LEGION: Best-First Concolic Testing

Dongge Liu*

The University of Melbourne
School of Computing and Information Systems
Melbourne, Victoria, Australia
donggel@student.unimelb.edu.au

Toby Murray

The University of Melbourne
School of Computing and Information Systems
Melbourne, Victoria, Australia
toby.murray@unimelb.edu.au

Gidon Ernst

LMU Munich
Software and Computational Systems Lab
Munich, Bavaria, Germany
gidon.ernst@lmu.de

Benjamin I.P. Rubinstein

The University of Melbourne
School of Computing and Information Systems
Melbourne, Victoria, Australia
benjamin.rubinstein@unimelb.edu.au

ABSTRACT

Concolic execution and fuzzing are two complementary coverage-based testing techniques. How to achieve the best of both remains an open challenge. To address this research problem, we propose and evaluate LEGION. LEGION re-engineers the *Monte Carlo tree search (MCTS)* framework from the AI literature to treat automated test generation as a problem of sequential decision-making under uncertainty. Its best-first search strategy provides a principled way to learn the most promising program states to investigate at each search iteration, based on observed rewards from previous iterations. LEGION incorporates a form of directed fuzzing that we call *approximate path-preserving fuzzing (APPFUZZING)* to investigate program states selected by MCTS. APPFUZZING serves as the Monte Carlo simulation technique and is implemented by extending prior work on constrained sampling. We evaluate LEGION against competitors on 2531 benchmarks from the coverage category of Test-Comp 2020, as well as measuring its sensitivity to hyperparameters, demonstrating its effectiveness on a wide variety of input programs.

KEYWORDS

Concolic execution, constrained fuzzing, Monte Carlo tree search

ACM Reference Format:

Dongge Liu, Gidon Ernst, Toby Murray, and Benjamin I.P. Rubinstein. 2020. LEGION: Best-First Concolic Testing. In *35th IEEE/ACM International Conference on Automated Software Engineering (ASE '20)*, September 21–25, 2020, Virtual Event, Australia. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3324884.3416629>

*This research was supported by Data61 under the Defence Science and Technology Group's Next Generation Technologies Program.

This work is licensed under a Creative Commons Attribution International 4.0 License.

ASE '20, September 21–25, 2020, Virtual Event, Australia

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6768-4/20/09.

<https://doi.org/10.1145/3324884.3416629>

1 INTRODUCTION

Theseus killed Minotauros in the furthest section of the labyrinth and then made his way out again by pulling himself along the thread. —Pseudo-Apollodorus, *Bibliotheca* E1. 7 - 1. 9 trans. Aldrich

The complexity of modern software programs are like labyrinths for software testers to wander: their program states and execution paths form a confusing set of connecting rooms and paths. Like the Minotaur, faults often hide deep inside. One might guess at a fault's possible location via static analysis, but in order to slay it Theseus needs to know the path to it for sure, and the software tester needs to know which input will trigger it.

In the myth of Theseus, the hero king finds Minotauros by accurately tracing past paths with a ball of thread, allowing him to learn and estimate the maze structure. We argue that the very same tricks, namely recording exact concrete execution traces and applying machine learning to estimate software structure and guide its exploration, can also benefit coverage-based testing.

The focus of this paper is the quest of coverage-based testing, which is to cover as many paths in as little time as possible, delegating Minotaur detection to separate tools (e.g. AddressSanitiser [25], UBSan [23], Valgrind [19], Purify [21]). Traditional methods for coverage-based testing have been dominated by the two complementary approaches of concolic execution (as exemplified by DART [13] and SAGE [14]) and coverage-guided greybox fuzzing (as exemplified by libFuzzer [24], AFL [36], its various extensions such as AFLFast [6], AFLGo [5], CollaAFL [12], Angora [10]).

Continuing the mythological metaphor, with concolic execution one spends a long time rigorously planning each path through the maze via constraint solving, to make the correct turn at each branching point and ensure that no path will ever be repeated. However, such computation is expensive and, for most modern software, the maze is so large that repeating it for every path is infeasible.

In contrast, a coverage-guided fuzzer like AFL blindly scurries around the maze, neither spending much time on planning nor accurately memorising the paths and structure traversed. Thus much time is inevitably spent unnecessarily repeating obvious execution paths.

Observing the complementary nature of these two methods, our research aims to generalise them with Theseus's strategy. Our

tool LEGION¹ traces observed execution paths to estimate the maze structure. Then it identifies the most promising location to explore next, plans its path to reach that location, and applies a form of directed fuzzing to explore extensions of the path to that location. LEGION precisely traces each concrete execution (i.e. fuzzing run) and gathers statistics to refine its knowledge of the maze to inform its decisions about where to explore next.

Statistics gathering is central to LEGION’s approach, unlike traditional fuzzers like AFL that eschew gathering detailed statistics to save time. Instead, LEGION aims to harness the power of modern machine learning algorithms, informed by detailed execution traces as in other contemporary tools [1].

Real-life coverage testing is more complicated than the Theseus myth, as it requires a universal strategy that can adjust itself according to different program structures: i.e. to fuzz the program parts that are more suitable to fuzz, and favour concolic execution elsewhere. However, how to determine the best balance between these two strategies remains an open question.

To address this challenge, LEGION adopts the *Monte Carlo tree search* (MCTS) algorithm from machine learning [17]. MCTS has proven to work well in complex games like Go [27], multiplayer board games [29, 32] and poker [30], as it can adapt to the game it is playing via iterations of simulations and reward analysis. Specifically, MCTS learns a policy by successively simulating different plays and tracking the rewards obtained during each. In LEGION, plays correspond to concrete execution (directed fuzzing), while rewards correspond to increased coverage (discovering new execution paths). More importantly, MCTS’s guiding principle of *optimism in the face of uncertainty* is appropriate for exploring a maze with an unknown structure, randomness, and large branching factors where rigorously analysing every detail is infeasible. Instead, MCTS balances *exploitation* of the branches that appear to be most rewarding based on past experience, against *exploration* of less-well understood parts of the maze where rewards (path discovery) are less certain.

With the two tricks of Theseus, LEGION provides a principled framework to generalise and harness the complementary strengths of concolic execution and fuzzing. In doing so, it sheds new light on the key factors for efficient and effective input search strategies for coverage-based testing, on diverse software structures.

We term LEGION’s directed fuzzing approach *approximate path-preserving fuzzing* (APPFUZZING), and are inspired by QUICKSAMPLER [11]. APPFUZZING aims to fuzz a specific location of the search space by launching several binary executions that, with high probability, follow the same path to that location (but might take different paths afterwards). It is described in Section 4.2.

LEGION treats coverage-based testing as progressively exploring a large space with uncertainty, by iteratively sampling inputs using cheap but less accurate APPFUZZING that targets specific program states selected by MCTS. Symbolic state information (i.e. a path constraint) is required to seed APPFUZZING for each program state. However, LEGION avoids unnecessary symbolic execution (and constraint solving) by performing symbolic execution *lazily*: specifically, LEGION computes symbolic successor states only for

¹The name is a homage to the Marvel fictional character who changes personalities for different needs. Our strategy can adjust its exploration preference under different metrics.

program states deemed promising (i.e. those for which existing statistics indicate are worthy of investigation).

LEGION’s adoption of MCTS is designed to ensure that computational power is used to explore the most beneficial program locations, as determined by the score function. In LEGION scores are evaluated according to a modularised reward heuristic, with interchangeable coverage metrics as discussed in Section 5.1.

Our contributions are:

- We propose a variation of Monte Carlo tree search (MCTS) that maintains a balance between concolic execution and fuzzing, to mitigate the path explosion problem of the former and the redundancy of the latter (Section 3).
- We propose *approximate path-preserving fuzzing*, which extends a constrained sampling technique, QUICKSAMPLER [11], to generate inputs efficiently that preserve a given path with high probability (Section 4.2).
- We conduct experiments to demonstrate that LEGION is competitive against other state-of-the-art approaches (Section 6.2.1), and evaluate the effect of different MCTS hyperparameter settings.

2 OVERVIEW

LEGION generalises the two traditional approaches to coverage-based testing: concolic execution and coverage-guided fuzzing.

