

# Running Symbolic Execution Forever

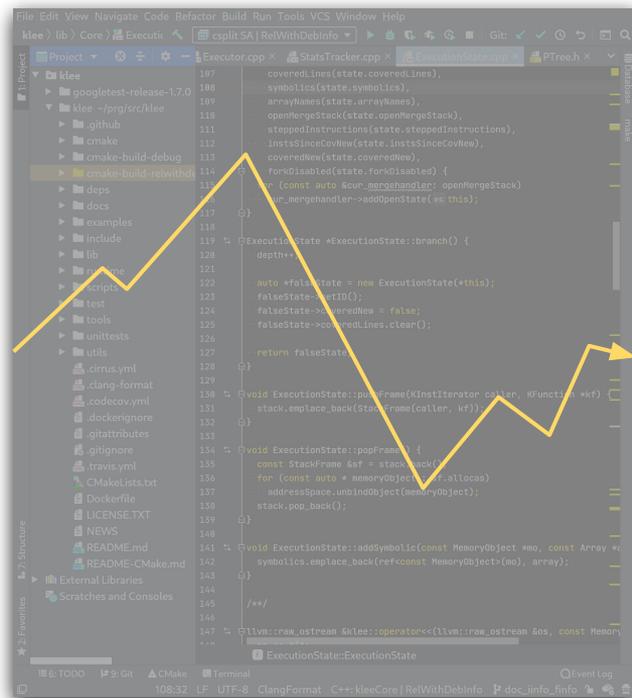
Frank Busse · Martin Nowack · Cristian Cadar  
Imperial College London

ISSTA 2020, 18-22 July, Virtual Conference, USA

# Concrete vs. Symbolic Execution

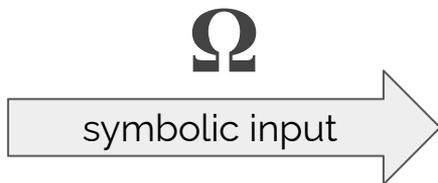


concrete input

A screenshot of a code editor showing C++ code for symbolic execution. The code includes a class `ExecutionState` with methods like `coveredLines`, `addSymbolic`, and `addOpenState`. A yellow arrow points from the hardware input to the code, and another yellow arrow points from the code to the output.

concrete output

# Concrete vs. Symbolic Execution

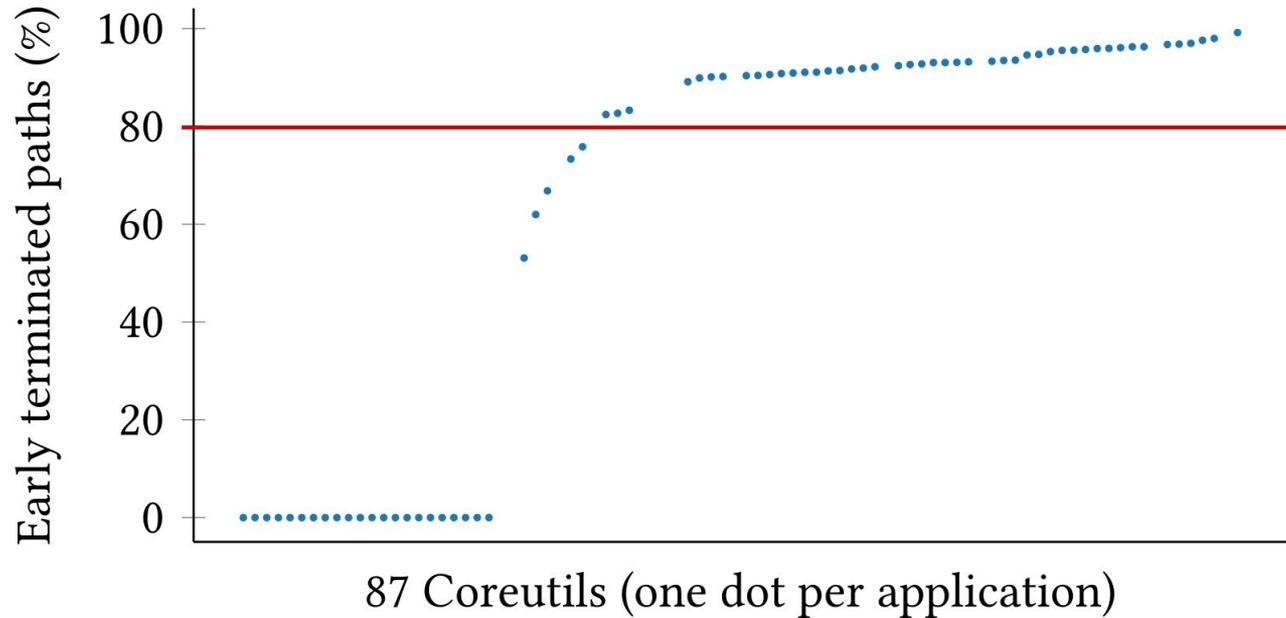


```
107 coveredLines(state.coveredLines),
108 symbolics(state.symbolics),
109 arrayNames(state.arrayNames),
110 openMergeStack(state.openMergeStack),
111 steppedInstructions(state.steppedInstructions),
112 instsSinceCovNew(state.instsSinceCovNew),
113 coverage(state.coverage),
114 forDisabled(state.forDisabled) {
115     if (const auto &cur_mergeHandler: openMergeStack)
116         mergeHandler->addOpenState(*this);
117 }
118
119 ExecutionState *ExecutionState::branch() {
120     depth++;
121     auto *falseState = new ExecutionState(*this);
122     falseState->id();
123     falseState->covNew = false;
124     falseState->coveredLines.clear();
125     return falseState;
126 }
127
128 void ExecutionState::pushFrame(KInstIterator &it, KFunction *kf) {
129     stack.emplace_back(StateFrame(caller, kf));
130 }
131
132 void ExecutionState::popFrame() {
133     const StackFrame &sf = stack.back();
134     sf(const auto &memoryObject: sf.allLocals)
135         addressSpace.unbindObject(memoryObject);
136     stack.pop_back();
137 }
138
139 void ExecutionState::pushSymbolics(const MemoryObject *mo, const Array *
140     symbolics) {
141     stack.emplace_back(ref(const MemoryObject *mo), array);
142 }
143
144 /**
145  *
146  */
147 llvm::raw_ostream &&ExecutionState::llvm::raw_ostream && const Memory
```

