

AUSTRALIAN NATIONAL UNIVERSITY

DOCTORAL THESIS

State Space Search in Fuzzing

Author:

Adrian HERRERA

Panel:

Dr. Antony L. HOSKING,
Dr. Mathias PAYER,
Dr. Michael NORRISH

*A thesis submitted for the degree of Doctor of Philosophy
of The Australian National University*

November 17, 2024

© Copyright by Adrian HERRERA 2024
All Rights Reserved

Declaration of Authorship

I, Adrian HERRERA, declare that this thesis titled, "State Space Search in Fuzzing" and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed: *Adrian Herrera*

Date: 19/11/2024

"Space is big. You just won't believe how vastly, hugely, mind-boggling big it is."

Douglas ADAMS

Acknowledgements

I am eternally indebted to my panel: Tony Hosking, Mathias Payer, and Michael Norrish. Each of them brought their own special *something* to my PhD journey. Tony's boundless optimism kept me going through all of those paper rejections. Mathias' technical chops are unparalleled, and he always provided insightful comments and feedback on my work. I am forever grateful for his unwavering dedication and support for a student on the other side of the world. Michael unquestionably made me a better writer. Those who say "*you can't polish a turd*" have not witnessed Michael's formidable wordsmithing abilities.

I am grateful to my friends and former colleagues at the *Defence Science and Technology Group* (DSTG) for providing a workplace where undertaking further study is not only supported, but actively encouraged. In particular: Shane Magrath, for encouraging me to do this in the first place; Jon Arnold, Nick McConnell, Adrian Caldwell, Julian Costa, and Gareth Parker, for supporting me while I context-switched between managing and PhD-ing; Ben Cheney, Danny Coscia, Josh Green, and Suneel Randhawa, for getting me into cyber security research; Jamie Harwood, David Jamieson, Vi Li, Chenni Shettigara, and Nicole Stephens, for allowing us DSTG interlopers to work in the *Australian Signal's Directorate's* (ASD) CoLab and being so welcoming; and finally the *Software Systems Security* (S3) team: Luke Croak, Lauren Nelson-Lee, James Shaker, and Melissa Turner, for keeping me sane (primarily by distracting me with management-y things).

Every morning—modulo a global pandemic—I would eagerly join my *Australian National University* (ANU) colleagues on campus at the *Little Pickle* café. In particular, I thank Steve Blackburn, Thomas Haines, and Peter Höfner for the chats about academia, research, and life. Thanks also to Luke for keeping me caffeinated ☕.

In addition to drinking coffee at the Little Pickle, I had the pleasure of working with several bright and inspiring ANU staff and students. I thank: Zixian Cai, for not only maintaining the Moma computing cluster and *not* crashing my car when learning to drive, but for being an impressive and inspiring young systems researcher; Joshua Corner, Kathryn Gill, and Amy Samson, for being smart, dedicated, and—above all—*fun* students; Hendra Gunadi, for his early experimental work on the *MoonLight* project, paving the way for the "Seed Selection for Successful Fuzzing" paper; and finally, Alwen Tiu, for trusting me to help develop and teach COMP2710 *Software Security* (and then asking me to return—multiple times—when it was rebadged as COMP3703).

Outside of ANU and DSTG, I had the opportunity to meet and work with some

super-smart and cool people. Arlen Cox and Andrew Ruef were always up for a chat, providing feedback on my papers, ideas, and sharing their (impressive) knowledge on all things program analysis and Datalog. I look forward to future discussions on the latest-and-greatest program analysis research.

I was fortunate to be on the organizing committee for the 9th *International Summer School on Information Security and Protection (ISSISP)* in Canberra. *Trail of Bits* generously sponsored this even and sent out Peter Goodman to give a workshop on binary lifting. Peter and I have kept in regular contact since; hacking on PASTA, lifting binaries to LLVM, and talking all things program analysis. Peter was also generous with his time and gave me invaluable feedback on my work. I hope to keep these conversations going well into the future.

Despite not physically being in Switzerland, I always felt like a member of the *HexHive* lab. In particular, the Slack workspace was a place of vibrant discussion, lively debate, and sharing of knowledge. I thank all of the HexHive crew; particularly Atri Bhattacharyya, Priyam Biswas, Manuel Egele, Ahmad Hazimeh, Yuseok Jeon, Zhiyuan Jiang, Hui Peng, Andrés Sanchez, Prashast Srivastava, Antony Vennard, and Chibin Zhang, for their passes on my papers and discussions on my research ideas.

Marcel Böhme greatly influences how I think about fuzzing. In particular, his statistical modeling of fuzzing has brought rigor to what I once considered a “wild-west” research field. The fuzzing research community is much healthier with people like Marcel.

As my PhD journey drew to a close, I started looking for the next adventure and joined the team at *Interrupt Labs*. Interrupt Labs have encouraged me to take the skills and techniques I developed during my PhD and apply them to “real-world” problems. I am very fortunate to work with a great team of world-class vulnerability researchers, and I thank James Loureiro, Mark Barnes, and Mat Ramsey for having me onboard.

I feel incredibly lucky to have worked with (and learned from) the team at *Infosect*, particularly Kylie McDevitt and Silvio Cesare. Not only are Kylie and Silvio first-class hackers, but they are also incredibly kind and supportive people. I thank them for allowing me to speak at their conferences and inviting me to their events (in particular, their excellent Christmas parties). I plan to attend many more of their conferences!

I thank my Mum, Julie Herrera; Dad, Carlos Herrera; Nan, Nancy Gallagher (RIP); and brother, Daniel Herrera. Despite not really understanding the technical intricacies of fuzzing, they always supported what I was doing. They encouraged me to study

what I wanted to study (at both school and university), ultimately positioning me where I am today. Thank you.

I met my partner Jessica Hughes halfway through my PhD. She made the second half of this PhD so much easier than it should have been; supplying me with not only coffee, icecream, and trashy reality TV shows, but unflinching support, time, and space for me to do my research. Meeting Jess also meant becoming a “bonus dad” to her two girls, Fleur and Willa. Fleur and Willa also provided a welcome distraction from the depths of this dissertation, albeit a very different kind of distraction! I love you all ♡.

AUSTRALIAN NATIONAL UNIVERSITY

Abstract

Doctor of Philosophy

State Space Search in Fuzzing

by Adrian HERRERA

Coverage-guided fuzzers are an indispensable tool in the software-testing toolbox. They uncover bugs in a target program by subjecting it to a large number of automatically-generated inputs. The fuzzer generates these inputs to search the corners of the target's (potentially vast) *state space*, where bugs are more likely to lurk. While fuzzers have successfully uncovered bugs in a range of targets, they struggle to discover "deep bugs" (i.e., bugs triggered under a complex set of control and data dependencies). Moreover, security professionals deploying fuzzers lack observability into fuzzers' state space search (and thus an understanding of *why* these deep bugs are missed).

This dissertation presents a set of techniques for reasoning about and improving a fuzzer's state space search, ultimately enhancing a fuzzer's bug-finding ability.

First, we consider the task of bootstrapping this search process. Fuzzers typically require an initial set of *seeds* (exemplar inputs accepted by the target) to kickstart their state space search. We empirically evaluate various methods for selecting these seeds, designing an optimal technique for reducing large seed sets in the process.

Second, we develop a new state space abstraction. Fuzzers traditionally abstract a target's state space based on *control-flow* features. We present an abstraction based on *data-flow* features and demonstrate how our data-flow-based abstraction uncovers bugs that traditional fuzzers fail to find.

Finally, we investigate how best to measure a fuzzer's state space search after a fuzzing campaign. We use static analysis to quantify a target's state space, allowing us to measure *how much* of this state space a fuzzer has explored. We empirically evaluate several modern static analysis frameworks and propose new approaches for assessing a fuzzer's state space search.