Concolic execution relies on a constraint solver to generate concrete inputs that can traverse a particular path of interest. New paths are selected by flipping path constraints of previously-observed execution traces, thereby attempting to cover all feasible execution paths. However, it suffers from the high computation cost of constraint solving and exponential path growth in large applications.

Coverage-guided fuzzing has become increasingly popular during the past decade due to its simplicity and efficiency [6, 33]. It generates randomised concrete inputs at low cost (e.g. via bit flipping), by mutating inputs previously observed to lead to new execution paths. However, the resulting inputs more often than not fail to uncover new execution paths, nor to satisfy complex constraints to reach deep program states.

The complementary nature of these two techniques [28], which LEGION’s design harnesses, is highlighted by considering the exploration of the program ackermann02 in Fig. 1. This program is drawn from the Test-Comp 2020 benchmarks² [2].

The program takes two inputs, m and n (lines 8/9), and then reaches a choke point in line 11, which likely takes the “common branch” when m and n are chosen at random and thus immediately returns in line 12.

Concolic execution can compute inputs to penetrate the choke point to reach the “rare branch” (lines 14–16), but generating sufficient inputs to cover all paths of the recursive function ackermann via constraint solving is unnecessarily expensive. In comparison, a hypothetical random fuzzer restricted to generating values of $m \in \{0, 1, 2, 3\}$ and $n \in \{0, \dots, 23\}$ will quickly uncover all paths of ackermann without the need to consider exponentially growing sets of constraints from the unfolding of its recursive calls. Note that in general finding such closed-form solutions is expensive

²<https://github.com/sosy-lab/sv-benchmarks/tree/testcomp20>

```

1  int ackermann(int m, int n) {
2    if (m==0) return n+1;
3    if (n==0) return ackermann(m-1,1);
4    return ackermann(m-1,ackermann(m,n-1));
5  }
6
7  void main() {
8    int m = input(), n = input();
9    // choke point
10   if (m < 0 || m > 3) || (n < 0 || n > 23) {
11     log(n,m);           // common branch
12     return;
13   } else {
14     int r = ackermann(m,n); // rare branch
15     assert(m < 2 || r >= 4);
16   }
17 }

```

Figure 1: Ackermann02.c

and/or undecidable for programs constraints that involve nonlinear arithmetic or bitwise operations. However, we note that it is instead sufficient if the inputs preserve a chosen path prefix with high probability.

Since they have complementary benefits, much prior work has sought to combine concolic execution and fuzzing [28, 37]. But how should they be combined and when should each be used? For instance, in the example would it be more efficient to quickly flood both branches with unconstrained random inputs, or to focus on uncovering the “rare branch” even though each input generation takes longer?

LEGION provides a general-purpose answer to this question by collecting and leveraging statistics about the execution of the two branches. LEGION collects statistics about program executions while simultaneously iteratively exploring the program’s execution paths and branching structure, unifying this information together into a common tree-structured search space on the fly. Fig. 2 illustrates such a tree representation of Ackermann02.c. The tree root corresponds to the program entry point and every child node represents a conditional jump target of its parent. Each node of the tree thus represents a partial program path. Each stores statistics about the concrete executions observed so far that follow that path (i.e. pass through that node).

At each iteration of the search, these statistics allow LEGION to decide which node is most worthy of further investigation. Having selected a node, LEGION unifies concolic execution (input-generation via solving path constraints) and fuzzing in the form of *approximate path-preserving fuzzing* (APPFUZZING), to target that part of the search space. APPFUZZING is designed to *sample* inputs that pass through the node, i.e. to generate inputs that, with high probability, cause the program to follow the execution path from the root to the node selected, but also distribute relatively randomly and uniformly among all child paths of the selected node. APPFUZZING uses a combination of constraint solving, applied to the node’s path condition, and controlled mutation applied to solver-generated inputs. Statistics are gathered from the concrete executions produced by APPFUZZING to further refine LEGION’s understanding of the search space of the program under test, for subsequent iterations of the search.

In essence, these statistics capture the value of sampling from a particular node. Since LEGION’s goal is to maximise coverage, its

Figure 2: A tree representation of Ackermann02

current implementation measures value in terms of new execution paths uncovered (during concrete execution of the program on those inputs). Thus statistics such as the ratio of observed paths over total paths serve to estimate the potential of finding new paths by sampling from a particular subtree. This enables LEGION to address the following two trade-offs:

- (1) **The trade-off in depth** between a parent program state and its child. Sampling under the path constraints of a parent node (e.g. the choke point in the example) has three benefits: a) Symbolic execution to the child states (common and rare branches) can be avoided; b) inputs generated from the parent may traverse execution paths of all children; and c) parents tend to have simpler constraints to solve. However, this may waste time on children with constraints that are easy to satisfy. For instance, in the example, sampling the entry state of the program may cover both branches, but many executions will repeatedly traverse the common branch, leaving the more interesting recursive function `ackermann` insufficiently tested. Sampling from the rare branch state can guarantee executions to pass through the recursive function, but it is more computationally expensive and will miss the function `log(n,m)`.
- (2) **The trade-off in breath** between siblings program states. Given two siblings, sampling from either can potentially gain the benefit of covering more paths and collecting more statistical information of the program structure underneath the selected node, at the opportunity cost of losing the same benefit on its sibling. For example, choosing to sample inputs for the “rare branch” will cover more paths of the function `ackermann`, but it comes with a higher cost for input generation via constraint solving and can neither cover the sibling nor learn about the subtree beneath it.

LEGION’s statistical approach allows it to address these trade-offs by adopting neither a breath- nor depth-first approach, but instead a *best-first* strategy.

LEGION’s best-first strategy is a variation of the Monte Carlo tree search algorithm [7], a popular AI search strategy for large search problems in the absence of full domain knowledge. A sequential decision-making framework, MCTS carefully balances the choice between selecting the strategy that appears to be most rewarding based on current information (*exploitation*), vs one that appears

to be suboptimal after few observations but may turn out to be superior in the long run (*exploration*).

We argue that MCTS is particularly well-suited for automated coverage testing, not only due to its utility on large search spaces without full domain knowledge or even domain heuristics, but also because it is known to perform well on asymmetric search spaces (as exhibited in the running example and ubiquitous in software generally).

To utilise the full potential of MCTS on coverage testing, LEGION addresses the following challenges:

- The selection policy of MCTS, controlled by its *score* function (explained later in Section 4.1), plays a key role in its performance. However there is no well-established strategy or heuristic in coverage testing to determine the most efficient program compartment to focus in different scenarios. Thus LEGION implements a *modular* score function, flexible to a range of heuristics (Section 4.1) that we show can be efficiently instantiated for coverage-based testing (Section 5.1).
- How should MCTS hyperparameters (Section 5.2) be chosen for a program under test, and what is their influence on LEGION’s performance? We answer this question by empirical evaluation, showing that LEGION can perform effectively with naive hyperparameter choices across a wide variety of programs, as well measuring the sensitivity of its performance to the choice of parameters (Section 6).
- To achieve maximal efficiency MCTS requires being able to randomly and uniformly simulate plays from the node selected at each iteration. For LEGION this means being able to generate many random and uniform inputs that can traverse the selected program state by mutating solutions from constraint solving. Neither fuzzing nor constraint solving alone are sufficient for this purpose, for which LEGION introduces approximate path-preserving fuzzing, explained in Section 4.2.

3 LEGION MCTS

The key insight of LEGION is to generalise concolic execution and fuzzing in a principled way with the Monte Carlo tree search algorithm. To do so, LEGION makes two modifications to traditional MCTS to make it more suitable for coverage-based testing, which we describe in this section. These concern respectively the MCTS tree nodes and the tree construction.

3.1 Tree Nodes

As mentioned in Section 2, LEGION treats testing as searching a tree-structured space that represents the reachable states of the program, an exemplar of which is depicted in Fig. 2. Each tree node is identified by the address of a code block in the program under test. The tree root corresponds to the program entry and every child node represents a conditional jump target of its parent.