- high-coverage test cases
- crashing inputs

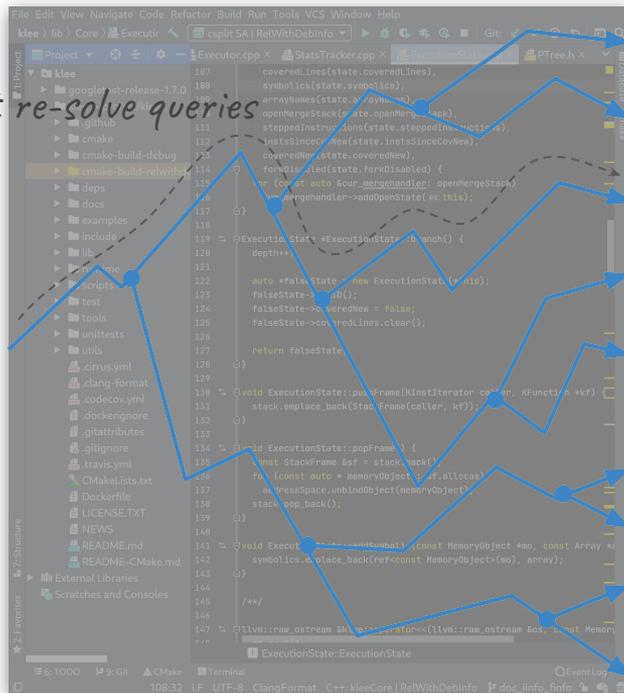


# Early Termination (Memory Pressure)



# Motivation

*don't re-solve queries*



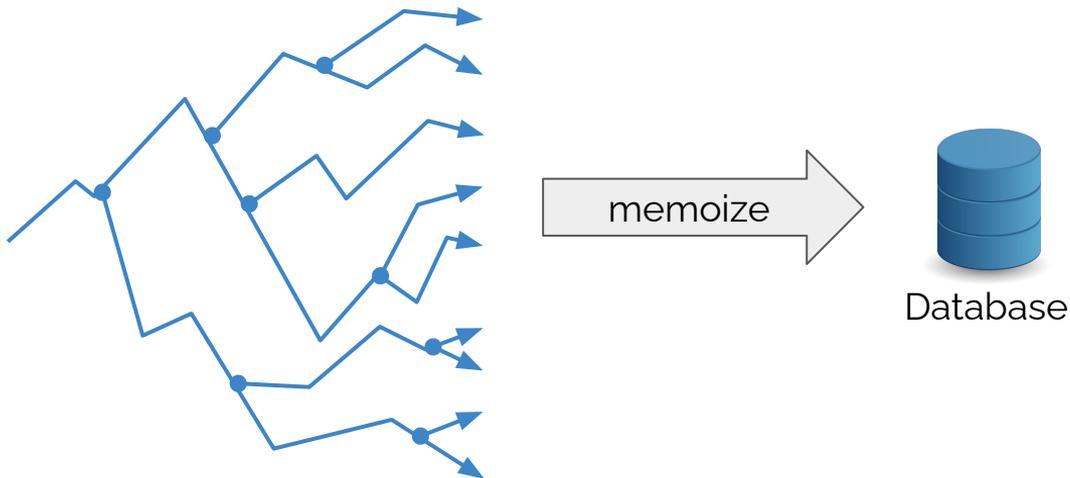
*large subtree*



*don't re-explore paths*

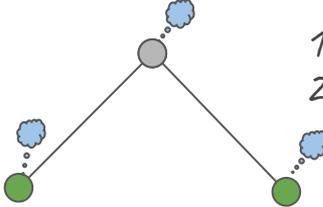
# Memoization

- trade time for space
- **store solver results** as metadata in execution tree nodes
- **persist tree to disk**
- **re-use results on re-execution**