Contents

Declaration of Authorship	iii
Acknowledgements	vii
Abstract	xi
1 Introduction	1
1.1 Thesis	3
1.2 Motivating Examples	4
1.3 Structure and Contributions	6
1.3.1 Publications	7
2 Background and Related Work	9
2.1 Introduction	9
2.2 Software Testing	9
2.3 Fuzz Testing	10
2.3.1 Blackbox Fuzzers	11
2.3.2 Greybox Fuzzers	11
2.3.3 Whitebox Fuzzers	11
2.4 Coverage-guided Greybox Fuzzing	12
2.5 Related Work	14
2.5.1 Benchmark Suites	14
2.5.2 Preprocessing	15
2.5.3 Seed Selection	15
2.5.4 Seed Scheduling	16
2.5.5 Energy Assignment	17
2.5.6 Mutation	17
2.5.7 Execution	18
2.5.8 Crash Detection	18
2.5.9 Stopping Condition	18
2.6 Ensuring Statistically-Sound Results	19
2.6.1 Repeated, Independent Trials	19

2.6.2	Bug Survival Time	19
2.6.3	Summary Statistics	20
2.6.4	Statistical Significance	20
2.7	Chapter Summary	20
3	Coverage Metrics	21
3.1	Introduction	21
3.2	Coverage Metrics in Theory and Practice	22
3.2.1	Basic Block	22
3.2.2	Edge	22
3.2.3	Path	24
3.2.4	Calling Context	25
3.2.5	Memory Access	25
3.2.6	Comparison Operator	26
3.2.7	Program Variable	26
3.3	Chapter Summary	27
4	Seed Selection	29
4.1	Introduction	29
4.2	Seed Selection Practices	30
4.2.1	An Ideal Initial Corpus	30
4.2.2	In Experimental Evaluation	30
4.2.3	In Deployment	33
4.3	Corpus Minimization	35
4.3.1	Prior Work	36
4.3.2	OPTIMIN	37
4.4	Evaluation	39
4.4.1	Methodology	41
	Fuzzer Selection	41
	Target Selection	41
	Sample Collection	42
	Experimental Setup	43
	Experiment	43
4.4.2	Minimization (RQ 1)	45
4.4.3	Bug Finding (RQ 2)	46
	FTS Coverage Benchmarks	46
	The EMPTY Seed	49
	PROVided Seeds	49
	Iteration Rates	50

	Comparison to Previous Magma Evaluations	51
	CVEs	51
4.4.4	Code Coverage (RQ 3)	52
4.4.5	Discussion	52
4.5	Chapter Summary	54
5	Data-Flow-Guided Fuzzing	57
5.1	Introduction	57
5.2	Motivation	57
5.3	Design of a Data-Flow-Guided Fuzzer	58
5.3.1	Coverage Sensitivity	58
	Def Site Sensitivity	59
	Use Site Sensitivity	61
	Composing Sensitivity Lattices	62
5.4	DATAFLOW Implementation	63
5.4.1	<i>Def-use</i> Site Identification	63
5.4.2	<i>Def-use</i> Tracking	64
	<i>Def</i> Site Instrumentation	64
	<i>Use</i> Site Instrumentation	65
5.4.3	Fuzzer Integration	68
5.4.4	Limitations	69
	<i>Def</i> Site Selection	69
	Custom Memory Allocators	69
	C++ Dynamic Memory Allocation	69
	Coverage Imprecision	70
5.5	Evaluation	70
5.5.1	Methodology	70
	Fuzzer Selection	70
	Target Selection	71
	Experimental Setup	71
5.5.2	Run-time Overheads (RQ 1)	72
5.5.3	Bug Finding (RQ 2)	74
5.5.4	Coverage Expansion (RQ 3)	78
	Accounting for Iteration Rates	83
5.5.5	Characterizing Data-Flow (RQ 4)	84
5.5.6	Discussion	86
5.6	Future Work	87
5.7	Chapter Summary	88

6	Quantifying State Space Search	89
6.1	Introduction	89
6.2	Program Analysis for State Space Search	90
6.2.1	Static Analysis	90
6.2.2	Dynamic Analysis	91
6.2.3	Implementation	91
6.3	Evaluation	92
6.3.1	Methodology	92
	Target Selection	92
	Experimental Setup	92
6.3.2	Static Analysis Performance (RQ 1)	93
	False Negatives	95
6.4	Discussion	97
6.5	Future Work	97
6.6	Chapter Summary	98
7	Comparison of Coverage Metrics	99
7.1	Introduction	99
7.2	Motivation	99
7.3	Evaluation	101
7.3.1	Methodology	101
	Fuzzer Selection	101
	Target Selection	102
	Experimental Setup	102
7.3.2	Comparison of Coverage Metrics (RQ 1)	102
	Context-Insensitive Edge Coverage	103
	Context-Sensitive Edge Coverage	105
7.3.3	Complementary Metrics (RQ 2)	106
7.4	Comparison to Previous Studies	111
7.5	Future Work	130
7.6	Chapter Summary	132
8	Conclusions and Reflections	133
8.1	Reflections	134
	Bibliography	137

List of Figures

1.1	High-level overview of a fuzzing campaign.	4
1.2	A simple dispatch table.	5
2.1	Examples of software testing techniques.	10
3.1	Greybox fuzzer genealogy.	23
3.2	Qualitative comparison of fuzzers' coverage metrics.	28
4.1	Code coverage for readelf.	34
4.2	Corpus minimization example.	38
5.1	<i>Def</i> and <i>use</i> site sensitivity lattices.	59
5.2	The simple dispatch table.	60
5.3	High-level overview of DATAFLOW.	63
5.4	Padding Area MetaData (PAMD) metadata.	65
5.5	DATAFLOW-instrumented dispatch table <i>def</i> site.	66
5.6	DATAFLOW-instrumented dispatch table <i>use</i> site.	67
5.7	Magma LUA003 bug.	78
5.8	Control-flow coverage.	81
5.9	Data-flow coverage.	82
7.1	AFL++ coverage expansion.	100
7.2	<i>lcms</i> context-insensitive probability maps.	108
7.3	<i>lcms</i> context-sensitive probability maps.	109
7.4	<i>sqlite3</i> context-insensitive probability maps.	110
7.5	Predicted improvements to control-flow coverage.	112

List of Tables

2.1	Fuzzer benchmark comparison.	16
2.2	Statistical techniques summary.	19
4.1	Past fuzzing evaluation, focusing on seed selection.	31
4.2	OPTIMIN minimization example.	40
4.3	Seed selection evaluation targets.	42
4.4	Corpora sizes.	45
4.5	Bug-finding results.	47
4.6	Assigned Common Vulnerabilities and Exposures (CVEs).	48
4.7	Coverage results.	53
5.1	Fuzzer configurations.	71
5.2	DDFuzz target dataset.	72
5.3	SPEC CPU2006 overheads.	73
5.4	Bug-finding results.	75
5.5	Coverage results.	80
5.6	Coverage results (last exec of the slowest fuzzer).	84
5.7	DD ratios.	85
6.1	FUZZBENCH target statistics.	93
6.2	Static analysis results.	94
6.3	Covered control-flow elements.	95
6.4	Static analysis false negatives.	96
7.1	Evaluated fuzzers.	102
7.2	Coverage results.	104

List of Acronyms

ASAN AddressSanitizer.

UBSAN UndefinedBehaviorSanitizer.

AFL American Fuzzy Lop.

BFF Basic Fuzzing Framework.

CFG control-flow graph.

CGC Cyber Grand Challenge.

CI confidence interval.

CNF conjunctive normal form.

CVE Common Vulnerabilities and Exposures.

DDG data dependency graph.

DTA dynamic taint analysis.

FTS Fuzzer Test Suite.

ICFG interprocedural control-flow graph.

ILP integer linear programming.

IR intermediate representation.

JIT just in time.

LoC lines of code.

LTO link-time optimization.

MaxSAT maximum satisfiability.

NSS Network Security Services.

OS operating system.

PAMD Padding Area MetaData.

PCA Principal Component Analysis.

PCSG probabilistic context-sensitive grammar.

RMST restricted mean survival time.

SAT boolean satisfiability.

SDLC software development lifecycle.

STADS software testing and analysis as discovery of species.

t-SNE t-distributed stochastic neighbor embedding.

To Jess, Fleur, and Willa.