As mentioned, each node represents the (partial) execution path from the program’s entry point to it, and stores statistics about the concrete program executions observed so far to follow that path, which are used by the MCTS selection policy to decide which part of the tree to investigate at each iteration of the search algorithm. When the MCTS algorithm selects a node for investigation, LEGION

then generates new program inputs that (with high probability) cause the program to traverse the path represented by that node. To do so, recall that LEGION uses a form of directed fuzzing called approximate path-preserving fuzzing (APPFUZZING), which is a hybrid of mutation fuzzing and input-generation by constraint solving. LEGION’s APPFUZZING implementation thus benefits from symbolic information about the program state that corresponds to the tree node (i.e. the state reached after traversing the partial path that the node represents), specifically the symbolic path condition. Therefore, certain nodes in the tree also carry a *symbolic state*, which encodes the corresponding path condition.

However, not all tree nodes carry a symbolic state. Recall, from Section 1 that LEGION performs symbolic execution lazily. This means that after observing a new concrete execution path, that path will be integrated into the tree but without symbolically executing it. Indeed LEGION defers symbolically executing a path until it has evidence that investigating that path could be beneficial for uncovering additional new execution paths (as determined by the MCTS selection algorithm). Until a path is symbolically executed, the tree nodes that represent it do not contain symbolic state information.

Thus, LEGION’s tree nodes come in a variety of types, depending on their purpose and the information they contain. For instance, hollow nodes are ones that represent observed concrete execution paths but do not (yet) carry symbolic state information, while solid nodes additionally do carry satisfiable symbolic state information; phantom nodes represent feasible program paths observed and validated during symbolic execution but not yet observed during concrete execution. Two further types also exist: redundant and simulation, as described below. We discuss each type of node, with occasional reference to Fig. 3.

Hollow nodes are basic blocks found by concrete, binary execution, but have not been selected for symbolic execution yet, and hence do not have their symbolic states attached. They are used to mark the existence (and reachability) of paths and to collect statistics of observed rewards. When a hollow node is selected for investigation by the MCTS selection step (explained shortly), LEGION invokes symbolic execution from its closest solid ancestor, to obtain a path condition to seed APPFUZZING. By doing so, the node will be re-categorised as one of the following two types. Each hollow node appears as a hollow round node in Fig. 3.

Solid nodes are hollow nodes with symbolic states attached and whose path condition has one more constraint over that of its closest solid ancestor, i.e. they are not redundant. Each solid node has a special child node called a *simulation child* (of square shape), described below. Each solid node appears as a solid round node in Fig. 3.

Redundant nodes represent states for which APPFUZZING is (a) impossible or (b) redundant: either (a) they were observed during concrete execution but not during symbolic execution (e.g. due to under-approximation and concretisation during symbolic execution), and so cannot be selected for APPFUZZING because they lack symbolic state information, or (b) have exactly the same symbolic constraint as their closest solid parent (e.g. jump targets of tautology conditions), for which APPFUZZING is redundant because

Figure 3: The four stages of (each iteration of) LEGION's MCTS algorithm.

identical information is already captured by their closest solid ancestor. Hence LEGION never simulates from them (i.e. never selects them for APPFUZZING). No redundant nodes are depicted in Fig. 3.

Phantom nodes denote symbolic states whose addresses have not yet been seen during concrete execution, but have been proved to exist by the symbolic execution engine. They are found during the selection stage, when the symbolic execution engine shows that the symbolic state of a hollow lone child has a sibling state. They act as place holders for as-yet unseen paths, waiting to be revealed by concrete, binary execution (i.e. by APPFUZZING). When that happens, a phantom node will be replaced by a solid node, when LEGION integrates the observed concrete execution path into the tree. No phantom nodes are depicted in Fig. 3.

Simulation nodes. Different from vanilla MCTS, LEGION allows sampling from intermediate nodes, which represents APPFUZZING applied to some point along a partial execution path. To incorporate this into MCTS, we add special leaf nodes to the tree whose shape is a square and whose parent is always solid. When the MCTS selection stage chooses a simulation node for APPFUZZING, this choice implies that LEGION believes that directing fuzzing towards the program state represented by the solid parent is more beneficial than fuzzing any of its child program states (where “beneficial” means the estimated likelihood of discovering a new path). Each simulation node appears as a square node in Fig. 3.

3.2 Tree Construction

Fig. 3 illustrates how LEGION uncovers the tree-structured search space in its variation of MCTS: Each iteration of the search algorithm proceeds in four stages. In Fig. 3, bold lines highlight the actions in each stage. Solid thin lines and nodes represent paths and code blocks that have been covered and integrated into the tree. Dotted thin lines and nodes represent paths and nodes not yet found. The four stages are as follows:

Selection. Shown by bold arrows, LEGION descends from the root by recursively applying a *selection policy* (Section 4.1), until it reaches a simulation node (square shape in Fig. 3). Upon visiting a hollow node (i.e. that carries no symbolic state, depicted as hollow circles), it performs symbolic execution from the nearest solid ancestor (depicted as filled nodes) to compute the symbolic state of the hollow node, adds the symbolic state to that node (turning it solid) and attaches to it a simulation child (the square).

Simulation. Having selected which program state to target for investigation, LEGION applies APPFUZZING (Section 4.2) on the program state represented by the (solid parent of) the selected simulation node to generate a collection of inputs that, with high probability, will cause the program to follow the path from the root to the solid parent. It then executes the program on those inputs and observes the resulting execution traces, which are represented by dashed circles and lines in Fig. 3.

Expansion. This step involves mapping each observed execution trace to a path of corresponding tree nodes: new traces cause new hollow (round) nodes to be added to the tree.

Back-propagation. This step updates the statistical information recorded in the tree, to take account of the executions observed during the simulation step. LEGION computes the *reward* (Section 5.1) from each simulation (i.e. from each concrete execution performed during the simulation step) and propagates the reward to all nodes along the execution trace. Note that the path(s) observed during concrete execution might differ from that chosen during the simulation step, because APPFUZZING is necessarily approximate. Therefore, in this step, LEGION also propagates the reward to each node on the path from the root to the node chosen during the selection step.

4 DESIGN CONSIDERATIONS

Having described LEGION's adaptation of MCTS at a high-level, we now discuss two of the most critical parts of its design, namely the design of the selection policy used during the selection phase of each MCTS iteration, and the design of LEGION's directed fuzzing implementation, APPFUZZING. We discuss further implementation details and considerations in Section 5.

4.1 Selection Policy

One of the central design goals of LEGION is to generalise a range of prior coverage-based testing methods. Doing so requires a general policy for selecting which tree node to investigate at each MCTS iteration, to make LEGION adaptable to different search strategies and coverage reward heuristics, with a modularised design.

At its simplest, the selection policy's job it to assign a *score* to each node. This score is used as follows, during the selection step

of each MCTS iteration. Recall that selection proceeds via recursive descent through the tree towards that part of the tree deemed most worthwhile to investigate. The score guides this descent. Selection starts at the root. Then the immediate child with the highest score is traversed (with ties being broken via uniform random selection). Traversal proceeds to that child's highest-scoring child, and so on.

In LEGION, the score of each node represents the optimistic estimation that APPFUZZING will discover new paths when applied to the (program state represented by the) node. Scores are computed by applying the *Upper Confidence Tree (UCT) algorithm* [7] on the rewards obtained during the prior simulation (i.e. during the prior APPFUZZING). We leave a discussion of rewards to Section 5.1 but for now it suffices to understand that rewards correspond to the discovery of new execution paths. For a node N we define its score, $UCT(N)$ as follows. For a newly created node, its score is initialised to ∞ , to ensure that uninvestigated subtrees are always prioritised over their siblings. Otherwise, $UCT(N)$ is:

$$UCT(N) = \bar{X}_N + \rho \sqrt{\frac{2 \ln P_{sel}}{N_{sel}}} \quad (1)$$

where $\bar{X}_N$ is the average past reward from N , ρ is a hyperparameter (described below) that defines the MCTS *exploration ratio*, N_{sel} and P_{sel} are, respectively, the number of times that node N and its parent have been traversed so far during prior selection stages.

This score is designed to balance two competing concerns when searching the tree, namely *exploitation* vs. *exploration*. Exploitation corresponds to investigating a part of the tree that has proved rewarding in the past (high $\bar{X}_N$), in the belief that because it has yielded new execution paths before it is likely to do so again in the future. Exploration, on the other hand, correspond to investigating a part of the tree that, so far, appears under-explored (low P_{sel}/N_{sel}), and where rewards are less certain. The second term derives from an upper bound of a confidence interval for the true mean, based on the multi-armed bandit algorithm UCB1 [17].