# Memoization

current execution tree

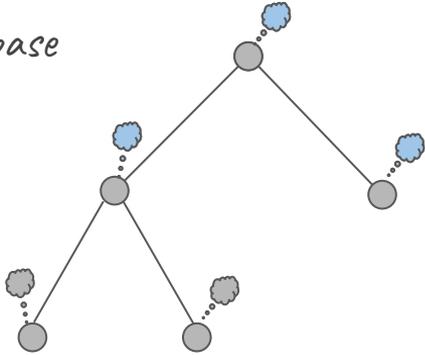


- 1. load metadata from database
- 2. re-use solver results

- 3. branch
- 4. load metadata
- 5. free metadata in parent

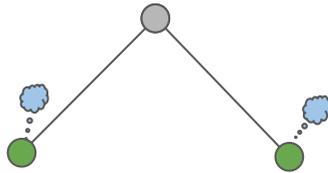


stored execution tree



# Memoization

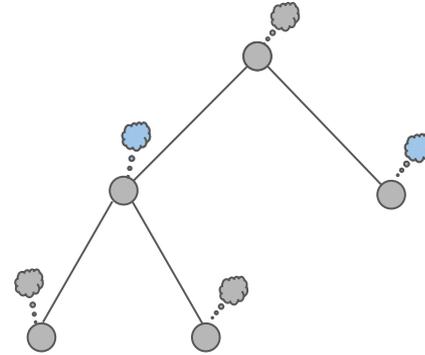
*current execution tree*



*path progresses beyond  
memoized data*

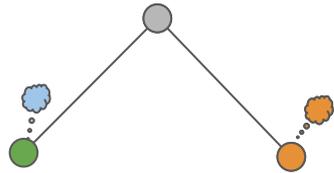


*stored execution tree*



# Memoization

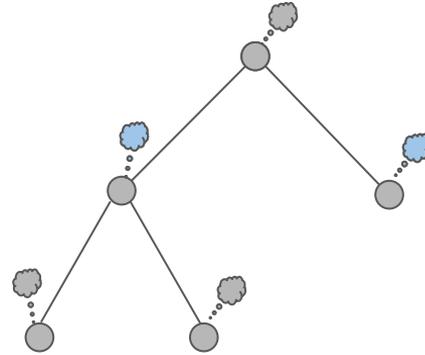
*current execution tree*



*path switches to  
recording mode*

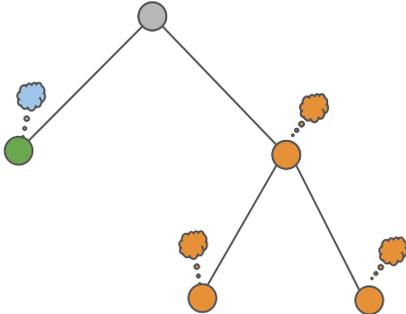


*stored execution tree*

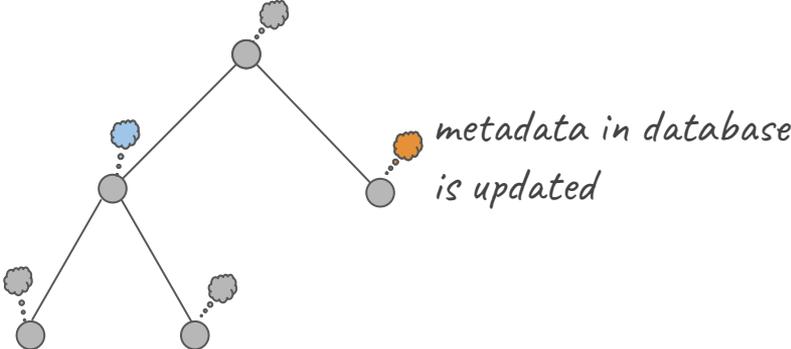


# Memoization

*current execution tree*

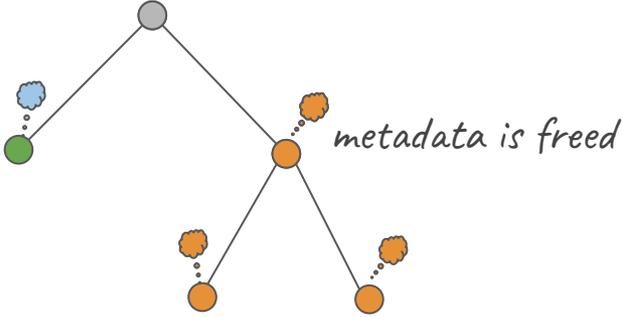


*stored execution tree*

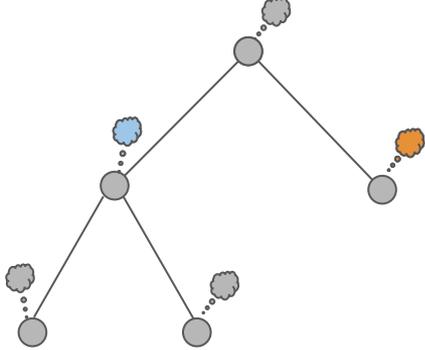


# Memoization

*current execution tree*

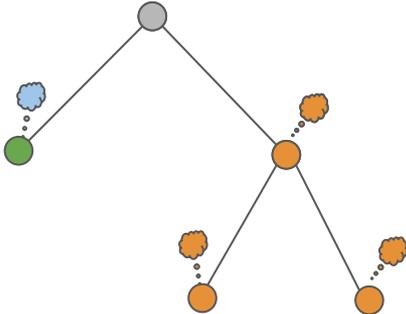


*stored execution tree*

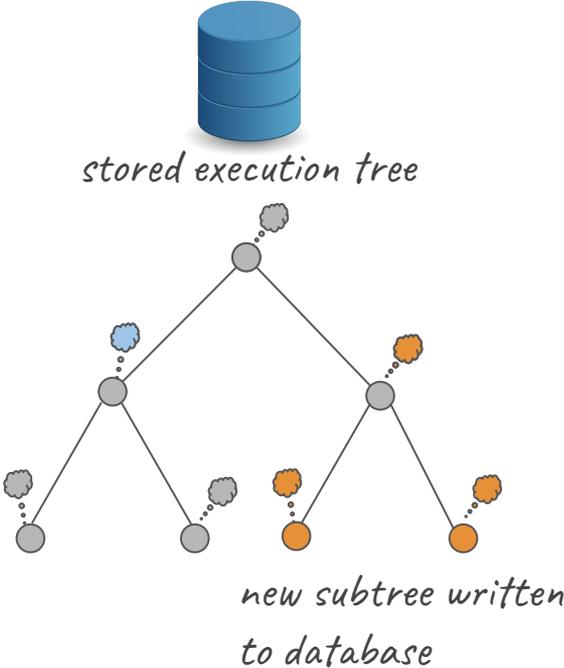


# Memoization

*current execution tree*

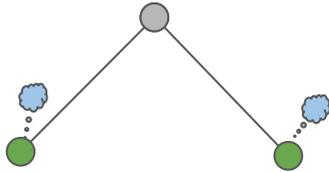


*stored execution tree*

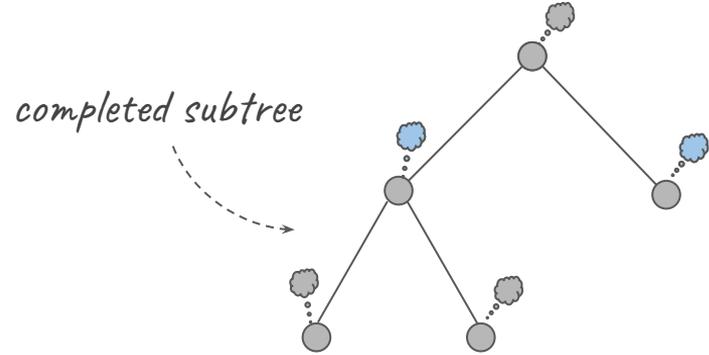


# Path Pruning

*current execution tree*

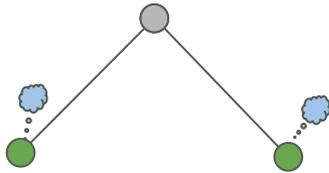


  
*stored execution tree*



# Path Pruning

*current execution tree*

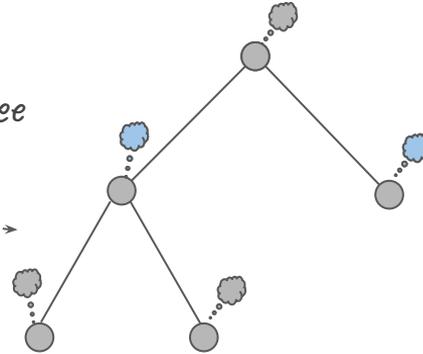


*on branch completeness  
immediately detected*



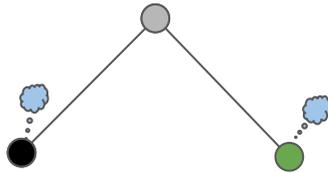
*stored execution tree*

*completed subtree*



# Path Pruning

*current execution tree*

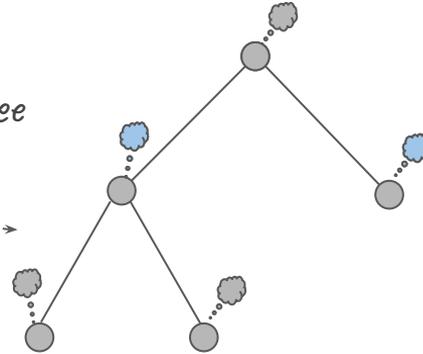


*path gets terminated*



*stored execution tree*

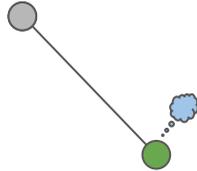
*completed subtree*



# Path Pruning

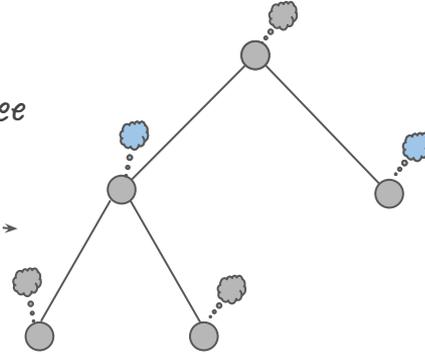
*current execution tree*

*... and removed from tree*



*stored execution tree*

*completed subtree*



# Memoized Symbolic Execution

Guowei Yang  
University of Texas at Austin  
Austin, TX 78712, USA  
guowei yang@utexas.edu

