- Conclusion


Given a program P , a **label** l is a pair (loc, φ) , where :

- φ is a well-defined predicate in P at location loc
- φ contains no side-effect expression


Basic definitions

- a test datum t **covers** l if $P(t)$ reaches loc and satisfies φ
- new criterion **LC** (label coverage) for annotated programs
- a criterion **C** **can be simulated by LC** if for any P , after adding “appropriate” labels in P , TS covers **C** \Leftrightarrow TS covers **LC**.


Goal : show the relative expressiveness of **LC**



Decision Coverage (**DC**)

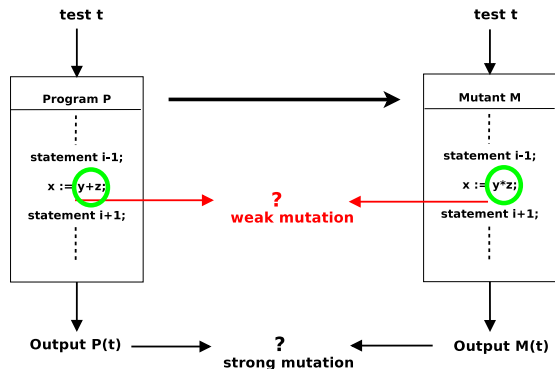


Condition Coverage (CC)



Multiple-Condition Coverage (**MCC**)

Weak Mutation (WM) testing in a nutshell



- mutant M = syntactic modification of program P
- **weakly covering** M = finding t such that $P(t) \neq M(t)$ just after the mutation

One label per mutant

Mutation inside a statement

- $lhs := e \quad \mapsto \quad lhs := e'$
 - ▶ add label : $e \neq e'$
- $lhs := e \quad \mapsto \quad lhs' := e$
 - ▶ add label : $\&lhs \neq \&lhs' \wedge (lhs \neq e \vee lhs' \neq e)$

Mutation inside a decision

- $if (cond) \quad \mapsto \quad if (cond')$
 - ▶ add label : $cond \oplus cond'$

Beware : no side-effect inside labels

Theorem

*The following coverage criteria can be simulated by **LC** : **IC**, **DC**, **FC**, **CC**, **MCC**, *Input Domain Partition*, *Run-Time Errors*.*


Theorem

*For any finite set O of side-effect free mutation operators, **WM** _{O} can be simulated by **LC**.*


Goals

- ✓ GOAL1 : generic specification mechanism for coverage criteria
- GOAL2 : efficient integration into DSE

- Introduction
- Simulation of coverage criteria by labels
- Efficient DSE for labels
- Experiments
- Conclusion



Covering label $I \Leftrightarrow$ Covering branch True




Covering label $l \Leftrightarrow$ Covering branch True

✓ sound & complete instrumentation w.r.t. LC

Direct instrumentation P' is not good enough

Direct instrumentation




Direct instrumentation P' is not good enough

Non-tightness 1

✗ P' has exponentially more paths than P

Direct instrumentation



Direct instrumentation P' is not good enough


Non-tightness 1

- ✗ P' has exponentially more paths than P

Non-tightness 2


- ✗ Paths in P' too complex
 - ▶ at each label, require to cover p or to cover $\neg p$
 - ▶ π' covers up to N labels

Direct instrumentation




Direct instrumentation P' is not good enough

Direct instrumentation




- ✓ sound & complete instrumentation w.r.t. LC
- ✗ dramatic overhead [theory & practice]

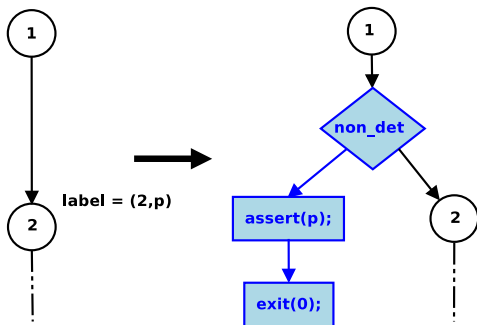


The DSE* algorithm

- Tight instrumentation P^* : totally prevents “complexification”
- Iterative Label Deletion : discards some redundant paths
- Both techniques can be implemented in a black-box manner