Chapter 1

Introduction

Modern software systems are growing exponentially larger and more complex. At the same time, society is becoming more dependent on these software systems. For example, the Linux kernel—which powers mobile phones, routers, servers, and more—has grown from 13 million lines of code (LoC) in 2016 to over 20 MLoC in 2021 (a 54 % increase in five years). This increased size and complexity results in a greater likelihood of bugs and security vulnerabilities simply through the added code. Moreover, efficiently analyzing these codebases to find bugs and security vulnerabilities is becoming increasingly difficult (e.g., due to their large size; complex interactions between multiple, heterogeneous components).

An abundance of automatic bug-finding techniques have been developed in an attempt to keep up with the growing size and complexity of modern codebases. These techniques are either *static* or *dynamic*; the former analyzes code without executing it, while the latter observes executed code. Static analyses can reason over *all* possible behaviors that may arise at run time, whereas dynamic analyses can only reason about behaviors observed over a (finite) set of inputs. However, Rice’s theorem [180] means that non-trivial semantic questions a static bug finder would like answered (e.g., “*is an arbitrary program free of security vulnerabilities?*”) are undecidable. This forces static analyses to be conservative, leading to false positives (i.e., claiming a bug exists when it does not) and false negatives (i.e., misdetecting a bug). Unfortunately, static analyses have garnered a reputation for being more harmful than helpful [45, 96, 109] (e.g., due to the prohibitively high number of false positives/negatives some static bug finders emit) and difficult to scale [54, 182]. This has led to a renewed interest in dynamic analyses, with a particular focus on *fuzz testing*.

Fuzz testing (“fuzzing”) has become the de facto dynamic analysis for discovering bugs and vulnerabilities in large, complex codebases. In contrast to other bug-finding and verification techniques—e.g., model checking [106], symbolic execution [14], and static analysis [11]—fuzzing is (a) highly scalable, (b) free of false positives, and

(c) unencumbered by models of the external environment. These advantages have led to the widespread adoption of fuzzing, and the subsequent discovery of tens of thousands of bugs in popular programs [3, 31, 33, 52, 181]. For example, Google’s OSS-Fuzz service has helped fix over 8,800 vulnerabilities and 28,000 bugs across 850 open-source projects [31].

A fuzzer finds bugs by executing the target codebase (the “target”) multiple times with a large number of automatically generated (and possibly malformed) inputs. These inputs are generated to maximize exploration of the target’s *state space*. A program *state* is a snapshot of memory and registers, as well as relevant aspects of the operating system (OS) affecting the program (e.g., file descriptors, network sockets) [7, 183]. The target’s state space is thus the set of all possible states a program can be in, and the target’s behaviors are described by paths through this state space. Intuitively, exploration of the target’s state space correlates with bug discovery [23]; after all, you can only trigger a bug in code executed.

Early fuzzers treated the target as a black box. In particular, these fuzzers had no internal view of the target’s state space; their only feedback was whether the target crashed when executing a given input. In contrast, modern fuzzers typically use lightweight instrumentation to track which parts of the target’s state space they have explored. This information is fed back to the fuzzer, helping to guide its state space search.

Unfortunately, working with the previous definition of a target’s state space is unwieldy and quickly makes any program analysis (static or dynamic) intractable. In particular, it is impractical for a fuzzer to measure a state with this level of detail; prior work has shown the importance of maintaining a fuzzer’s *iteration rate* (the number of inputs the target executes per unit of time), and this level of detail will significantly reduce this iteration rate, harming fuzzing outcomes.

Abstraction is a solution to this problem, improving an analysis’ tractability [113]. In fuzzing and symbolic execution, *control* abstraction is common, while *data* abstraction is customary in abstract interpretation and model checking. In the former, a target program’s state space is described by the set of computation sequences and control flow; e.g., basic blocks/edges in the target’s control-flow graph (CFG) [8, 9, 16, 19, 21, 36, 37, 39, 40, 51, 61, 68, 142, 167]. In the latter, the target’s state space is described by interpretations (i.e., valuations) of variables [48, 56, 106]. These abstractions provide complementary views of the target’s state space.

Despite these abstractions, even trivially small programs can have innumerably

large state spaces.¹ This is particularly problematic for coverage-guided greybox fuzzers—the focus of this dissertation—which randomly walk the target’s state space.² Notably, the fuzzer must satisfy increasingly complex control- and data-flow-based dependencies to explore every corner of the target’s state space. Unfortunately, despite improvements in fuzzing techniques, fuzzers still have difficulty solving complex control- and data-flow constraints, leading to “deep” bugs continuing to evade long-running fuzzing campaigns. Moreover, security professionals deploying fuzzers lack *observability* into the fuzzer’s state space search, making it difficult to understand *why* these bugs remain undiscovered.

1.1 Thesis

This dissertation describes techniques, frameworks, and methodologies for measuring and improving a fuzzer’s ability to explore a target program’s state space. My thesis is:

Thesis statement

Fuzzing outcomes (i.e., bug discovery) are enhanced by improving fuzzers’ state space search.

We present a high-level overview of a typical fuzzing campaign in Fig. 1.1. Given the target program to fuzz, a set of well-formed, exemplar inputs (of the format accepted by the target) are sourced (①). As previously described, modern fuzzers instrument the target to track state space exploration (②). Following this, the fuzzer is run (③), mutating inputs to explore the target’s state space.³ Finally, the results are analyzed after completing the campaign (④). This analysis includes triaging crashes (to determine a bug’s root cause) and examining which parts of the state space were explored.

We can improve a fuzzer’s state space search in several ways. First, rather than starting the search “from scratch”, we can cover a subset of the search space before beginning a fuzzing campaign (i.e., at ①). Second, we can avoid revisiting previously-visited states, giving the fuzzer more time to discover new states (during ③). Third, we can adopt different views of a target’s state space (at ②), revealing states that may not be explicit in conventional abstractions. Finally, we can use more powerful

¹I.e., the state space cannot be enumerated before the heat death of the universe.

²We leave details on how coverage-guided fuzzers operate to Chapter 2.

³We use the rabbit icon as a homage to the American Fuzzy Lop (AFL) fuzzer, named after the rabbit breed.




FIGURE 1.1: High-level overview of a fuzzing campaign.

techniques to gain deeper insight into a fuzzer’s state space search (at ④). This dissertation explores all of these improvements. Before outlining our specific contributions (Section 1.3), we provide several examples to motivate our work.

1.2 Motivating Examples

CVE-2021-43527 [158] is a heap buffer overflow in Mozilla’s Network Security Services (NSS) library that occurs when handling particular cryptographic signatures. While the bug itself is unremarkable, the reasons why it remained undiscovered in widely-fuzzed code for almost ten years [162] makes it a motivating example for this dissertation.

How did this bug remain undiscovered for so long? One reason was “*misleading [coverage] metrics*” [162]. A coverage metric guides a fuzzer’s state space search by abstracting the target’s state space (we present a more thorough discussion of fuzzer coverage metrics in Chapter 3). Moreover, fuzzers are commonly evaluated and compared by how much of the target’s state space they explore (under a given coverage metric). Despite “*good test coverage for the vulnerable areas*” [162], the bug was only found when fuzzing with an “unconventional” coverage metric (specifically, context-sensitive edge coverage, discussed further in Section 3.2.4).


Similarly, we use the program in Fig. 1.2 (adapted from Ormandy [162]) to demonstrate the need to rethink fuzzers’ state space search. The bug at line 3 (in Fig. 1.2a) is only triggered when `foo` is called directly from `main` (i.e., when `data[i] == 0x66` at line 34); `buf` cannot overflow if `bar` or `baz` are called (lines 4 and 5, respectively). However, this is problematic for fuzzers driven by control-flow coverage alone. Consider a function call graph abstraction of the target’s state space (Fig. 1.2b) and two inputs (byte arrays represented as integer sequences, read from the file descriptor `fd` at line 27): $\iota_1 = \{0xa, 0xb, 0xc, 0x61, 0x40, 0x40\}$ and $\iota_2 = \{0xa, 0xb, 0xc, 0x66, 0x40, 0x40\}$. Under this abstraction, ι_1 ’s coverage is a superset

```

1  static char buf[128];
2
3  void foo(int a, size_t b) { memset(buf, a, b); }
4  void bar(int a, size_t b) { foo('A', sizeof(buf)); }
5  void baz(int a, size_t b) { bar(a, sizeof(buf)); }
6
7  const struct handler_t {
8      char code;
9      void (*handler)(int, size_t);
10 } handlers[] = {{0x66, foo}, {0x61, bar}, {0x7a, baz}};
11
12 int sig_cmp(void *data, size_t start, size_t len) {
13     if (start + 3 > len)
14         return 1;
15
16     static const char sig[3] = {0xa, 0xb, 0xc};
17     return memcmp(&data[start], sig, 3);
18 }
19
20 int main(int argc, char *argv[]) {
21     struct stat st;
22
23     int fd = open(argv[1], O_RDONLY);
24     fstat(fd, &st);
25     size_t size = st.st_size;
26     char *data = malloc(size);
27     read(fd, data, size);
28
29     if (sig_cmp(data, 0, size))
30         return 1;
31
32     for (unsigned i = 4; (i + 3) <= size; i += 3)
33         for (unsigned j = 0; j < 3; ++j)
34             if (data[i] == handlers[j].code)
35                 handlers[j].handler(data[i + 1], data[i + 2]);
36
37     free(data);
38     close(fd);
39
40     return 0;
41 }

```

(A) Source.