Corina S. Păsăreanu  
Carnegie Mellon Silicon Valley  
NASA Ames, M/S 269-2  
Moffett Field, CA 94035, USA  
corina.s.pasareanu@nasa.gov

Sarfraz Khurshid  
University of Texas at Austin  
Austin, TX 78712, USA  
khurshid@ece.utexas.edu

## ABSTRACT

This paper introduces *memoized symbolic execution (Memoise)*, a new approach for more efficient application of forward symbolic execution, which is a well-studied technique for systematic exploration of program behaviors based on bounded execution paths. Our key insight is that application of symbolic execution often requires several successive runs of the technique on largely similar underlying problems, e.g., running it once to check a program to find a bug, fixing the bug, and running it again to check the modified program. Memoise introduces a *trie*-based data structure that stores the key elements of a run of symbolic execution. Maintenance of the trie during successive runs allows re-use of previously computed results of symbolic execution without the need for re-computing them as is traditionally done. Experiments using our prototype implementation of Memoise show the benefits it holds in various standard scenarios of using symbolic execution, e.g., with iterative deepening of exploration depth, to perform regression analysis, or to enhance coverage using heuristics.

## Categories and Subject Descriptors

D.2.5 [Software Engineering]: Testing and Debugging—*Symbolic execution*

Off-the-shelf constraint solvers are used to reason about the formulas to discard those paths whose conditions are unsatisfiable. In practice, the technique can be costly to apply due to its inherent high time and space complexity. There are two key factors that determine its cost: (1) the number of paths that need to be explored and (2) the cost of constraint solving.

Recent years have seen substantial advances in raw computation power and constraint solving technology [1], as well as in basic algorithmic approaches for symbolic execution [4, 25]. These advances have made symbolic execution applicable to a diverse class of programs and enable a range of analyses, including bug finding using automated test generation – a traditional application of this technique – as well as other novel applications, such as program equivalence checking [23], regression analysis [17], and continuous testing [27]. All these applications utilize the same path-based analysis that lies at the heart of symbolic execution. As such, their effectiveness is determined by the two factors that determine the cost of the symbolic execution, and at present, reducing the cost of symbolic execution remains a fundamental challenge.

This paper introduces *memoized symbolic execution (Memoise)*, a new approach that addresses both factors to enable more efficient applications of symbolic execution. Our key insight is that applying symbolic execution often requires several successive runs of

The second assumption maintains the correspondence of the executions of program paths across different runs of symbolic execution, and makes feasible the reuse of symbolic execution results. As long as the same search order is used during re-execution, the symbolic execution tree corresponding to the same program executions remain the same, and this assures the correctness of trie-guided symbolic execution. Merging is correct since the executions corresponding to the removed parts remain the same in re-execution and will yield to the same sub-trie, and thus the removed parts be brought back from the old trie.

*(same search strategy)*

*(loads complete tree)*

### 3.2.1 Node Marking

The first step in memoized execution is to *mark nodes of interest*. Specifically, we characterize parts of the old trie that may be updated using *candidate* nodes, which represent roots of sub-trees potentially updated during memoized execution. Given the candidate nodes, we mark nodes on paths that need re-execution – all nodes on any path from the trie root to a *candidate* node are marked (while the rest of the nodes remain unmarked). The exact classification of candidate nodes depends on the particular analysis that is performed. For example, for iterative deepening, the boundary nodes are the candidate nodes (e.g., *n9* in Figure 3). In regression analysis the nodes that are impacted by the program change are considered as candidate ones (the impacted nodes are found by an impact analysis as described in Section 4.1.2). The node marking is reset at the beginning of memoized analysis.

The second assumption maintains the correspondence of the executions of program paths across different runs of symbolic execution, and makes feasible the reuse of symbolic execution results. As

longer  
symbolic  
run  
symbolic  
results  
and  
be

**Table 1: Iterative Deepening Results**

Depth		Sym Exe at Depth A				Sym Exe at Depth B											
A	B	Time (ss)		Mem (MB)		States		#Solver calls		Time (ss)			Mem (MB)			Trie (MB)	
		Reg	ID	Reg	ID	Reg	ID	Reg	ID	Reg	ID-p	ID-c	Reg	ID-p	ID-c	ID-p	ID-c
24	25	16	21	305	419	349272	252952	335358	77312	20	24	22	242	474	474	18.8	13.4
29	30	34	60	246	260	644184	171784	629758	32256	37	36	27	214	500	486	35.4	9.3

(a) WBS Example

Depth		Sym Exe at Depth A				Sym Exe at Depth B											
A	B	Time (ss)		Mem (MB)		States		#Solver calls		Time (ss)			Mem (MB)			Trie (MB)	
		Reg	ID	Reg	ID	Reg	ID	Reg	ID	Reg	ID-p	ID-c	Reg	ID-p	ID-c	ID-p	ID-c
24	25	35	38	304	367	17103	16756	12252	2942	47	46	45	413	263	395	0.9	0.8
29	30	86	87	419	333	33273	15250	25684	1540	92	45	45	413	345	263	2.0	0.9
34	35	96	97	419	345	35359	1476	27636	18	102	9	9	292	404	243	2.1	0.1

(b) MerArbiter Example

Depth		Sym Exe at Depth A				Sym Exe at Depth B											
A	B	Time (ss)		Mem (MB)		States		#Solver calls		Time (ss)			Mem (MB)			Trie (MB)	
		Reg	ID	Reg	ID	Reg	ID	Reg	ID	Reg	ID-p	ID-c	Reg	ID-p	ID-c	ID-p	ID-c
9	10	127	125	425	352	674	647	255	121	195	76	77	421	343	296	0.03	0.03
11	12	538	549	413	414	2243	2113	2160	966	1033	490	483	390	316	420	0.12	0.11

(c) Apollo Example

*(short runtimes)*  
*No divergence detection.*

considered as candidate ones (the impacted nodes are found by an impact analysis as described in Section 4.1.2). The node marking is reset at the beginning of memoized analysis.