Covering label `l` \Leftrightarrow Covering `exit(0)`




Covering label `l` \Leftrightarrow Covering `exit(0)`

✓ sound & complete instrumentation w.r.t. LC


Direct instrumentation




Tight Instrumentation



Direct instrumentation




Tight Instrumentation

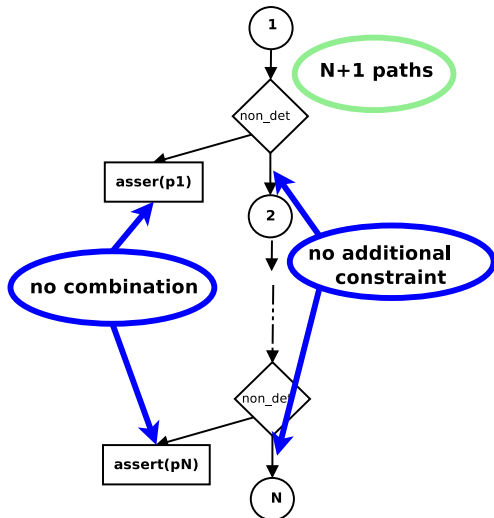


DSE* : Direct vs tight instrumentation, P' vs P^*

Direct instrumentation




Tight Instrumentation




DSE* : Direct vs tight instrumentation, P' vs P^*

Direct instrumentation





Tight Instrumentation




Tightness

- ✓ P^* has (only) linearly more paths than P
- ✓ paths in P^* are simple : covers ≤ 1 label

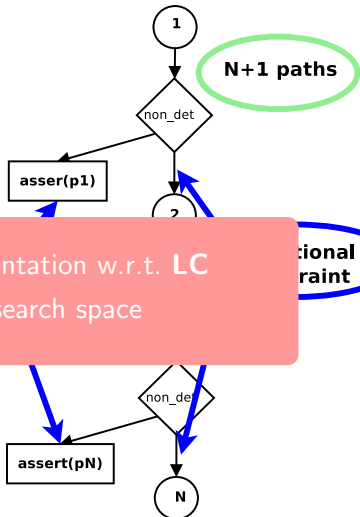


DSE* : Direct vs tight instrumentation, P' vs P^*

Direct instrumentation



Tight Instrumentation



combi

- ✓ sound & complete instrumentation w.r.t. LC
- ✓ no complexification of the search space

itional
raint

Observations

- we need to cover each label only once
- yet, DSE explores paths of P^* ending in already-covered labels
- we burden DSE with “useless” paths w.r.t. **LC**

Observations

- we need to cover each label only once
- yet, DSE explores paths of P^* ending in already-covered labels
- we burden DSE with “useless” paths w.r.t. **LC**

Solution : Iterative Label Deletion

- keep a *covered/uncovered* status for each label
- symbolic execution ignores paths ending in a covered label
- dynamic execution updates the status [truly requires DSE]

Implementation

- symbolic part : a slight modification of P^*
- dynamic part : a slight modification of P'

Observations

- we need to cover each label only once
- yet, DSE explores paths of P^* ending in already-covered labels
- we burden DSE with “useless” paths w.r.t. **LC**

Solution : Iterative Label Deletion

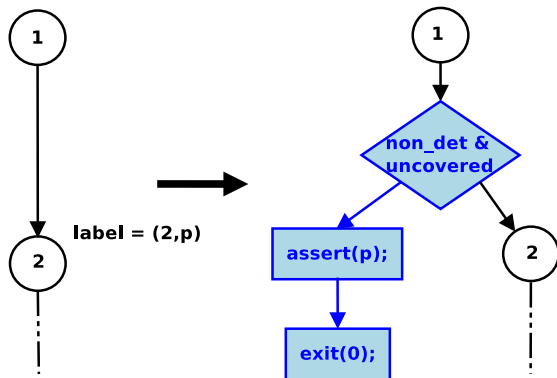
- keep a *covered/uncovered* status for each label
- symbolic execution ignores paths ending in a covered label
- dynamic execution updates the status [truly requires DSE]

Implementation

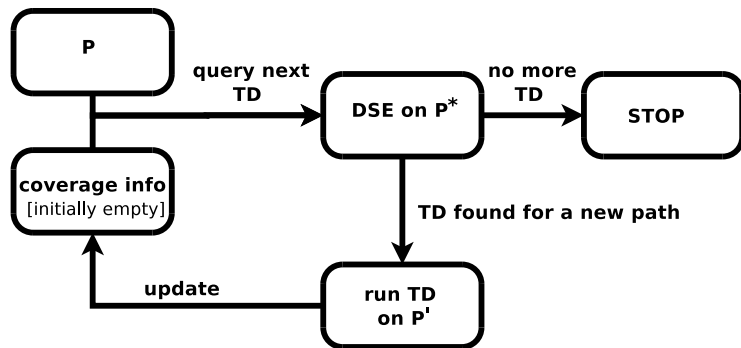
- symbolic part : a slight modification of P^*
- dynamic part : a slight modification of P'

Iterative Label Deletion is relatively complete w.r.t. **LC**

DSE* : Iterative Label Deletion (2)



DSE* : Iterative Label Deletion (3)