(B) Call graph. The dashed lines indicate indirect calls that are resolved at run time.

(C) Call graph for the input $\iota_1 = \{0xa, 0xb, 0xc, 0x61, 0x40, 0x40\}$ (covered nodes are green).

(D) Call graph for the input $\iota_2 = \{0xa, 0xb, 0xc, 0x66, 0x40, 0x40\}$.

FIGURE 1.2: A simple dispatch table, adapted from Ormandy [162].

of ι_2 's coverage (see Figs. 1.2c and 1.2d).⁴ This means that even before fuzzing begins, *corpus minimization* (the process of minimizing the number of inputs that maximize coverage of the target's state space; discussed further in Chapter 4) would discard inputs covering only foo in favor of inputs covering bar and baz (because covering bar or baz also covers foo).

Observation

Conventional fuzzer coverage metrics (e.g., those based off of control flow) may lead to inaccurate and misleading measurements of program state.

We also use Fig. 1.2 to demonstrate the need to avoid starting a fuzzing campaign “from scratch”. In particular, `sig_cmp` (line 12) rejects inputs not starting with the signature `{0xa, 0xb, 0xc}`. While this is a relatively simple check to overcome, one can imagine more intricate file formats with more complex checks (e.g., a JavaScript interpreter). In general, these checks are difficult for a fuzzer to “guess” (e.g., via random mutation). However, if the fuzzer is bootstrapped with knowledge of the input format (e.g., with exemplar inputs starting with the three-byte signature), then the fuzzer can spend more time exploring the handler (rather than wasting cycles overcoming the signature check).

Observation

A fuzzer's ability to reach “deep” states in a target may depend on how the fuzzer's search was bootstrapped.

These examples demonstrate the need to rethink and improve fuzzers' state space search. This includes appropriately bootstrapping the state space search, adopting coverage metrics that go beyond traditional control-flow-based metrics, and improving insights into what states the fuzzer has covered. The following section outlines how we achieve this.

1.3 Structure and Contributions

This dissertation describes how we reason about and improve the state space exploration capabilities of coverage-guided fuzzers. We start in Chapter 2 by providing the background material necessary to motivate and understand our contributions.

⁴The same is true for a CFG abstraction; we use call graphs to simplify presentation.

We also summarize the current state-of-the-art fuzzing research. The structure of the remaining chapters echoes a fuzzing campaign (Fig. 1.1). In these chapters we make the following contributions:

Systematization of coverage metrics. In Chapter 3 we explore the different ways a fuzzer represents a program’s state space, presenting a systematization of coverage metrics and taxonomy of greybox fuzzers.

The impact of seed selection and optimal corpus minimization. In Chapter 4 we focus on the initial set of inputs required to bootstrap the fuzzing process (① in Fig. 1.1). We systematically (a) investigate how these initial input sets are constructed, and (b) evaluate how they affect a fuzzer’s ability to find bugs. We also present OPTIMIN, a fuzzing corpus minimization tool. Unlike other corpus minimizers, OPTIMIN is capable of deriving an *optimal* corpus by encoding corpus minimization as a boolean satisfiability (SAT) problem.

Data-flow-guided fuzzing. In Chapter 5 we present DATAFLOW, a greybox fuzzer guided by lightweight data-flow profiling. Unlike most greybox fuzzers (i.e., those discussed in Chapter 3), DATAFLOW’s instrumentation eschews control-flow coverage metrics in favor of *def-use chain* coverage (②).

Quantifying state space search. In Chapter 6 we look at measuring a fuzzer’s state space search after a fuzzing campaign (④). We use static analysis to quantify a target’s state space, aiming to produce a more-thorough analysis and evaluation of a fuzzer’s ability to explore a target’s state space.

Comparing coverage metrics. Finally, in Chapter 7 we conduct a large-scale fuzzing campaign to evaluate the full range of coverage metrics discussed throughout this dissertation.

We support our contributions with extensive empirical evaluation on real-world software. In total, we report on over 44 CPU-yr of fuzzing campaigns. Finally, we release as open source all code associated with our contributions (details are provided in each chapter).

1.3.1 Publications

While this dissertation is self-contained, its contributions draw from the following published material:

- Ahmad Hazimeh, **Adrian Herrera**, and Mathias Payer. *Magma: A Ground-Truth Fuzzing Benchmark* [84]. SIGMETRICS. 2020.
- **Adrian Herrera**, Hendra Gunadi, Shane Magrath, Michael Norrish, Mathias Payer, and Antony L. Hosking. *Seed Selection for Successful Fuzzing* [92]. International Symposium on Software Testing and Analysis (ISSTA). 2021.
- **Adrian Herrera**, Mathias Payer, and Antony L. Hosking. *Registered Report: DATAFLOW– Toward a Data-Flow-Guided Fuzzer* [91]. International Fuzzing Workshop (FUZZING). 2022.
- Zhiyuan Jiang, Shuitao Gan, **Adrian Herrera**, Flavio Toffalini, Lucio Romerio, Chaojing Tang, Manuel Egele, Chao Zhang, and Mathias Payer. *EVOCATIO: Conjuring Bug Capabilities from a Single PoC* [107]. Computer and Communications Security (CCS). 2022.
- **Adrian Herrera**, Mathias Payer, and Antony L. Hosking. *DATAFLOW: Toward a Data-Flow-Guided Fuzzer* [90]. Transactions on Software Engineering and Methodology (TOSEM). 2023.
- Simon Luo, **Adrian Herrera**, Paul Quirk, Michael Chase, Damith C. Ranasinghe, Salil S. Kanhere. *Make out like a (Multi-Armed) Bandit: Improving the Odds of Fuzzer Seed Scheduling with T-Scheduler* [132]. Asia Computer and Communications Security (AsiaCCS). 2024.

Chapter 2

Background and Related Work

2.1 Introduction

The term “fuzz testing” first appeared in an assignment in Prof. Barton Miller’s 1988 advanced operating system (OS) class¹ at the University of Wisconsin [146]. Since then, fuzz testing (“fuzzing”) as both a research field and industrial practice has progressed in leaps and bounds. Nowadays, you can browse the proceedings of the top security and software engineering conferences and find numerous fuzzing papers, or read countless blog posts about how software companies are scaling-up fuzzing to find bugs in their code (before it reaches production). Here, we distill this prior work and describe both fuzzing fundamentals and the current state-of-the-art.

Chapter outline. This chapter provides the background material required to motivate and understand our research contributions. We begin by describing fuzzing and how it relates to other software testing techniques. This is followed by an outline of the different “shades” of fuzzer: black-, grey-, and white-box (Section 2.3) and a more in-depth discussion of *coverage-guided greybox fuzzing* (Section 2.4), the focus of this dissertation. We discuss related prior work in Section 2.5 and conclude with a discussion on the experimental methodologies we use to ensure statistically-sound empirical evaluation in Section 2.6.

2.2 Software Testing

Most software requires testing to ensure it operates as intended. A primary goal of testing is to find bugs *before* the software is released to consumers. Figure 2.1 summarizes different testing techniques. *Static* techniques include formal verification (mathematically proving the system is correct with respect to a formal specification),

¹CS763, to be precise.