er-  
be  
es  
an-  
all  
ed  
ifi-  
nat  
ary  
on  
are

# Divergences

## Causes

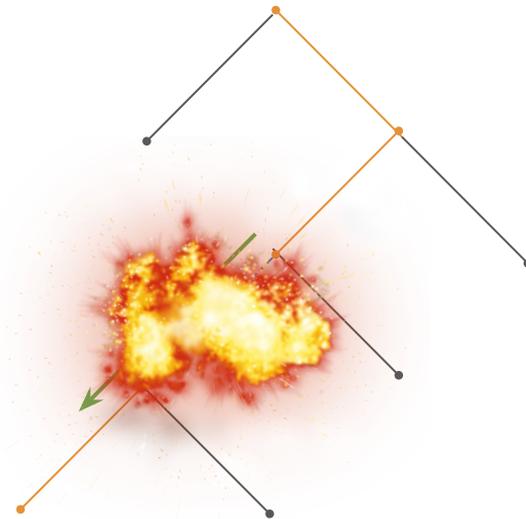
- changes in external environment (disk layout, date, environment variables)
- shared address space between execution states

## Problem

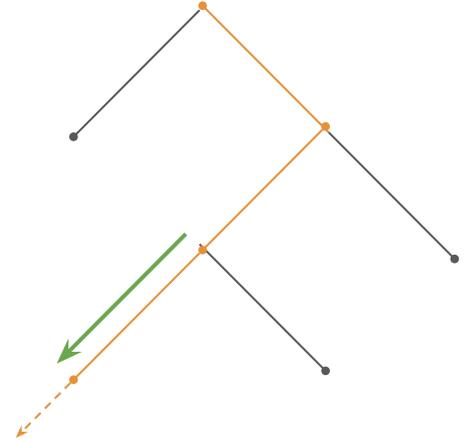
- exploration of infeasible paths
- false negatives



*current execution tree*



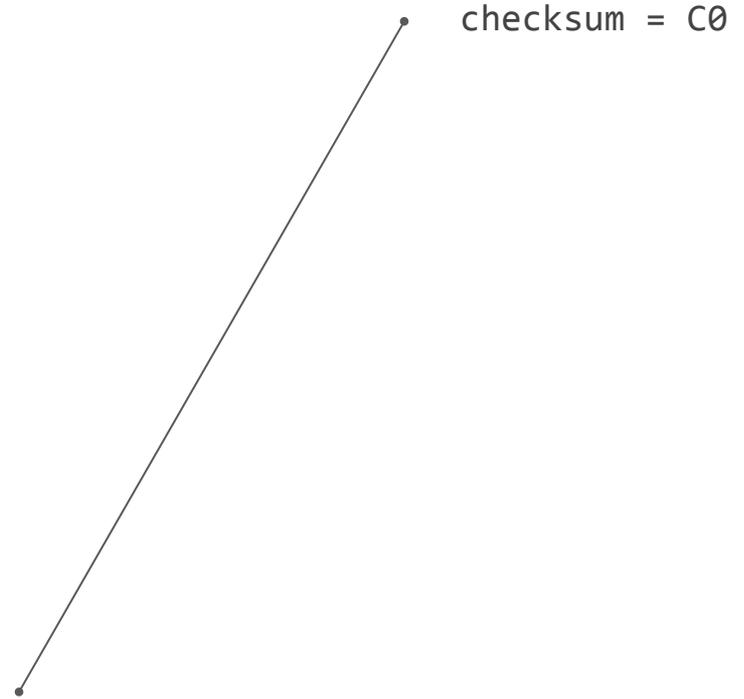
*stored execution tree*



# Divergence Detection

## Mitigation

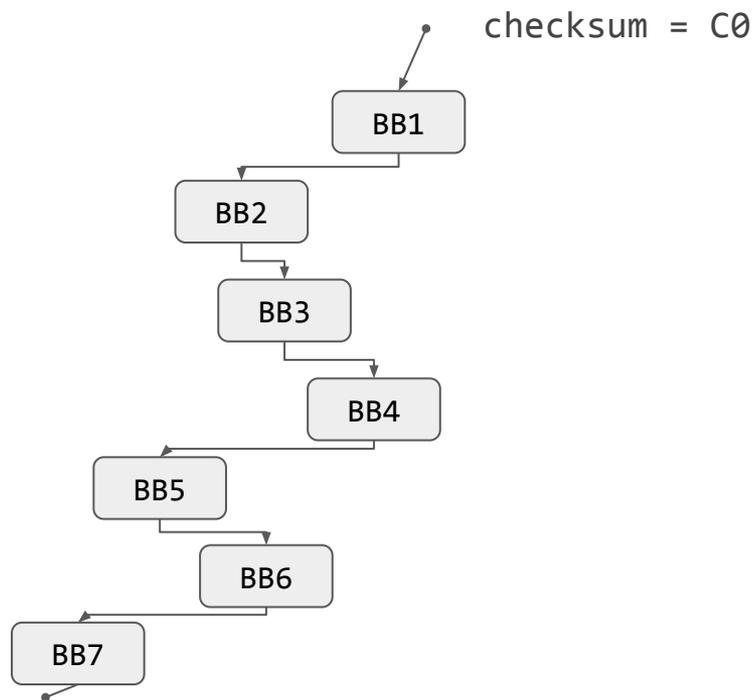
- checksum over sequence of basic blocks validated on each branch
- affected paths are reset



# Divergence Detection

## Mitigation

- checksum over sequence of basic blocks validated on each branch
- affected paths are reset

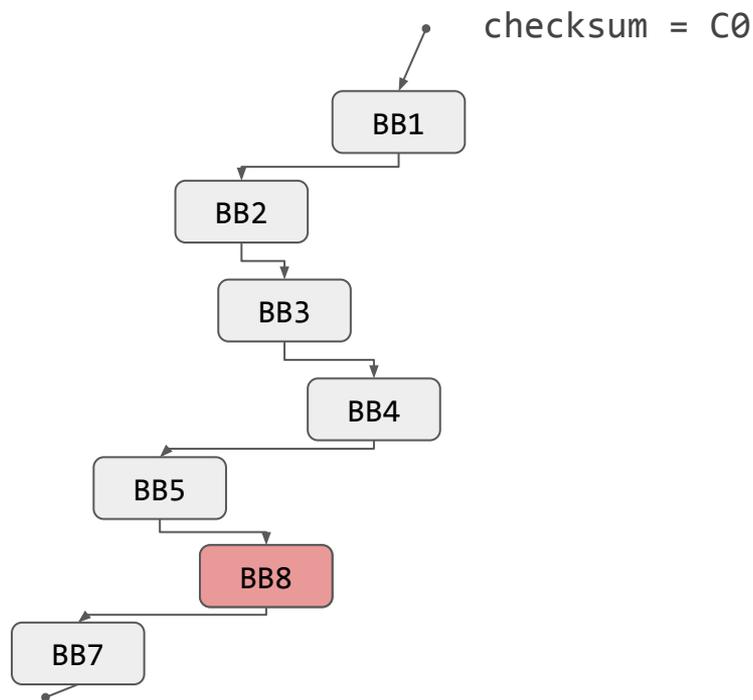


$$\text{checksum} = \text{hash}(\text{BB7}) \otimes \text{hash}(\text{BB6}) \otimes \dots \otimes \text{C0}$$