The DSE* algorithm

- Tight instrumentation P^* : totally prevents “complexification”
- Iterative Label Deletion : discards some redundant paths
- Both techniques can be implemented in black-box

The DSE* algorithm

- **Goals**

- ✓ GOAL1 : generic specification mechanism for coverage criteria
- ✓ GOAL2 : efficient integration into DSE

- Introduction
- Simulation of coverage criteria by labels
- Efficient DSE for labels
- Experiments
- Conclusion

Implementation

- inside PATHCRAWLER
- follows DSE*
- search heuristics : “label-first DFS”
- run in deterministic mode

Goal of experiments

- evaluate DSE* versus DSE'
- evaluate overhead of handling labels

Benchmark programs

- 12 programs taken from standard DSE benchmarks (Siemens, Verisec, MediaBench) [beware : small programs]
- 3 coverage criteria : **CC**, **MCC**, **WM**
[uncoverable labels not discarded]

Results

- DSE' : 4 timeouts (TO), max overhead 122x [excluding TO]
 - DSE* : no TO, max overhead 7x (average : 2.4x)
 - on one example, 94s instead of a TO [1h30]
 - DSE* achieves very high **LC**-coverage [$> 90\%$ on 28/36]
 - after a static analysis step for detection of uncoverable labels, it becomes even higher [$> 99\%$]
-

Results

- DSE' : 4 timeouts (TO), max overhead 122x [excluding TO]
- DSE* : no TO, max overhead 7x (average : 2.4x)
- on one example, 94s instead of a TO [1h30]
- DSE* achieves very high LC-coverage [$> 90\%$ on 28/36]
- after a static analysis step for detection of uncoverable labels, it becomes even higher [$> 99\%$]

Results

- DSE' : 4 timeouts (TO), max overhead 122x [excluding TO]
- DSE* : no TO, max overhead 7x (average : 2.4x)
- on one example, 94s instead of a TO [1h30]
- DSE* achieves very high **LC**-coverage [$> 90\%$ on 28/36]
- after a static analysis step for detection of uncoverable labels, it becomes even higher [$> 99\%$]

Results

- DSE' : 4 timeouts (TO), max overhead 122x [excluding TO]
- DSE* : no TO, max overhead 7x (average : 2.4x)
- on one example, 94s instead of a TO [1h30]
- DSE* achieves very high **LC**-coverage [$> 90\%$ on 28/36]
- after a static analysis step for detection of uncoverable labels, it becomes even higher [$> 99\%$]

Conclusion


- DSE* performs significantly better than DSE'
- The overhead of handling labels is kept reasonable
- still room for improvement

A few detailed results

			DSE	DSE'	DSE*
utf8-5 108 loc	wm 84 /	#paths time cover	680 2s	11,111 40s 82/84	743 8.1s 82/84
utf8-7 108 loc	wm 84 /	#paths time cover	3,069 5.8s	81,133 576s 82/84	3,265 35s 82/84
tcas 124 loc	wm 111 /	#paths time cover	4,420 5.6s	300,213 662s 101/111	6,014 27s 101/111
replace 100 loc	wm 79 /	#paths time cover	866 2s	87,498 245s 70/79	2,347 14s 70/79
get_tag-6 240 loc	cc 20 /	#paths time cover	76,456 3,011s	TO	76,468 1,512s 20/20
	wm 47 /	#paths time cover	76,456 3,011s	TO	76,481 1,463s 44/47
gd-5	wm 63 /	#paths time cover	14,516 50s	TO	14,607 94s 62/63
gd-6	wm 63 /	#paths time cover	107,410 3,740s	TO	107,521 2,232s 63/63

Implementation on top of FRAMA-C

- FRAMA-C is a toolset for analysis of C programs
 - ▶ an extensible, open-source, plugin-oriented platform
 - ▶ offers value analysis (VA), weakest precondition (WP), specification language ACSL,...
- LTEST is open-source except test generation
 - ▶ based on the PATHCRAWLER test generation tool



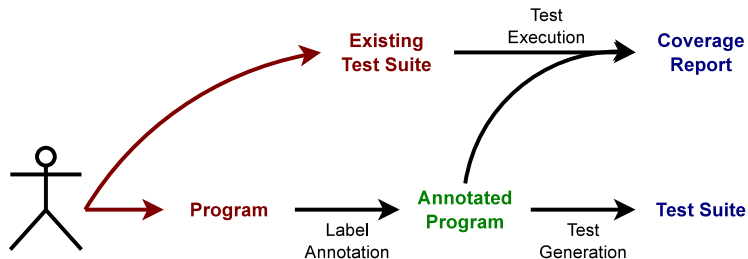
Supported criteria

- DC, CC, MCC
- FC, IDC, WM

Encoded with labels [ICST 2014]