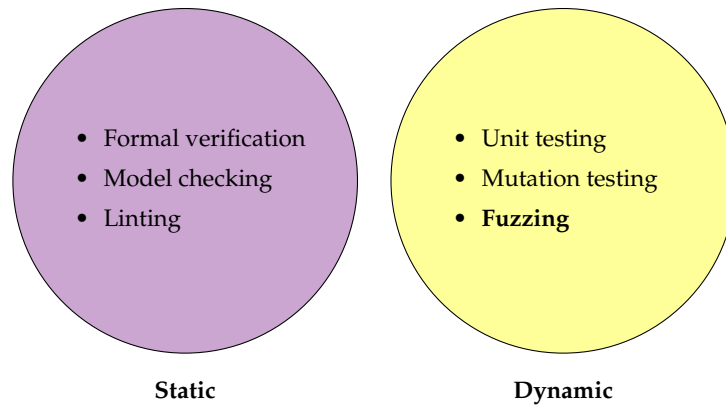


FIGURE 2.1: Examples of software testing techniques.

model checking (checking a finite-state model of the system meets a given specification), and linting (pattern matching against known bug classes and suspicious programming constructs). *Dynamic* techniques include unit tests (checking isolated code—typically at the function level—exhibits expected behavior), mutation testing (modifying the program—typically to cause an error—to ensure the test suite can detect the modifications), and fuzzing (the focus of this dissertation).

While it is often impossible to identify *all* bugs and errors in software before it is released, testing is an important component of the software development lifecycle (SDLC). In particular, testing reduces the risk of *security critical* bugs that can be exploited by attackers for nefarious purposes.

2.3 Fuzz Testing

Fuzzing is a dynamic software testing technique for finding bugs. While some forms of testing (e.g., unit testing) attempt to verify code behavior based on valid inputs (“positive” testing), fuzzing investigates code behavior on invalid or unexpected inputs (“negative” testing). Fuzzers do this by executing code with automatically-generated inputs and triggering crashes [138]. How a fuzzer produces these inputs depends on whether it is *generational* or *mutational*.

Generational fuzzers (e.g., QuickFuzz [77], NAUTILUS [8], Gramatron [199], and Favocado [53]) require a specification/model of the input format. They use this specification to synthesize inputs. In contrast, mutational fuzzers (e.g., libFuzzer [188], Angora [37], REDQUEEN [9] and INVSCOV [61]) use an initial corpus of seed inputs (e.g., files, network packets, and environment variables) to bootstrap the input-generation process. New inputs are then generated by mutating existing seeds in this corpus. Mutational fuzzers are more common than generational fuzzers, so when

discussing fuzzers we assume the mutational kind (rather than generational) unless otherwise stated.

In addition to how inputs are produced, fuzzers are further categorized by the “*granularity of semantics [i.e., behaviors] a fuzzer observes in each fuzz run*” [138]. These categories—black-, grey-, and white-box—are summarized below.

2.3.1 Blackbox Fuzzers

Blackbox fuzzers have no internal view of the target; they can only observe input/output relationships. Without any introspection, blackbox fuzzers are unencumbered by run-time overheads (e.g., associated with instrumentation) and thus achieve high *iteration rates* (i.e., the number of inputs executed per unit of time). However, they cannot measure/reason about how much/which parts of the target’s state space have been explored. Example blackbox fuzzers include Radamsa [87], the CERT Basic Fuzzing Framework (BFF) [196], and CodeAlchemist [82].

2.3.2 Greybox Fuzzers

A greybox fuzzer uses lightweight instrumentation (typically compiler-based instrumentation that avoids the need for heavyweight program analysis at run time) to collect run-time information about the target. This run-time information is fed back to the fuzzer to guide it toward unexplored parts of the state space (and, ultimately, uncover bugs). Example greybox fuzzers include American Fuzzy Lop (AFL) [231], NAUTILUS [8], and DDFuzz [141]. We discuss greybox fuzzing in greater detail in Section 2.4.

2.3.3 Whitebox Fuzzers

Whitebox fuzzers rely on heavyweight program analysis to guide target exploration. This program analysis is typically provided by a symbolic execution engine (“symbolic executor”) [14]. While capable of overcoming complex conditionals and other fuzzer roadblocks, symbolic executors incur a high run-time cost (e.g., the KLEE [27] and angr [194] symbolic executors are $\sim 3,000\times$ and $\sim 321,000\times$ slower than native execution, respectively [230]). Moreover, *path explosion*—where branches and symbolic memory accesses lead to a large number of program states that cannot be feasibly explored—makes systematically exploring a target’s state space intractable.

These performance and scalability issues have led to the emergence of concolic (concrete + symbolic) execution engines (“concolic executor”). Rather than exploring the target’s state space (and following multiple execution paths simultaneously), a

concolic executor traces the constraints for a single execution path and uses these constraints to generate new inputs that follow new paths (that deviate from the single execution path in a given branch). Example whitebox fuzzers include KLEE [27], *BuzzFuzz* [69], and SAGE [73]. Example concolic fuzzers include Driller [201], QSYM [230], and SYMCC [173].

2.4 Coverage-guided Greybox Fuzzing

Coverage-guided greybox fuzzers are the most pervasive type of fuzzer: they have successfully uncovered tens of thousands of bugs in a range of widely-used software [3, 33, 52, 181]. We describe their operation in greater detail in this section.

Algorithm 1: Mutational greybox fuzzing.

```

Input: Target  $\mathcal{P}$ , starting seeds  $\mathcal{S}$ 
Output: Crashing inputs  $\mathcal{S}_X$ 
1  $\mathcal{S}_X \leftarrow \emptyset, \mathcal{C}_S \leftarrow \emptyset$ 
2 if  $\mathcal{S} = \emptyset$  then /* Use the empty seed if no seeds are provided */
3    $\mathcal{S} \leftarrow \{\text{EmptySeed}\}$ 
4  $\mathcal{R} \leftarrow \text{Preprocess}(\mathcal{P})$ 
5 foreach  $\iota \in \mathcal{S}$  do /* Save starting seeds' coverage */
6    $\mathcal{C}_\iota \leftarrow \text{Execute}(\mathcal{P}, \iota)$ 
7    $\mathcal{C}_S \leftarrow \mathcal{C}_S \cup \mathcal{C}_\iota$ 
8 repeat /* The main fuzzing loop */
9    $\iota \leftarrow \text{ChooseNext}(\mathcal{S}, \mathcal{R})$ 
10   $e \leftarrow \text{AssignEnergy}(\iota, \mathcal{R})$ 
11  foreach  $i$  from 1 to  $e$  do
12     $\iota' \leftarrow \text{MutateInput}(\iota, \mathcal{R})$ 
13     $\mathcal{C}_{\iota'} \leftarrow \text{Execute}(\mathcal{P}, \iota')$ 
14    if CrashDetected then
15       $\mathcal{S}_X \leftarrow \mathcal{S}_X \cup \{\iota'\}$ 
16    else if IsInteresting( $\mathcal{C}_{\iota'}, \mathcal{C}_S, \mathcal{R}$ ) then
17       $\mathcal{S} \leftarrow \mathcal{S} \cup \{\iota'\}$ 
18     $\mathcal{C}_S \leftarrow \mathcal{C}_S \cup \mathcal{C}_{\iota'}$ 
19 until timeout reached or abort signal

```

Algorithm 1—adapted from Böhme, Pham, and Roychoudhury [22] and Böhme et al. [24]—provides an overview of a generic coverage-guided greybox fuzzer. The user provides an instrumented target program to fuzz, \mathcal{P} , and an optional set of starting inputs (“seeds”), \mathcal{S} . An “empty seed” is used if \mathcal{S} is not provided (lines 2 to 3). The target may first undergo a preprocessing/analysis stage (line 4) before the fuzzer enters the main fuzzing loop (lines 8 to 19), consisting of the following steps:

1. The fuzzer selects an input $\iota \in \mathcal{S}$ (line 9) and assigns it an *energy*, e (line 10). This energy “specifies the number of inputs to be generated from that seed [i.e., input]” [22].

The assignment of e depends on the fuzzer (e.g., AFL assigns energy based on the inputs execution time, coverage, and creation time).