# Divergence Detection

## Mitigation

- checksum over sequence of basic blocks validated on each branch
- affected paths are reset

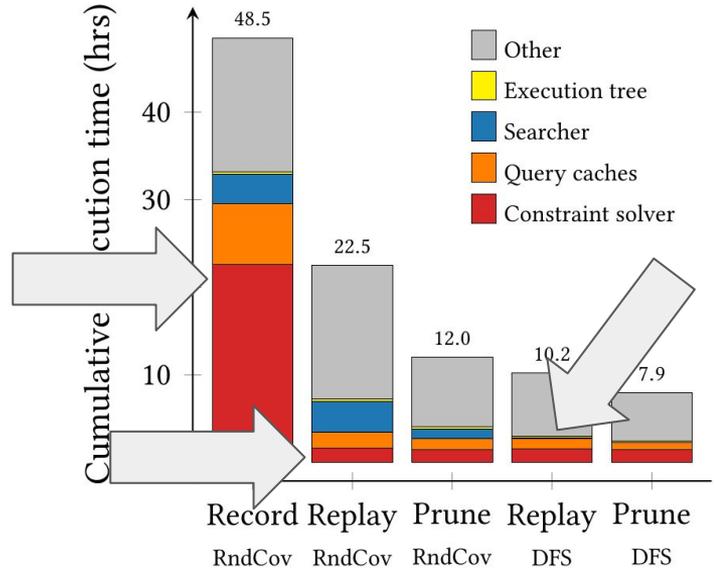
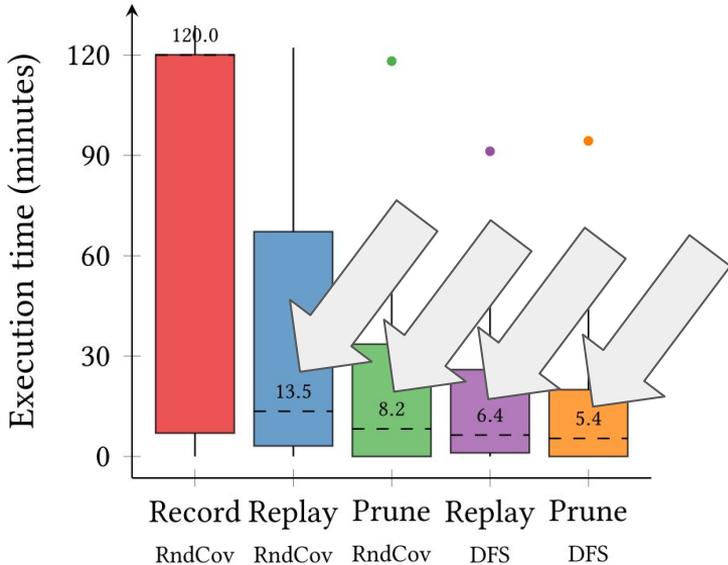


$$\text{checksum} \neq \text{hash}(\text{BB7}) \otimes \text{hash}(\text{BB8}) \otimes \dots \otimes C0$$

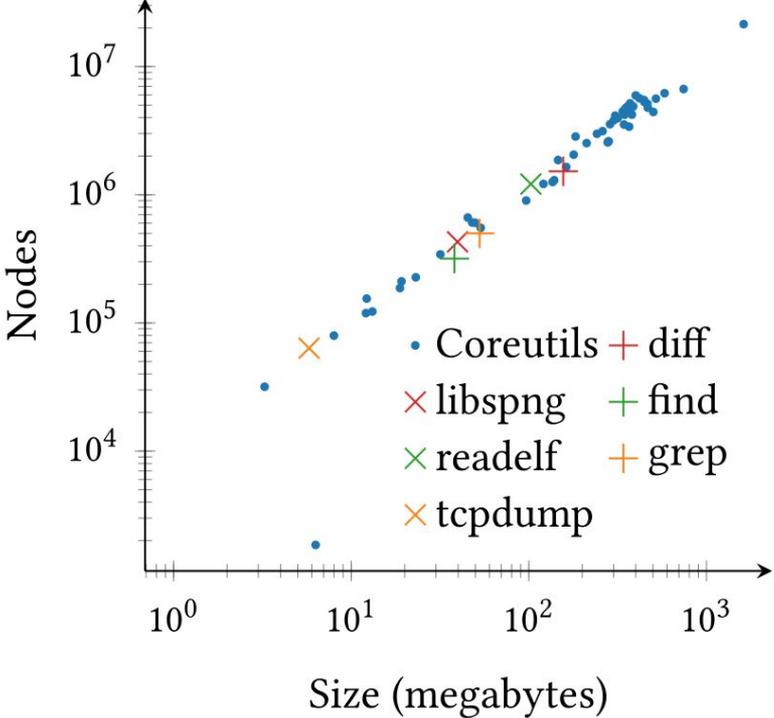
# Evaluation

- MoKlee is implemented on top of KLEE 1.4
- evaluated on 93 benchmarks:
  - readelf (Binutils)
  - 87 Coreutils
  - diff (Diffutils)
  - find (Findutils)
  - grep
  - libspng
  - tcpdump