The precise balance between these two concerns is controlled by the choice of non-negative hyperparameter ρ . Were ρ zero (or very large), LEGION would base its selection decisions only on past rewards (resp. visitation counts), corresponding to pure exploitation (resp. exploration). Instead ρ should be chosen to be some small positive value, to permit exploration. We investigate in more detail how ρ affects LEGION's performance. For any fixed choice, as a node's parent is traversed without the node visited, the fraction $\sqrt{\frac{2 \ln P_{sel}}{N_{sel}}}$ grows, causing LEGION to favour the child.

4.2 Approximate Path-Preserving Fuzzing

Having selected a node (i.e. program state) to target, LEGION then applies its approximate path-preserving fuzzing algorithm APPFUZZING to generate inputs that when supplied to the program under test will, with high probability, cause the program to execute from the entry point to reach that state, following the path from the tree root to the selected node.

APPFUZZING can be seen as a hybrid of input generation via constraint solving and mutation fuzzing, and is seeded by the symbolic path condition stored in the selected simulation node. LEGION's APPFUZZING is inspired by QUICKSAMPLER [11], a recent algorithm for sampling likely solutions to Boolean constraints that mixes SAT

solving and bit mutation. LEGION's APPFUZZING on the other hand operates on SMT constraints, specifically the bit-vectors theory of the constraints produced by LEGION's symbolic execution engine, angr [26].

```
def APPFuzzGen(constraint):
    Σ = {}
    σ = solve(constraint)
    yield {σ}
    for each bit  $b_i$  of σ do
         $\tilde{\Sigma} = \{\}$ 
         $\sigma' = solve(constraint \wedge \neg b_i)$ 
         $\tilde{\Sigma} = \tilde{\Sigma} \cup \{\sigma'\}$ 
        for  $\sigma''$  in Σ do
             $\tilde{\Sigma} = \tilde{\Sigma} \cup \{(\sigma \oplus ((\sigma \oplus \sigma') \vee (\sigma \oplus \sigma'')))\}$ 
        end
        yield  $\tilde{\Sigma}$ 
        Σ = Σ ∪  $\tilde{\Sigma}$ 
    end
```

Algorithm 1: Input Generation for APPFUZZING

```
def APPFuzz( $N_{samples}$ , node):
    results = []
    while len(results) <  $N_{samples}$  do
        results.append(APPFuzzGen(node.path_constraint))
    end
```

Algorithm 2: Approximate Path-Preserving Fuzzing

The APPFUZZING algorithm, APPFuzz, is depicted in Algorithm 2. Its inputs comprise the selected node to which APPFUZZING is being applied, as well as constant $N_{samples}$ the *minimum* number of new inputs that APPFUZZING should attempt to generate. Like ρ , $N_{samples}$ is a hyperparameter, discussed further in Section 5.2.

Input generation is delegated to a helper *generator* function, depicted in Algorithm 1, which is repeatedly called until at least $N_{samples}$ have been generated. By design, the helper generator, APPFuzzGen, can yield a variable number of inputs each time it is called. Hence, APPFuzz can often end up generating more than the minimal number of required inputs for each simulation.

Input generation, handled by APPFuzzGen in Algorithm 1, is performed using a mixture of constraint solving and mutation of prior solver-generated solutions. Thus APPFuzzGen takes as a parameter the path constraint *constraint* of the node in question. On its first invocation it simply solves the node's constraint and returns that solution σ . However, when subsequently invoked for the same node (e.g. if multiple inputs are required for simulation, when $N_{samples} > 1$) its execution restarts from the point directly after the first **yield** statement, i.e. at the entry to the outer loop. Each iteration of this outer loop generates an ever-larger set of new inputs, and the generator returns the newly-generated inputs after each iteration. Subsequent invocations of the generator resume execution from the point just after the second **yield** statement, i.e. at the point where the next iteration of the outer-loop proceeds.

The outer loop iterates over each bit b_i of the initial solver-generated solution σ . In each outer loop iteration, the generator invokes the constraint solver once to generate a solution σ' to the path constraint that, if one exists, is guaranteed to be distinct from σ by differing in bit b_i . Then, for each input σ'' generated in previous iterations of the outer loop, σ' is *mutated* with σ'' and σ to produce a new input. The σ' and all of the new mutation-generated inputs are then returned.

Thus each invocation of the generator causes a single invocation of the constraint solver; however repeated invocations yield exponentially more inputs. This process continues until the outer-loop terminates. The generator will begin afresh on subsequent invocations.

Not depicted in Algorithm 1 is the fact that previously generated inputs are remembered, to avoid returning duplicate inputs. Once the solver can return no new solutions, a node is marked as *exhausted* and no further simulation (APPFUZZING) will be performed on it (see Section 5.1.1).

The mutation operator combining σ , σ' and σ'' , via bitwise exclusive-or $\oplus$ simulates the mutation operator used in a similar fashion by QUICKSAMPLER [11] on Boolean constraints, lifting it to bit-vectors. That operator was shown, with high probability, to produce results that each satisfy the Boolean constraint [11] and, in aggregate, are uniformly distributed.

For this reason, we conjecture that our APPFUZZING implementation is likely to produce inputs that preserve the path to the selected node and that satisfy i.i.d. requirements of MCTS simulations.

In scenarios where one wishes to produce more inputs per constraint solution, our implementation can be adjusted to do so by having it perform more mutations per outer loop iteration. However doing so is likely to reduce the accuracy of the results, i.e. the probability that they will satisfy the given path condition.

5 PRACTICAL CONSIDERATIONS

With the major design considerations out of the way, we now turn to the salient details of LEGION's current implementation. These concern respectively the reward heuristic used in the UCT score function of its MCTS implementation, as well as the choices of the various hyperparameters like ρ and $N_{samples}$.

5.1 A Reward Evaluation Heuristic

In MCTS, rewards quantify the benefits observed during each simulation. In LEGION's MCTS, simulation corresponds to concrete execution of the program under test, via APPFUZZING. An important implementation consideration is what should be the reward function? Put another way, what should rewards correspond to? When should a simulation (concrete execution) be considered rewarding?

Recall that the average past rewards associated with a node N was denoted $\bar{X}_N$, in the UCT score function defined in Eq. (1) in Section 4.1. By choosing different instantiations for this term, one can naturally adapt LEGION to various exploitation strategies.

However, since the goal of LEGION's present implementation is to discover the maximum number of execution paths in the shortest possible time, in this paper we implement and evaluate a simple

reward heuristic: the reward of a simulation is the number of new paths found by APPFUZZING.

Recall that for a node N , N_{sel} denotes the number of times N has been traversed during selection, and that P_{sel} does likewise for N 's parent. Then we denote by N_{win} the number of distinct execution paths discovered so far that pass through node N .

The average past reward associated with node N is then $\frac{N_{win}}{N_{sel}}$, the ratio of the number of paths found so far that pass through this node, compared to the number of times it (or a child) has been selected for APPFUZZING. This ratio is chosen under the assumption that the more new paths were found by running APPFUZZING on a node in the past, the more potential the node has in increasing coverage in the future.

Plugging this into Eq. (1), and remembering that the initial score of a node is set to ∞ , we arrive at the following instantiation of the UCT score in LEGION's current implementation:

$$UCT(N) = \begin{cases} \infty & N_{sel} = 0 \\ \frac{N_{win}}{N_{sel}} + \rho \sqrt{\frac{2 \ln P_{sel}}{N_{sel}}} & N_{sel} > 0 \end{cases}$$

To better understand this heuristic, let us return to the example of Fig. 1. Note that only a specific pair of m and n can violate the assertion on line 15, but uncovering that assertion requires being able to get past the recursion of the ackermann function. The recursion of ackermann is therefore an attractive nuisance for uncovering the assert statement, since recursion necessarily produces new execution paths. This is why the exploration component of the UCT score is critical. At the same time, the use of such a fine-grained coverage metric, which tracks individual execution paths, instead of more coarse-grained metrics like statement and branch coverage or AFL's coverage maps, ensures that LEGION does not overlook paths that can arise only after deep recursion.