2. The fuzzer mutates ι , producing ι' (line 12). These mutation operations can be *structure agnostic* or *structure aware*, depending on how much knowledge (if any) the fuzzer has of \mathcal{P} 's input format. Structure-agnostic mutations operate on the raw input string and include bit-flipping, byte substitution, and string concatenation. Structure-aware mutations are typically applied when \mathcal{P} 's input format is expressible by a *grammar*. In this case, mutations occur on an input's *parse tree*² and include node insertion/deletion, replacing nodes with those derived by a different set of grammar rules, and splicing parse trees together.
3. The target \mathcal{P} is executed with ι' (line 13). During execution, the fuzzer records coverage information and stores this information in a coverage map $C_{\iota'}$. This coverage approximates a program's state space and takes various forms (e.g., basic block, edges between basic blocks). For performance reasons, coverage is typically based on an approximation of *control-flow* information.
4. The fuzzer detects crashes and updates the global coverage map (lines 14 and 18). Crash-inducing inputs are saved for further (offline) triage and analysis. Further, an input is added to \mathcal{S} if the fuzzer considers it "interesting"; i.e., it satisfies some objective function (e.g., generates new coverage, gets "closer" to a target location). Otherwise, the input is discarded (line 16). If interesting, the coverage obtained by ι' is added to the global coverage map $C_{\mathcal{S}}$, which maintains coverage for *all* inputs.
5. Return to 1 (or terminate if the timeout is reached, line 19). The fuzzer may increase its aggressiveness (i.e., increase the energy e) on already-explored inputs.

The analysis report \mathcal{R} may guide seed selection, energy assignment, mutation, and/or assist in determining if ι' is interesting (Steps 1, 2 and 4, respectively). For example, directed fuzzers (e.g., AFLGo [24], CAFL [121], BEACON [99]) precompute a function call graph to guide the fuzzer toward specific target locations. An input is thus interesting if it reduces the "distance" to a target location.

The software testing and analysis as discovery of species (STADS) model [17] provides an alternative, probabilistic model of fuzzing. STADS is an ecology-inspired

²A parse tree—also known as a *derivation tree* or *concrete syntax tree*—represents the syntactic structure of ι according to a specific grammar.

statistical framework that enables accurate extrapolation of a fuzzer’s future performance (based on past performance). This is important to fuzzing because it “*provides statistical correctness guarantees with quantifiable accuracy*” [17].

Under the STADS model, greybox fuzzing is considered a random walk through \mathcal{P} ’s state space [22]. This walk is achieved by probabilistically sampling (with replacement) from \mathcal{P} ’s input space, which is divided into subdomains called *species*. These species are defined based on behavior observed when executing \mathcal{P} with input l' . Observed behaviors are approximated by one or more *coverage metrics* and are recorded in C_S . Thus, improving the state space exploration capabilities of a greybox fuzzer is equivalent to improving its *ability to discover new species* under STADS.

2.5 Related Work

Significant attention—from both academia and industry—has been given to many of the elements and steps in Algorithm 1. We summarize this work in the following sections.

2.5.1 Benchmark Suites

Appropriate benchmarks—from which \mathcal{P} is commonly sourced—are required for fair and meaningful fuzzer comparisons. Popular open-source software (e.g., GNU Binutils [72]) are frequently used to demonstrate a fuzzer’s ability to discover *zero-day* vulnerabilities (i.e., bugs that have not been patched). However, determining the root cause of a “real-world” zero-day is a time-consuming and largely manual process, increasing the time and effort required to evaluate and compare multiple fuzzers. To rectify this, benchmarks such as LAVA [55], the Cyber Grand Challenge (CGC) [131], UNIFUZZ [124], Fuzzer Test Suite (FTS) [75], FUZZBENCH [144], Magma [84], and REVBUGBENCH (from FIXREVERTER [234]) have been developed. These benchmarks evaluate and compare fuzzers across a range of performance metrics, including: coverage profiles (FUZZBENCH); synthetic bug counts (LAVA, CGC, and REVBUGBENCH); and real-world bug counts (UNIFUZZ, FTS, and Magma). Is there a preferred set of metrics for measuring fuzzing outcomes?

Code coverage has long been used as a proxy measure for a fuzzer’s bug-finding ability: a bug cannot be found in code never executed. Thus, it is common for fuzzer evaluations to use an independent code coverage metric to enable comparisons; commonly, lines of code (LoC) in the target (we revisit and expand on this approach in Chapter 7). However, controversy surrounds the strength of the correlation between

coverage profiles and bugs [102], with seemingly contradictory results across multiple works due to the incorrect use of statistical analyses (i.e., stratification) [42]. Even when the correlation is strong [23, 76, 117], the *agreement* on which fuzzer is superior may be weak; i.e., “the fuzzer best at achieving coverage may not be best at finding bugs” [23].

The potentially weak correlation between code coverage and bug finding has led to the development of “bug-based” benchmarks. A fuzzer’s ultimate goal is to find bugs, so why not measure this directly? Unfortunately, realistic bug-based benchmarks are challenging to build and maintain: they require ground-truth knowledge of bugs [84] and must be regularly updated to prevent overfitting [18, 23]. Moreover, the distribution of bugs may not be uniform or related to the actual possibility of bugs in the target [74].

The LAVA, CGC, and REVBUGBENCH benchmarks contain both a large number of programs and bugs; however, these bugs—and, in the case of the CGC, programs—are synthetic and are considered unrepresentative of “real-world” codebases (e.g., because of their small size). In contrast, Goerz et al. [74] use *mutation analysis* to assess a fuzzer’s performance, creating (faulty) variants of real-world programs. In contrast, Magma and UNIFUZZ contain both real-world targets and bugs, but the former contains only a small number of each, while the latter has no ground truth to validate results.

Table 2.1 summarizes our discussion on benchmark suites. We use both coverage- and bug-based benchmarks throughout this dissertation; primarily FUZZBENCH and Magma.

2.5.2 Preprocessing

Directed fuzzers (e.g., AFLGo [24], CAFL [121], and BEACON [99]) use a preprocessing stage to determine program locations to drive the fuzzer toward. This preprocessing stage computes a distance metric to the target location(s). The distance is calculated using static control- and data-flow analyses, with the result stored in \mathcal{R} .

We extend these static control- and data-flow analyses (and augment them with analogous dynamic analyses) in Chapters 6 and 7 to more accurately measure and reason about fuzzers’ state space exploration.

2.5.3 Seed Selection

Mutational fuzzers require an initial seed set to bootstrap the fuzzing process. Typically, these seed sets contain a small number of inputs that \mathcal{P} is known to accept (i.e.,

TABLE 2.1: Comparison of fuzzer benchmark suites. We characterize benchmarks across three dimensions: (i) the targets making up the benchmark; (ii) the bugs that exist across these targets (Goerz et al. [74] use mutants rather than bugs); and (iii) the coverage measure used. For both “targets” and “bugs” we count the number of targets/bugs (“#”) and classify their type as “Real” or “Synthetic”. Bug density is the mean number of bugs per target. Finally, ground truth may be (a) available (“✓”), (b) available but not easily accessible (“▶”), or (c) unavailable (“✗”).

Benchmark	Targets		Bugs				Coverage
	#	Type	#	Type	Density	Ground Truth	
LAVA-M	4	R	2,265	S	566.25	✓	–
CGC	131	S	590	S	4.50	▶	–
FTS	24	R	47	R	1.96	▶	Known source location
FUZZBENCH	32	R	8	R	0.25	▶	LoC
Magma	9	R	138	R	15.33	✓	–
UNIFUZZ	20	R	?	R	?	✗	LoC
REVBUGBENCH	10	R	7,910	S	791	✓	–
Goerz et al. [74]	7	R	141,278	S	18,754	✓	–
Open-source software	–	R	?	R	?	✗	LoC

the seeds are “well-formed” and do not trigger errors). Perhaps counter-intuitively, the “empty seed”—i.e., a zero-length seed—is also a popular starting seed because it “is an easy way to baseline a significant variable in the input configuration” [115].