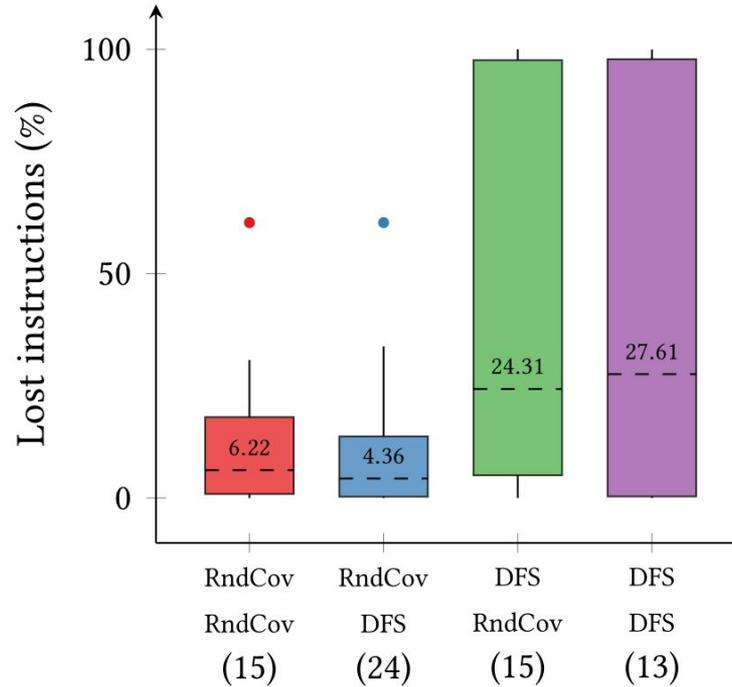
# Evaluation - Runtime



# Evaluation - Storage Size



# Evaluation - Divergences



# Evaluation - Long Running Symbolic Execution

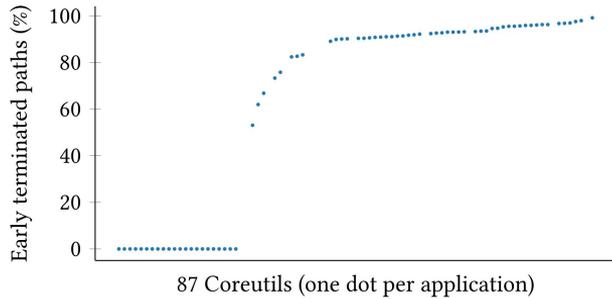
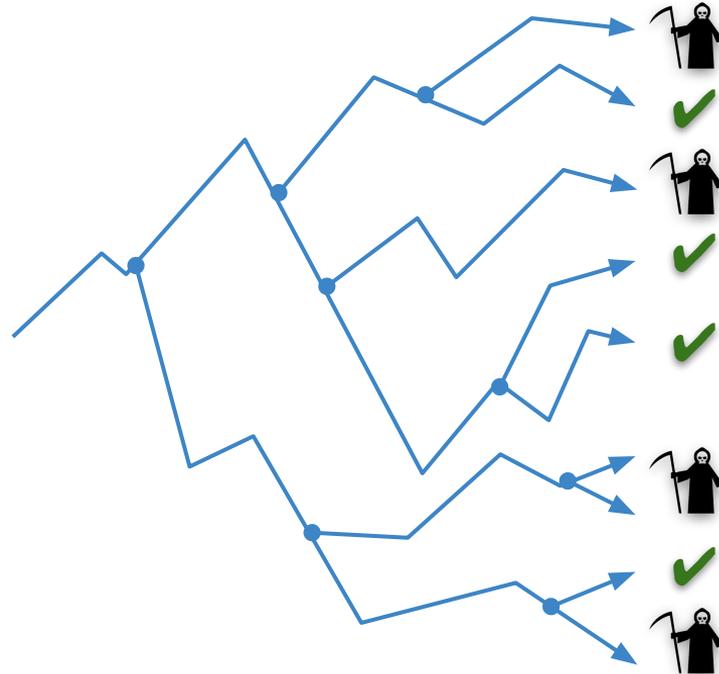
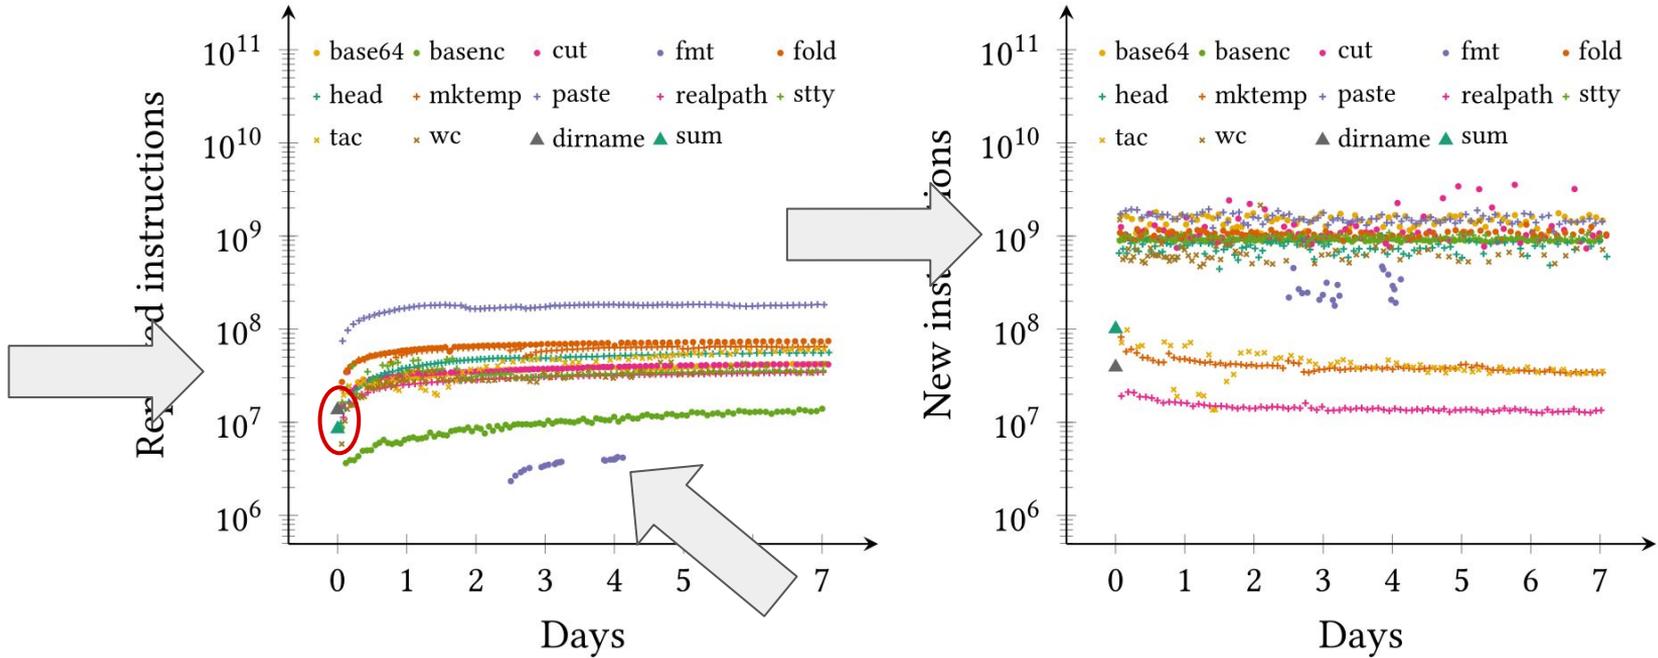


Figure 1: When running KLEE<sup>1</sup> on 87 *Coreutils* for 2 h each with the default search heuristic and memory limit (2 GB), most paths are terminated early due to memory pressure.