As we show later in Section 6.2, even with this simple heuristic, LEGION performs surprisingly well, and the heuristic appears relatively robust. Yet there of course exists much scope to consider heuristics that take into account other relevant information, such as time consumption, subtree size or other static properties of the program. We leave the investigation of such for future work.

5.1.1 Optimisations. This heuristic permits a number of optimisations on node selection, which we have implemented in LEGION. Essentially these allow LEGION to decide when a node will no longer produce any future rewards (i.e. that no new paths can be found via APPFUZZING applied to the node). When doing so, LEGION overrides the node's score to $-\infty$. In this case we say that the node is *pruned* from the tree. However note that the node is never physically removed from the tree: its score is just overridden to ensure it will never be selected.

Fully explored nodes. A node is fully explored if there is no undiscovered path beneath it. For example, final nodes of complete concrete execution paths are fully explored. A parent node will be pruned if all non-simulation children are fully explored. Given a fully explored node may have hidden siblings, LEGION will prune it only after identifying all of its siblings via symbolic execution.

Useless simulation nodes. A simulation node is useless if it has less than two not fully explored siblings. Recall that simulation nodes are children of (solid) nodes and that solid nodes represent

a point in a concrete execution path and contain a symbolic path condition. When a simulation node is selected during the MCTS selection stage (Section 3) this represents the decision to target its solid parent node with APPFUZZING. Doing so essentially will yield some paths from the full set of paths S in the subtree under the solid parent. On the other hand, sampling from one of the simulation node's siblings (i.e. from another child of the solid parent) will yield paths that are drawn from a strict subset T of S . Sampling from S can be beneficial if it can yield paths from multiple such T . However, if there is only one such T remaining, it is better to sample from it than from S . Hence, in this case, the simulation node is pruned.

Exhausted simulation nodes. A simulation node is exhausted if no input can be found when applying APPFUZZING on it. This happens when all inputs that the solver is capable of producing under the constraint of the node have been generated.

Under-approximate symbolic execution. Due to under-approximation, the symbolic execution engine might incorrectly classify a feasible state as infeasible, e.g. due to early concretisation. This creates a mismatch between the symbolic execution results and concrete execution traces. In this case LEGION cannot apply APPFUZZING to such a node or any of its decedents, but must instead target its parents. Under-approximation can cause LEGION to erroneously conclude that a subtree has been fully explored. To mitigate this issue, LEGION can be run in “persistent” mode, which continues input generation even for apparently fully-explored subtrees (see Section 6.1.1).

Recall from Section 2 that to construct the search tree LEGION uses binary instrumentation to collect the addresses of conditional jump targets traversed during concrete execution. Doing so produces a trace (i.e. a list) of conditional jump target addresses. LEGION can be configured to limit the length of such traces to a fixed bound, mitigating the overheads of instrumented binary execution [1] and, hence, the depth of the search tree. The number is controlled by a hyperparameter named *tree depth* (see Section 5.2).

5.2 Hyperparameters

Another important practical consideration is the choice of the various hyperparameters that control LEGION's behaviour. Some of these, like ρ and N_{samples} we have already encountered. However we summarise them all below for completeness.

Naturally, different choices of the hyperparameters will bias LEGION's performance in favour of certain kinds of programs under test. We investigate this effect in Section 6.2.2.

Exploration ratio ρ . The exploration ratio of the MCTS algorithm controls the amount of exploration performed, in comparison to exploitation (see Section 4.1), by the selection phase of MCTS.

Number of cores. LEGION supports the *leaf parallelisation* of MCTS, wherein simulations are run in parallel. For LEGION this corresponds to running multiple concrete executions in parallel, when APPFUZZING returns multiple inputs. The maximum number of such simultaneous parallel executions is controlled by this hyperparameter.

Tree depth. LEGION can limit the depth of its search tree to this parameter by forcibly terminating concrete executions once the length of the trace of conditional jump targets produced by the binary instrumentation reaches this value.

Concrete execution timeout. This hyperparameter controls when LEGION will forcibly terminate concrete executions (produced by APPFUZZING) that take too long.

Symbolic execution timeout. Recall that when LEGION decides to target a node for APPFUZZING that does not yet contain its symbolic path condition, it then symbolically executes to that node from its nearest symbolic ancestor. The symbolic execution timeout is used to forcibly terminate such symbolic executions if they take too long. If such a timeout occurs, LEGION will proceed to select the nearest symbolic ancestor produced by the terminated symbolic execution.

Minimum number of samples per simulation N_{samples} . Recall that this parameter controls the minimum number of inputs to generate for each APPFUZZING invocation, and that each such invocation might return more than this minimum (see Section 4.2).

Persistent. When set, this boolean flag causes LEGION to continue input generation even for apparently fully-explored subtrees to mitigate under-approximate symbolic execution (Section 5.1.1) and expensive constraint solving.

6 EVALUATION

We designed two kinds of experiments to evaluate LEGION's current design and implementation. We sought to answer two fundamental questions, respectively: (1) is LEGION effective at generating high coverage test suites in fixed time, as compared to other state-of-the-art tools? and (2) what is the effect of the choices of LEGION's various hyperparameters on its performance, and do there exist suitable default choices for these that are robust across a variety of input programs?

Peer competition. To answer the first question, LEGION competed in the Cover -Branches category of Test-Comp 2020 test generation competition. LEGION has evolved since the competition; this paper reports the result of running the latest version of LEGION on the same benchmark suite³ on the same host machine with the same resources controlled by the same benchmarking framework as was used in Test-Comp 2020, thereby allowing our results to be compared against results obtained during the competition.

Sensitivity evaluation. To answer the second question, we selected a carefully chosen subset of the benchmark programs from the Test-Comp 2020 benchmarks, and then we evaluated the effect of varying LEGION's various hyperparameters on its ability to perform on these programs. Again, we report the results based on the most recent version of LEGION.

6.1 Experiment Setup

To maximise reproducibility, we carried out both evaluations using the BenchExec framework⁴ [4]. It allows the usage of 8 cores and 15GB memory. Our evaluations used Test-Comp 2020 benchmarks in which, for each, coverage percentages computed by testcov⁵ were reported after 15 minutes. There are 2531 programs in total in

³<https://github.com/sosy-lab/sv-benchmarks/tree/testcomp20>

⁴<https://gitlab.com/sosy-lab/software/benchexec/-/tree/2.5>

⁵<https://gitlab.com/sosy-lab/software/test-suite-validator/-/tree/testcomp20>

Settings	Score function	ρ	Core	Tree Depth	ConEx Timeout	SymEx Timeout	$N_{samples}$	Score	p-value
Baseline	UCT	$\sqrt{2}$	8	10^5	–	–	1	48.9018	
Score Function = Random	Random	$\sqrt{2}$	8	10^5	–	–	1	48.1632	0.004404148
$\rho = 0$	UCT	0	8	10^5	–	–	1	53.9418	0.001752967
$\rho = 0.0025$	UCT	0.0025	8	10^5	–	–	1	53.7572	0.00298362
$\rho = 100$	UCT	100	8	10^5	–	–	1	49.0115	0.600806429
Core = 4	UCT	$\sqrt{2}$	4	10^5	–	–	1	49.0633	0.355824737
Core = 1	UCT	$\sqrt{2}$	1	10^5	–	–	1	49.4586	0.000250096
Tree Depth = 10^2	UCT	$\sqrt{2}$	8	10^2	–	–	1	45.3216	0.001406423
Tree Depth = 10^3	UCT	$\sqrt{2}$	8	10^3	–	–	1	49.0667	0.340845887
Tree Depth = 10^7	UCT	$\sqrt{2}$	8	10^7	–	–	1	48.8048	0.529105141
ConEx Timeout = 0.5	UCT	$\sqrt{2}$	8	10^5	0.5	–	1	48.9121	0.950464912
ConEx Timeout = 1	UCT	$\sqrt{2}$	8	10^5	1	–	1	49.0451	0.422595538
SymEx Timeout = 5	UCT	$\sqrt{2}$	8	10^5	–	5	1	49.0368	0.3715003
SymEx Timeout = 10	UCT	$\sqrt{2}$	8	10^5	–	10	1	49.1712	0.068867779
$N_{samples} = 3$	UCT	$\sqrt{2}$	8	10^5	–	–	3	49.0362	0.569679957
$N_{samples} = 5$	UCT	$\sqrt{2}$	8	10^5	–	–	5	49.1774	0.31321928

Table 1: The sensitivity experiment: compared hyperparameter settings. Only one setting is changed in each alternative, indicated by the first column, including: the score function, exploration ratio (ρ), the number of cores (Core), concrete execution timeout (ConEx Timeout), symbolic execution timeout (SymEx Timeout), and the minimum number of samples per simulation ($N_{samples}$). The significant p-values (< 0.05) are in bold font

the Test-Comp 2020 set, sorted into 13 *suites*. LEGION⁶ requires two dependencies, *angr*⁷ and *claripy*⁸.