In contrast, generational fuzzers generate their initial seeds. For example, QuickFuzz [77] generates seeds from file-format-handling libraries available in the Haskell programming language. Finally, some prior work combines both mutational and generational techniques. For example, Skyfire [215] learns a probabilistic context-sensitive grammar (PCSG) from both an initial seed set and ANTLR [169] grammar, and then uses the PCSG to generate uncommon inputs with diverse grammar structures (i.e., those structures with low probability). Similarly, SLF [228] is a “seedless fuzzer” capable of generating valid seeds from scratch using mutational fuzzing.

We analyze both (a) prior work in seed selection, and (b) the impact seed selection has on fuzzing outcomes in more detail in Chapter 4.

2.5.4 Seed Scheduling

“Interesting” inputs (i.e., leading to new coverage) are added to \mathcal{S} . Given \mathcal{S} grows rapidly, what is the best approach for selecting the next input to mutate in ChooseNext? To answer this question, CollAFL [67] introduced a suite of seed selection techniques prioritizing inputs which (a) exercise edges in $\mathcal{C}_{\mathcal{S}}$ with many unexercised *neighbor edges*, (b) exercise edges in $\mathcal{C}_{\mathcal{S}}$ with many unexercised *descendant paths*, and (c) exercise paths containing many memory access operations. TortoiseFuzz [217] similarly

prioritizes inputs exercising sensitive memory operations at different granularities (function, loop, and basic block levels) to improve the discovery of memory safety vulnerabilities. Finally, K -Scheduler [190] applies centrality measure from graph analysis to approximate the likelihood of executing new code by mutating a given input (and subsequently scheduling that input for mutation).

2.5.5 Energy Assignment

Once an input $\iota_i \in \mathcal{S}$ has been selected for mutation, how many times should this input be mutated before selecting a new input $\iota_j \in \mathcal{S}$? This count is determined by how much energy e is assigned to ι_i in `AssignEnergy`. Both AFLFast [22] and AFLGo [24] demonstrated improvements when more energy was assigned to inputs exercising low-frequency coverage elements in \mathcal{C}_S (compared to inputs exercising high-frequency elements). This was achieved by modeling coverage-guided fuzzing as a Markov chain. Similar improvements have been made using other statistical models (e.g., Multi-Armed Bandit [212, 229]).

2.5.6 Mutation

The fuzzer must first decide *where* to mutate. Random selection of input bits/bytes is highly inefficient, because most of the input may be unrelated to changes in \mathcal{P} 's behavior. Dynamic taint analysis (DTA) is commonly used to address this inefficiency (e.g., in VUzzer [176], Angora [37], and ParmeSan [163]). DTA improves mutation accuracy by tracking the flow of input bytes³ in \mathcal{P} , informing the fuzzer which bytes affect control flow (and hence which bytes to mutate). Unfortunately, accurate DTA is computationally expensive and requires significant manual effort (e.g., to specify taint policies) [47, 95, 192]. This has led to fuzzers avoiding DTA in favor of *approximate taint tracking*. For example, REDQUEEN [9] and AFL++ [63] use “input-to-state correspondence”, a technique based on the intuition that “*parts of the input directly correspond to the memory or registers at run time*” [9]. Similarly, GREYONE [68] infers taint by monitoring variable values during mutation. These approaches sacrifice precision for run-time performance.

After deciding where to mutate, the fuzzer must choose *how* to mutate. Following Section 2.4, this depends on how much knowledge the fuzzer has of the input structure. For example, NAUTILUS [8] uses grammar-aware mutations that operate on a parse tree representation of s . Grammartron [199] improves the performance of grammar-aware mutation by restructuring the parse tree as a *finite state automaton*.

³The exact granularity (e.g., bits, bytes) depends on the specific *taint policy*. The taint policy specifies the relationship between an instruction's input and output operand(s).

Fuzzers typically resort to bit flipping, byte substitution, and string concatenation when the input structure is unknown. Even if the complete input grammar/structure is unknown, fuzzers may be able to leverage “dictionaries” containing keywords from the input format (e.g., the PNG chunk specifiers IDAT, IEND, and PLTE). Rather than blindly substituting random bytes, the fuzzer can instead substitute keywords from the dictionary.

To alleviate the inefficiencies brought on by the randomness of the mutation process, MOPT [134] uses particle swarm optimization to derive an optimal probabilistically distribution from which a mutation operator is selected. Similarly, EMS [133] captures mutation strategies from previous fuzz runs and uses this history to find interesting inputs that trigger unique paths and crashes.

2.5.7 Execution

Blackbox fuzzers blindly execute \mathcal{P} without any feedback mechanism. In contrast, greybox fuzzers use positive feedback to drive the fuzzer toward unexplored parts of \mathcal{P} 's state space. A coverage map $C_{s'}$ is used to measure which parts of \mathcal{P} 's state space was executed with (mutated) input ι' .

We analyze different coverage metrics used by greybox fuzzers in Chapter 3. Moreover, we introduce a new coverage metric based on data flow coverage (specifically, *def-use* chains) in Chapter 5.

2.5.8 Crash Detection

The simplest form of crash detection is to rely on the underlying hardware/OS to report a violation (e.g., segmentation faults, divide-by-zero). Unfortunately, this approach may miss bugs (e.g., an out-of-bounds memory access may not necessarily trigger a segmentation fault). Thus, *sanitizers* are often deployed to detect a wider variety of bug classes more accurately. For example, AddressSanitizer (ASAN) [187] is used for detecting memory safety violations, UndefinedBehaviorSanitizer (UBSAN) [210] is used for detecting undefined behavior, and type confusion errors are detectable with HexType [105]. Song et al. [197] provide a more detailed summary of other sanitizers.

2.5.9 Stopping Condition

How long should a fuzzing campaign last? Klees et al. [115] found campaign timeouts ranging from 1 h to days and weeks, with 24 h being the most common campaign length. Further, Böhme, Liyanage, and Wüstholz [20] developed estimators for

TABLE 2.2: Techniques used to ensure statistically-sound results. “Summary” techniques are used to summarize a set of observations (i.e., over a set of N trials). “Comparison” techniques enable comparisons of observations.

	Bug finding	Coverage profiles
Summary	Restricted mean survival time (RMST)	Arithmetic mean
Comparison	log-rank test	Mann-Whitney U -test

determining when to stop a fuzzing campaign such that the *residual risk* of a bug remaining undiscovered is minimized.

2.6 Ensuring Statistically-Sound Results

Proper benchmarking is essential in systems security [118]. In particular, the highly-stochastic nature of fuzzing (e.g., due to random mutations [223]) requires the use of appropriate statistical techniques when analyzing results [5, 115]. We measure fuzzing outcomes across two axes: *bug finding* and *coverage profiles* (the two most common metrics used to analyze the results of a fuzzing campaign, per Section 2.5.1). We adopt the following approaches (and terminology) to ensure statistical significance when analyzing results across these two axes (summarized in Table 2.2).

2.6.1 Repeated, Independent Trials

We run one fuzzing *campaign* per {target \times fuzzer $\times y$ } combination, where y is a *dependent variable*⁴ that changes depending on the given evaluation (e.g., in Chapter 4 we examine the effect different seed selection strategies have on fuzzing outcomes; here, y is the set of seed selection strategies). Each campaign consists of N independent *trials* of length T hours.

2.6.2 Bug Survival Time

We statistically analyze and compare time-to-bug based on the recommendations of Böhme and Falk [19]. Following prior work [5, 211], we apply *survival analysis* [116] to time-to-bug events. This allows us to handle *censoring*: individual trials where a given bug is not found.

For each campaign (i.e., set of N repeated T hour trials) we model a bug’s *survival function*—the probability of a bug being found over time—using the Kaplan-Meier estimator [110]. Integrating this survival function with upper-bound T gives a bug’s

⁴An evaluation may have more than one dependent variable.

RMST for a given campaign. Shorter bug survival times indicate better performance. We report the RMST and 95 % confidence interval (CI) across each campaign.

We also use the *log-rank test* [140] to statistically compare bug survival times. The log-rank test is computed under the null hypothesis that two {target \times fuzzer \times y } combinations share the same survival function. Thus, we consider two fuzzing campaigns to have statistically-equivalent bug survival times if the log-rank test's p -value > 0.05 . We use the log-rank test because it is non-parametric and thus makes no assumptions about the underlying distribution of bugs.