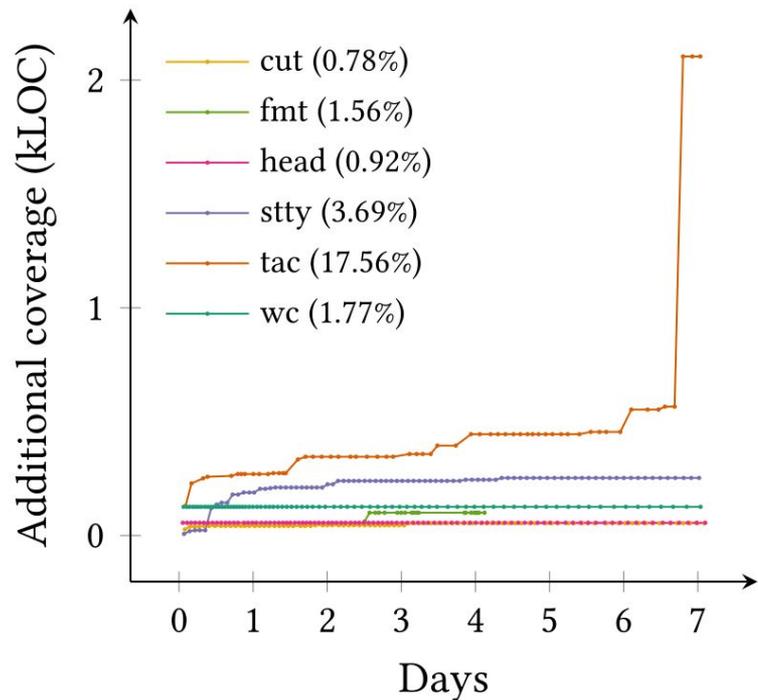


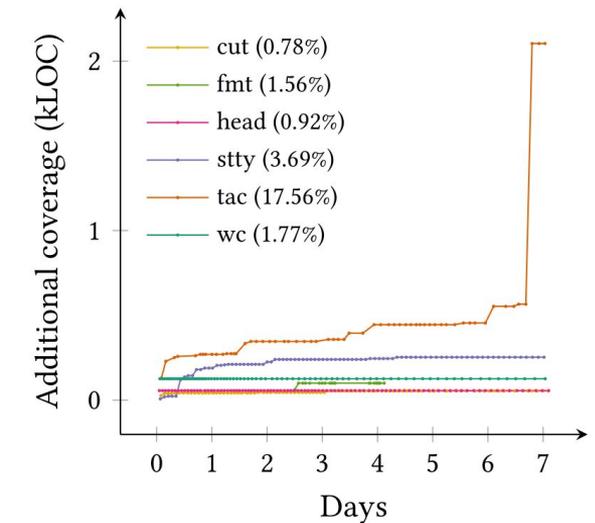
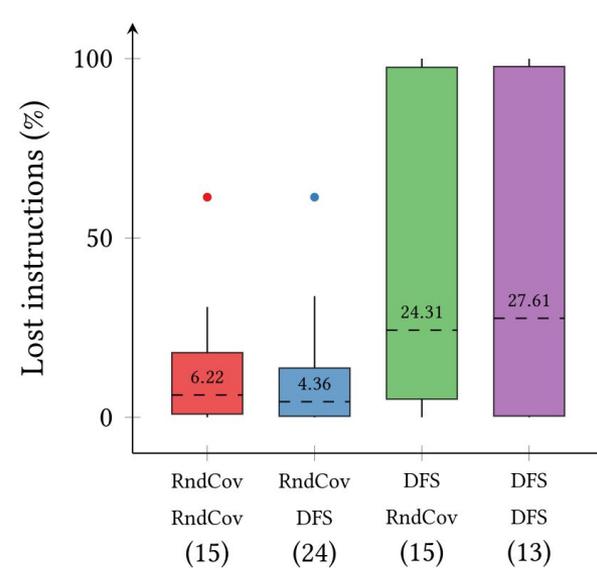
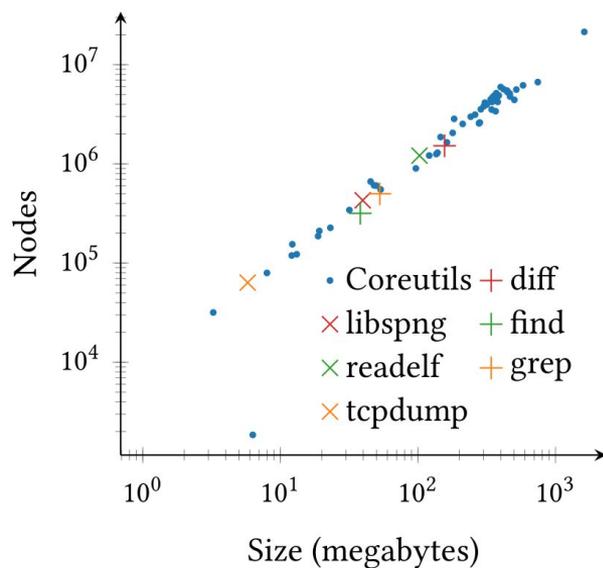
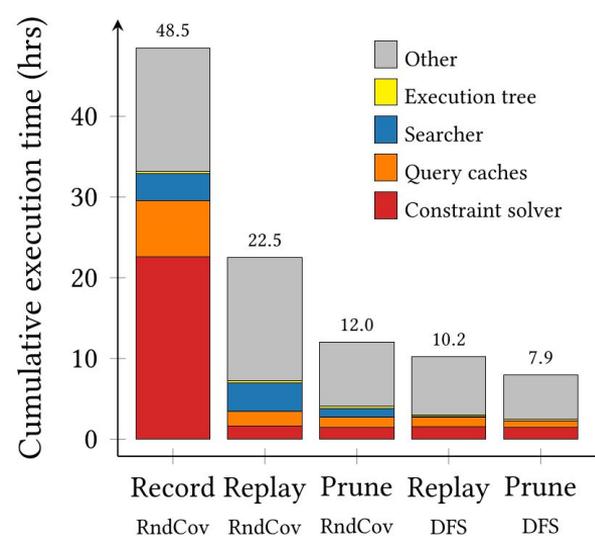
*14 applications terminate ran out of states before the 2h limit!*

# Evaluation - Long Running Symbolic Execution



# Evaluation - Long Running Symbolic Execution





MoKlee Artefact: <https://srg.doc.ic.ac.uk/projects/moklee/>

KLEE: <https://klee.github.io/>

2<sup>nd</sup> KLEE Workshop: 22-23 April 2021 in London