6.1.1 Peer competition. All tools to which we compared LEGION in the peer competition were seeded with the same initial dummy input: 0. We purposefully did not fine-tune the hyperparameters for the competition. For instance, we used the value of $\sqrt{2}$ for ρ , which is commonly used for UCT scores [7]. We limited the tree depth to be within 100000, and the number of cores to be 8, as they are sufficient for LEGION in most cases. We did not time out symbolic or concrete execution, and run at least one input in each simulation stage ($N_{samples} = 1$) like the original MCTS. The results of our sensitivity experiments would later suggest that the performance of LEGION is not sensitive to most of these parameters when they were chosen reasonably. (Section 6.2.2).

Each experiment was conducted on the same host machine as the Test-Comp 2020 competition, with a 3.40 GHz Intel Xeon E3-1230 v5 CPU, running at 3800 MHz, with Turbo Boost disabled.

To support a fair and meaningful comparison, we excluded two benchmark suites, *BusyBox-MemSafety* and *SQLite-MemSafety*. All programs in the former cause source code compilation errors due to name conflicts. All tools scored close to 0 on the latter suite, because of an unknown issue.

Since LEGION’s goal is maximising coverage, our evaluation was restricted to the coverage (as opposed to the error-finding) category (i.e. Cover-Branches) of Test-Comp 2020.

For the peer competition we ran LEGION in “persistent” mode (discussed in Section 5.1.1) to ensure that it would make full use of the time given to it for each benchmark program.

6.1.2 Sensitivity Experiments. We evaluated LEGION’s sensitivity to the various hyperparameters by measuring its performance on the Sequentialized benchmark suite from the Test-Comp 2020 benchmarks. This benchmark suite was selected following the competition results, which showed that LEGION had some room to improve its coverage score for this suite, yet the competition result for it was not abysmal either. Hence, effects on LEGION’s performance here would be clearly visible in the final coverage score obtained for this benchmark suite. This suite is also of sufficient size, containing 81 programs that together comprise 26166 branches (323 branches on average per program).

The various hyperparameter settings that we compared are listed in Table 1. For these experiments, LEGION was run on a machine with a 2.50 GHz Intel Xeon Platinum 8180M CPU at 2494 MHz. The “persistent” mode flag was not enabled during these sensitivity experiments to more accurately demonstrate how other hyperparameters affect the results.

6.2 Results

6.2.1 Peer Competition. The normalised⁹ results from the peer competition appear in Table 2. Each benchmark suite is listed in the first column, with the number of programs it contains in brackets. That number is also a theoretical upper bound on the total coverage score that can be achieved for each category. The following columns list the score of each competitor on each suite, where the score is

⁶<https://github.com/Alan32Liu/Legion/tree/TestComp2020-ASE2020v3>

⁷<https://github.com/Alan32Liu/angr/tree/TestComp2020-ASE2020>

⁸<https://github.com/Alan32Liu/claripy/tree/TestComp2020-ASE2020>

⁹<https://test-comp.sosy-lab.org/2020/rules.php#meta>

the average of the coverage score (between 0 and 1) achieved on each program in the suite.

The peer experiment demonstrates that LEGION is effective across a wide variety of programs, despite the simplicity of its current heuristics and limitations of its implementation. LEGION is competitive in many suites of Test-Comp 2020, achieving within 90% of the best score in 5 of the 11 suites in the Cover-Branches category. LEGION outperformed all other tools on some benchmarks (discussed below), which highlights the general effectiveness of the approach.

6.2.2 Sensitivity Experiments. For each of the hyperparameter settings investigated, Table 1 reports the Sequentialized suite coverage result. It also reports the p-value obtained by applying Student’s t-test (paired, two-tailed) to determine whether the difference from the baseline was statistically significant. Conventionally, we claim the difference is significant when that value is less than 0.05.

As to be expected, These results indicate UCT scoring superior to the baseline random selection. LEGION’s performance is also sensitive to the choice of its exploration ratio ρ . Small values of ρ near 0.0025 appear to work best, based on these results, and outperform pure exploitation ($\rho = 0$). For currently unknown reasons, using 1 core appears to have a slightly better performance. The choices of other hyperparameters, including the tree depth, concrete execution timeout, symbolic execution timeout, and the minimum number of inputs to generate for each APPFUZZING ($N_{samples}$), appear to be not influential when chosen within a reasonable range.

However, the value of ρ still exhibited very limited effect across all suites when we were measuring LEGION’s overall performance to rule out overfitting. The most preferable setting above ($\rho = 0.0025$) improved the overall competition score by merely 1.04%.

6.3 Discussion

LEGION inherits limitations from angr and MCTS.

angr, as symbolic execution backend of LEGION, eagerly concretises size values used for dynamic memory allocations, which then causes it to conclude erroneously that certain observed concrete execution paths are infeasible. The design of LEGION’s MCTS allows it to work around this limitation: LEGION first detects this case from the mismatch between symbolic and concrete execution, then omits the erroneous programs states from the selection stage. When MCTS predicts new paths under them, LEGION invokes APPFUZZING on their parents instead. With this mitigation, LEGION achieved full coverage on `loops/sum_array-1.c` (in contrast to all other tools) despite the dynamic allocations.

angr uses VEX IR which prior work [35] has noted leads to increased overheads in symbolic execution. LEGION mitigates this by performing symbolic execution lazily, only computing the symbolic states of nodes selected by MCTS.

angr cannot model arrays that have more than 4096 (0x1000) elements. In particular, it failed on 16 benchmarks in suite `Arrays` that involves multidimensional arrays.

angr only supports 22 out of the 337 system calls in the Linux kernel 2.6 [35, 37], and is known to scale poorly to large programs [20].

The nature of MCTS makes LEGION less suitable for exploring spaces where the reward is rare. For example, benchmarks in suite `ControlFlow` are full of long sequences of equality constraints that are satisfied by few inputs, making new path discovery via sampling

(i.e. APPFUZZING) very rare. This benchmark suite happens to be an extreme negative example of the trade-off in depth (discussed in Section 2) that LEGION intends to balance where fuzzing the parent gives no reward. Expensive constraint solving also adds extra cost to sampling. In particular, `claripy` spent 14.82 seconds on average to solve each constraint in this suite. Generally, invoking APPFUZZING at intermediate nodes with less complicated constraints to generate mutations can amortise this cost. However, although on average LEGION generated 17.37 mutations per program at negligible costs (10^{-5} seconds per mutation) and 99.39% of them preserved their paths, none of them found any new path due to the equality constraints. To a certain extent, this could possibly be mitigated by decreasing LEGION’s exploration ratio (constant ρ in the UCT score) but such fine-tuning would overfit LEGION to the competition benchmark suites.

To analyse the impact of this extreme example, we re-normalised the scores of the peer competition with this suite excluded. It shows that the overall normalised score of LEGION (1502.98) is higher than KLEE (1492.94) and close to HybridTiger (1504.49).

MCTS works much better on less extreme examples. Benchmarks `seq-mthreaded/pals_lcr-var-start-time*.c` interleave equality constraints with range constraints, making new path discovery less rare (11.77% of the new paths are found via sampling at a path-preserving rate of 93.49%) and less expensive (each constraint solving took 0.81 seconds on average). As a result, LEGION performed roughly the same as CoveriTest on these 12 benchmark programs and largely outperformed all other competitors. Similarly, LEGION fully covered the sample program `Ackermann02.c` (shown in Fig. 1) with 77.63% of the new paths uncovered by mutations from APPFUZZING (71.88% of them preserving their path prefix) at an average cost of 1.10×10^{-5} seconds, approximately 1000 times faster than constraint solving via `claripy`.