- treated in a unified way
- easy to add new criteria

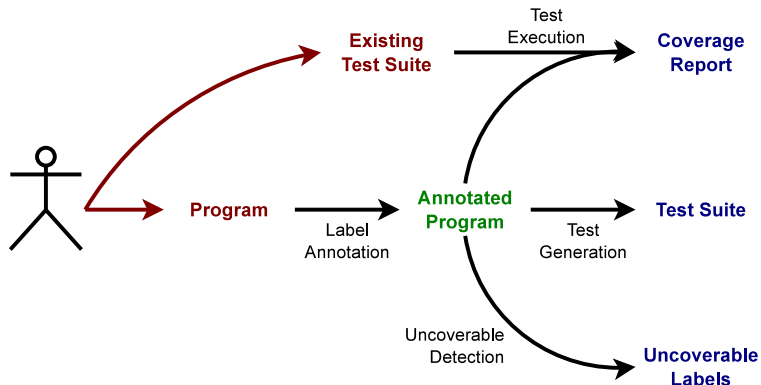




DSE* procedure [ICST 2014]

- DSE with native support for labels
- extension of PATHCRAWLER

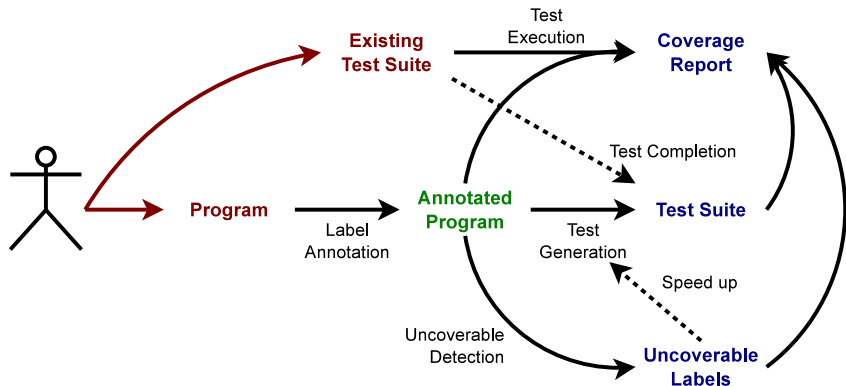
The LTEST toolset for labels [TAP 14]



Uses static analyzers from FRAMA-C

- sound detection of uncoverable labels

The LTEST toolset for labels [TAP 14]



Uses static analyzers from FRAMA-C

- sound detection of uncoverable labels

Service cooperation

- share label statuses
- Covered, Infeasible, ?

- Introduction
- Simulation of coverage criteria by labels
- Efficient DSE for labels
- Experiments
- Conclusion

Goal = extend DSE to a large class of coverage criteria

Results

- Labels : a well-defined and expressive specification mechanism for coverage criteria
 - DSE* : an efficient integration of labels into DSE
 - ▶ no exponential blowup of the search space
 - ▶ only a low overhead [huge savings w.r.t. related work]
-

Goal = extend DSE to a large class of coverage criteria

Results

- Labels : a well-defined and expressive specification mechanism for coverage criteria
- DSE* : an efficient integration of labels into DSE
 - ▶ no exponential blowup of the search space
 - ▶ only a low overhead [huge savings w.r.t. related work]

Dynamic Symbolic Execution [dart, cute, exe, sage, pex, klee, ...]

- ✓ very powerful approach to (white box) test generation
- ✓ arguably one of the most wide-spread use of formal methods in “common software”

Goal = extend DSE to a large class of coverage criteria

Results

- Labels : a well-defined and expressive specification mechanism for coverage criteria
- DSE* : an efficient integration of labels into DSE
 - ▶ no exponential blowup of the search space
 - ▶ only a low overhead [huge savings w.r.t. related work]

Dynamic Symbolic Execution [dart, cute, exe, sage, pex, klee, ...]

- ✓ very powerful approach to (white box) test generation
- ✓ arguably one of the most wide-spread use of formal methods in “common software”
- ✗ support only basic coverage criteria

Goal = extend DSE to a large class of coverage criteria

Results

- Labels : a well-defined and expressive specification mechanism for coverage criteria
- DSE* : an efficient integration of labels into DSE
 - ▶ no exponential blowup of the search space
 - ▶ only a low overhead [huge savings w.r.t. related work]

Dynamic Symbolic Execution [dart, cute, exe, sage, pex, klee, ...]

- ✓ very powerful approach to (white box) test generation
- ✓ arguably one of the most wide-spread use of formal methods in “common software”
- ✓ can be efficiently extended to a large class of coverage criteria