7 RELATED WORK

We consider recent work that adopts similar ideas to LEGION, at the intersection of fuzzing and concolic execution.

Like LEGION, Driller’s [28] design is also inspired by wanting to harness the complementary strengths of concolic execution and fuzzing. It augments AFL [36] by detecting when fuzzing gets stuck and then applying concolic execution (constraint solving) to generate inputs to feed back to the fuzzer, to allow it to traverse choke points. Unlike our APPFUZZING, however, little attempt is made to ensure that subsequent mutations made by the fuzzer to each solver-generated input preserves the input’s paths condition. Thus AFL might undo some of the hard work done by the solver.

LEGION uses rewards from past APPFUZZING to predict where to perform concolic execution, modelling automated test generation as Monte Carlo Tree Search. The authors in [31] formally model programs as Markov Decision Processes with Costs. Their model is used to decide when to perform input generation via constraint solving vs. when to generate inputs via pure random fuzzing. Like LEGION, their approach tracks the past rewards from concrete execution to decide when to perform fuzzing. However, unlike LEGION their fuzzing approach is purely random. LEGION applies the UCT algorithm to make choices about which parts of the program to target at each iteration, while the approach of [31] instead applies

Participants	LEGION *	KLEE [8]	CoVeriTest [3]	HybridTiger [22]	LibKluzzer [18]	PRTest	Symbiotic [9]	Tracer-X [16]	VeriFuzz [1]
Arrays (243)	123	121	123	126	201	93.4	75.4	116	200
Floats (212)	105	59.7	107	107	107	51.2	56.4	50.8	97.7
Heap (139)	102	101	91.6	98.4	104	51.6	82.7	94.7	103
Loops (123)	97.6	93	95.8	92.5	101	54.1	71.3	87.3	97.1
Recursive (38)	34.3	32.1	30.9	30.8	34.3	9.85	21.3	27.6	34.2
DeviceDriversLinux64 (290)	22.9	21.7	23.2	22.1	22.1	9.24	21	22	22.7
MainHeap (231)	161	153	181	180	195	53.3	129	175	195
BitVectors (36)	23.2	25.7	27.7	27	28.2	3.43	18.6	27.6	28.1
ControlFlow (19)	6.71	12.6	13.9	13.9	13.8	0.558	1.68	12.7	13.4
ECA (1046)	729	957	811	766	962	489	639	932	963
Sequentialized (81)	49.7	52.5	57.4	43.5	58.2	5.28	14.5	15.7	62.5
Normalised Average**	1455.908002	1515.97673	1588.862835	1541.850942	1759.050529	584.57209	969.1994602	1382.674571	1746.261043

Table 2: Normalised accumulated coverage scores for each benchmark suite (The higher the better). LEGION outperformed KLEE in 7 out of 11 benchmark suites (in bold font). * The results are reproduced with the latest version of LEGION; ** Excluded two benchmark suites to avoid noise: BusyBox-MemSafety, SQLite-MemSafety

a greedy algorithm to estimate the rewards of random input generation. The authors’ approach [31], unlike LEGION, applies a static heuristic to estimate the cost of constraint solving in order to make an informed decision between the precision and expense of constraint solving vs. the imprecision and efficiency of pure random fuzzing. LEGION eschews this all-or-nothing choice altogether, by blending fuzzing and concolic execution via APPFUZZING to avoid the drawbacks of both while retaining their strengths.

DigFuzz [37] extends Driller and, similar to LEGION, collects statistics to decide which program states would benefit most from concolic execution, via Monte Carlo simulations.

While LEGION collects coverage statistics, DigFuzz instead focuses on learning which parts of the program are most difficult for the fuzzer to traverse, and only then does it decide to invoke the solver, on the assumption that solving is expensive and should only be used when necessary. Whereas LEGION estimates the likelihood of reward (uncovering new execution paths via directed APPFUZZING) for each node, based on past rewards from APPFUZZING, DigFuzz instead estimates the probability that the undirected fuzzer will be able to traverse each path, based on the fuzzer’s past performance. It collects statistics about how often each branch is traversed by the fuzzer. Then the probability of traversing a particular path is computed by multiplying together estimates of the probabilities on each of the branches along that path.

So while both methods use a form of Monte Carlo simulation to track statistics to decide how to generate inputs for the program, they do so for quite different purposes and using methods that appear relatively complementary. This comes from their different design philosophies: DigFuzz wants to minimise solver use as much as possible, and is happy to spend more time on repeated traversal of the paths by the cheap fuzzer. LEGION instead is willing to use the solver more often (although not as often as in traditional concolic testing), to avoid concrete executions that often repeat the same execution paths.

MCTS has also been used recently to guide traditional concolic execution [34]. The authors use MCTS to pick which branch to concolically execute. Unlike LEGION, this requires symbolic execution along the entire path, which LEGION can avoid since it only ever performs symbolic execution lazily (Section 3.1). It also requires one solver call for each concrete execution. LEGION avoids this by using APPFUZZING instead, which generates increasingly more inputs

for each solver invocation, at the price of not guaranteeing that all generated inputs will satisfy the path constraint. Finally, while LEGION’s score function applies UCT to estimate rewards in terms of discovering new execution paths, it is not clear what is the score function used in [34].

Finally, we note that other work has also attempted to improve fuzzing by adopting methods from machine learning. For instance, Learn&Fuzz [15] trains a neural network to represent the program under test, to aid fuzzing. The Angora [10] fuzzer applies gradient descent as an input mutation strategy to help it traverse conditional branches, while AFLGo [5] uses simulated annealing to guide directed fuzzing. In contrast, LEGION adopts MCTS with APPFUZZING as a principled unification of fuzzing and concolic execution.

8 CONCLUSION

How can we maximise code coverage by generalising concolic execution and fuzzing? LEGION shows a possible solution with its variation of the Monte Carlo Tree Search algorithm. By tracking past rewards from APPFUZZING, LEGION learns where concolic execution effort should be spent, namely on those parts of the program in which new execution paths have been observed in the past. Symbolic execution then returns the favour by computing new promising symbolic states to support accurate APPFUZZING. Which states to prefer is estimated using the UCT algorithm, which navigates the exploration versus exploitation trade-off.

We demonstrated a simple reward heuristic to maximise discovery of new execution paths, and showed how this could be encoded in MCTS’s traditional UCT score function. We evaluated LEGION, both to measure its performance against its peers and to understand its sensitivity to hyperparameters. We found that LEGION was effective on a wide variety of benchmark programs without being overly sensitive to reasonable choices for its hyperparameters.

LEGION provides a new framework for investigating trade-offs between traditional testing approaches, and the incorporation of further statistical learning methods to assist automated test generation. Our results demonstrate the practicality and the promise of the ideas it embodies.

REFERENCES

- [1] Animesh Basak Chowdhury, Raveendra Kumar Medicherla, and Venkatesh R. 2019. VeriFuzz: Program Aware Fuzzing. In *Tools and Algorithms for the Construction and Analysis of Systems*, Dirk Beyer, Marieke Huisman, Fabrice Kordon, and Bernhard Steffen (Eds.). Springer International Publishing, Cham, 244–249.
- [2] D. Beyer. 2020. Second Competition on Software Testing: Test-Comp 2020. In *Proc. FASE (LNCS)*. Springer. https://www.sosy-lab.org/research/pub/2020-FASE.Second_Competition_on_Software_Testing_Test-Comp_2020.pdf
- [3] Dirk Beyer and Marie-Christine Jakobs. 2019. CoVeriTest: Cooperative verifier-based testing. In *International Conference on Fundamental Approaches to Software Engineering*. Springer, 389–408.
- [4] Dirk Beyer, Stefan Löwe, and Philipp Wendler. 2019. Reliable benchmarking: Requirements and solutions. *International Journal on Software Tools for Technology Transfer* 21, 1 (2019), 1–29.
- [5] Marcel Böhme, Van-Thuan Pham, Manh-Dung Nguyen, and Abhik Roychoudhury. 2017. Directed greybox fuzzing. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 2329–2344.
- [6] Marcel Böhme, Van-Thuan Pham, and Abhik Roychoudhury. 2017. Coverage-based greybox fuzzing as Markov chain. *IEEE Transactions on Software Engineering* 45, 5 (2017), 489–506.
- [7] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton. 2012. A Survey of Monte Carlo Tree Search Methods. *IEEE Transactions on Computational Intelligence and AI in Games* 4, 1 (March 2012), 1–43. <https://doi.org/10.1109/TCLIAIG.2012.2186810>
- [8] Cristian Cadar, Daniel Dunbar, Dawson R Engler, et al. 2008. KLEE: Unassisted and Automatic Generation of High-Coverage Tests for Complex Systems Programs.. In *OSDI*, Vol. 8. 209–224.
- [9] Marek Chalupa, Tomáš Jašek, Lukáš Tomovič, Martin Hruška, Veronika Šoková, Paulína Ayaziová, Jan Strejček, and Tomáš Vojnar. 2020. Symbiotic 7: Integration of Predator and More. In *Tools and Algorithms for the Construction and Analysis of Systems*, Armin Biere and David Parker (Eds.). Springer International Publishing, Cham, 413–417.
- [10] Peng Chen and Hao Chen. 2018. Angora: Efficient fuzzing by principled search. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 711–725.
- [11] R. Dutra, K. Laeufer, J. Bachrach, and K. Sen. 2018. Efficient Sampling of SAT Solutions for Testing. In *2018 IEEE/ACM 40th International Conference on Software Engineering (ICSE)*. 549–559. <https://doi.org/10.1145/3180155.3180248>
- [12] Shuitao Gan, Chao Zhang, Xiaojun Qin, Xuwen Tu, Kang Li, Zhongyu Pei, and Zuoning Chen. 2018. Collafl: Path sensitive fuzzing. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 679–696.
- [13] Patrice Godefroid, Nils Klarlund, and Koushik Sen. 2005. DART: Directed Automated Random Testing. In *Proceedings of the 2005 ACM SIGPLAN Conference on Programming Language Design and Implementation* (Chicago, IL, USA) (PLDI '05). Association for Computing Machinery, New York, NY, USA, 213–223. <https://doi.org/10.1145/1065010.1065036>
- [14] Patrice Godefroid, Michael Y. Levin, and David Molnar. 2012. SAGE: Whitebox Fuzzing for Security Testing. *Queue* 10, 1 (Jan. 2012), 20–27. <https://doi.org/10.1145/2090147.2094081>
- [15] Patrice Godefroid, Hila Peleg, and Rishabh Singh. 2017. Learn&fuzz: Machine learning for input fuzzing. In *2017 32nd IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 50–59.
- [16] Joxan Jaffar, Rasool Maghareh, Sangharatna Godbole, and Xuan-Linh Ha. 2020. TracerX: Dynamic symbolic execution with interpolation (competition contribution). In *International Conference on Fundamental Approaches to Software Engineering*. Springer, 530–534.
- [17] Levente Kocsis and Csaba Szepesvári. 2006. Bandit based Monte-Carlo planning. In *European Conference on Machine Learning*. Springer, 282–293.
- [18] Hoang M Le. 2020. Llvm-based hybrid fuzzing with LibKluzzer (competition contribution). In *International Conference on Fundamental Approaches to Software Engineering*. Springer, 535–539.
- [19] Nicholas Nethercote and Julian Seward. 2007. Valgrind: a framework for heavy-weight dynamic binary instrumentation. *ACM Sigplan notices* 42, 6 (2007), 89–100.
- [20] H. Peng, Y. Shoshitaishvili, and M. Payer. 2018. T-Fuzz: Fuzzing by Program Transformation. In *2018 IEEE Symposium on Security and Privacy (SP)*. 697–710.
- [21] Bob Joyce Reed Hastings. 1991. Purify: Fast detection of memory leaks and access errors. In *In Proc. of the Winter 1992 USENIX Conference*. Citeseer.
- [22] Sebastian Ruland, Malte Lochau, and Marie-Christine Jakobs. 2020. HybridTiger: Hybrid model checking and domination-based partitioning for efficient multi-goal test-suite generation (competition contribution). In *International Conference on Fundamental Approaches to Software Engineering*. Springer, 520–524.
- [23] Andrey Ryabinin. 2014. UBSan: run-time undefined behavior sanity checker. *E-mail publié sur la liste de développement du noyau Linux* 20 (2014).
- [24] Kostya Serebryany. 2015. libFuzzer—a library for coverage-guided fuzz testing. *LLVM project* (2015).
- [25] Konstantin Serebryany, Derek Bruening, Alexander Potapenko, and Dmitriy Vyukov. 2012. AddressSanitizer: A fast address sanity checker. In *Presented as part of the 2012 {USENIX} Annual Technical Conference ({USENIX}{ATC} 12)*. 309–318.
- [26] Yan Shoshitaishvili, Ruoyu Wang, Christopher Salls, Nick Stephens, Mario Polino, Andrew Dutcher, John Grosen, Siji Feng, Christophe Hauser, Christopher Kruegel, et al. 2016. Sok: (State of) The Art of War: Offensive techniques in binary analysis. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, 138–157.
- [27] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529, 7587 (2016), 484.
- [28] Nick Stephens, John Grosen, Christopher Salls, Andrew Dutcher, Ruoyu Wang, Jacopo Corbetta, Yan Shoshitaishvili, Christopher Kruegel, and Giovanni Vigna. 2016. Driller: Augmenting Fuzzing Through Selective Symbolic Execution.. In *NDSS*, Vol. 16. 1–16.
- [29] István Szita, Guillaume Chaslot, and Pieter Spronck. 2010. Monte-Carlo Tree Search in Settlers of Catan. In *Advances in Computer Games*, H. Jaap van den Herik and Pieter Spronck (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 21–32.
- [30] Guy Van den Broeck, Kurt Driessens, and Jan Ramon. 2009. Monte-Carlo tree search in poker using expected reward distributions. In *Asian Conference on Machine Learning*. Springer, 367–381.
- [31] Xinyu Wang, Jun Sun, Zhenbang Chen, Peixin Zhang, Jingyi Wang, and Yun Lin. 2018. Towards Optimal Concolic Testing. In *Proceedings of the 40th International Conference on Software Engineering (Gothenburg, Sweden) (ICSE '18)*. Association for Computing Machinery, New York, NY, USA, 291–302. <https://doi.org/10.1145/3180155.3180177>
- [32] C.D. Ward and Peter Cowling. 2009. Monte Carlo search applied to card selection in Magic: The Gathering. 9 – 16. <https://doi.org/10.1109/CIG.2009.5286501>
- [33] Xiaofei Xie, Lei Ma, Felix Juefei-Xu, Minhui Xue, Hongxu Chen, Yang Liu, Jianjun Zhao, Bo Li, Jianxiong Yin, and Simon See. 2019. DeepHunter: A Coverage-Guided Fuzz Testing Framework for Deep Neural Networks. In *Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis (Beijing, China) (ISSTA 2019)*. Association for Computing Machinery, New York, NY, USA, 146–157. <https://doi.org/10.1145/3293882.3330579>
- [34] C. Yeh, H. Lu, J. Yeh, and S. Huang. 2017. Path Exploration Based on Monte Carlo Tree Search for Symbolic Execution. In *2017 Conference on Technologies and Applications of Artificial Intelligence (TAAL)*. 33–37. <https://doi.org/10.1109/TAAL.2017.26>
- [35] Insu Yun, Sangho Lee, Meng Xu, Yeongjin Jang, and Taesoo Kim. 2018. {QSYM}: A practical concolic execution engine tailored for hybrid fuzzing. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*. 745–761.
- [36] Michal Zalewski. 2017. American fuzzy lop (AFL) fuzzer.
- [37] Lei Zhao, Yue Duan, Heng Yin, and Jifeng Xuan. 2019. Send Hardest Problems My Way: Probabilistic Path Prioritization for Hybrid Fuzzing.. In *NDSS*.