2.6.3 Summary Statistics

We use both the arithmetic mean and 95 % CI to summarize non-bug-related fuzzing results (e.g., coverage profiles generated by different fuzzers). We use *bootstrapping* [57] to compute CIs, again because it is a non-parametric technique.

2.6.4 Statistical Significance

We use the non-parametric Mann-Whitney U -test [139] to compare and determine the ranking of A and B , where A and B are non-bug-related fuzzing results (e.g., coverage profiles). A p -value > 0.05 implies a statistically-equivalent ranking.

2.7 Chapter Summary

This chapter gives an overview of fuzzing, including both fundamental concepts and the state-of-the-art. We discussed coverage-guided greybox fuzzers in depth because they are the most pervasive type of fuzzer and the focus of this dissertation.

In the next chapter, we delve deeper into the coverage metrics fuzzers use to measure state space exploration. After this, we focus on our contributions to improve fuzzers' state space exploration: (i) how we bootstrap the fuzzing process to lower the cost of the initial state space search (Chapter 4); (ii) a data-flow-based coverage metric, providing an alternative view of a program's state space (Chapter 5); (iii) an investigation of techniques for measuring a fuzzer's state space search (Chapter 6); and (iv) a comprehensive evaluation of the coverage metrics discussed throughout this dissertation (Chapter 7).

Chapter 3

Coverage Metrics

Greybox fuzzers use one or more *coverage metrics* to abstract and approximate a target’s state space. In doing so, the fuzzer maintains a notion of “progress”: expanding coverage implies exploring new program behaviors. In this chapter, we survey and analyze widely-used coverage metrics to understand their advantages and disadvantages.

3.1 Introduction

A fuzzer finds bugs by maximizing its coverage of the target’s state space. Greybox fuzzers are guided by a feedback loop based on one or more *coverage metrics*. Per Algorithm 1 (Chapter 2), the coverage measured (in C_s) tells the fuzzer whether to keep or discard a seed. An effective greybox fuzzer requires coverage metrics that are: (i) **accurate**: program behaviors are accurately approximated; and (ii) **performant**: run-time overheads are minimal. Accurate metrics lead to more-informed decisions in `IsInteresting` (line 16 in Algorithm 1), guiding the fuzzer to explore more of the target’s state space (and increasing the likelihood of finding bugs). In contrast, performant metrics ensure that the fuzzer spends most cycles executing inputs. There is a clear tension between these requirements: more accurate metrics require more intrusive instrumentation to capture subtle changes in program state. But these introduce higher run-time overheads.

Requirement

The *accuracy* of a fuzzer’s coverage metric(s) must be balanced with *performance*, ensuring available cycles are spent finding as many bugs as possible.

Chapter outline. Keeping these accuracy and performance requirements in mind, the remainder of this chapter analyzes the coverage metrics used in 79 coverage-guided greybox fuzzers published from 2013–2022.

3.2 Coverage Metrics in Theory and Practice

Figure 3.1 shows the genealogy of 79 fuzzers published from 2013–2022 and the coverage metrics they use. To realize a coverage metric in practice, fuzzers introduce inaccuracy to reduce run-time overheads (in both time and space). In the remainder of this chapter we describe ideal coverage metrics and how they are realized in practice.

3.2.1 Basic Block

Basic-block (“block”) coverage is the simplest form of control-flow coverage,¹ relying exclusively on block labels. There is no universal approach for labeling a block: UnTracer [151] uses static binary analysis to identify each block in the target and assign it an integer label, while VUzzer [176] uses a block’s run-time address. The simplicity of block coverage means that it can be universally applied with low run-time overhead. However, this simplicity comes at a cost: fuzzers cannot distinguish block sequences and simply count how many times they execute each block. This limits the fuzzer’s ability to differentiate between a loop’s forward and backward edges.

3.2.2 Edge

Edge coverage allows the fuzzer to differentiate between different orderings of the same two blocks. An idealized edge can be described by pairing the previously-executed block label l_{prev} with the current block label l , capturing both intra- and inter-procedural edges. Recording edge coverage this way requires double the memory needed to record the same trace using block coverage. This, combined with the fact that a target typically has many more edges than blocks, means that *approximate* edge coverage is used in practice. We categorize techniques for approximating edge coverage as one of: (i) coarse-grained probabilistic; (ii) coarse-grained link-time optimization (LTO); and (iii) fine-grained control-flow graph (CFG) transformation.

Coarse-grained probabilistic. Randomly-assigned block labels are hashed together to identify an edge between blocks. American Fuzzy Lop (AFL) [231]—perhaps the

¹Function level coverage is even simpler but is too coarse-grained for fuzzers and is never used in practice.

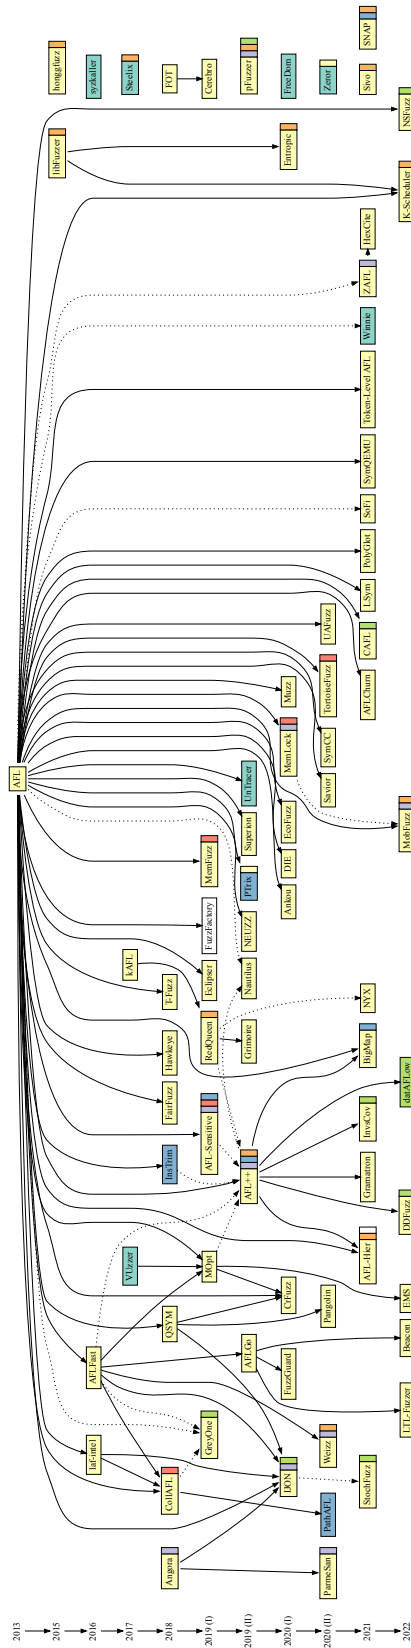


FIGURE 3.1: Greybox fuzzer genealogy for 79 fuzzers published over the last ten years. Coverage metrics are colored: **basic block**, **edge**, **calling context**, **memory access**, **path**, **comparison operator**, and **program variable**. The size of a rectangle within a node differentiates primary and secondary metrics. Solid edges indicate the fuzzer is built on top of the parent fuzzer, while dotted edges indicate the fuzzer reimplements a feature from the parent fuzzer. FuzzFactory is left uncolored because it allows for user-specified coverage feedback [164], while AFL-HIER [213] uses function coverage, which we do not include in our metrics.

most widely-used greybox fuzzer—uses the following hash algorithm:

$$\begin{aligned} p &\leftarrow ((l \gg x) \oplus (l_{\text{prev}} + z)) \\ l_{\text{prev}} &\leftarrow l \gg y \end{aligned} \tag{3.1}$$

Where l is the randomly-generated block identifier (assigned at compile time), l_{prev} is the identifier of the previously-executed block, $x = z = 0$, and $y = 1$. The result p is used as an index into \mathbb{C} . Right-shifting l allows AFL to differentiate between different orders of two blocks. While performant, this approach is prone to hash collisions, resulting in miscounted edges and thus reduced precision [2, 67, 127]. Notably, this problem is not unique to AFL: all fuzzers built on top of AFL [22, 24, 121, 170, 184, 216] or inspired by AFL [8, 37, 163] share the same implementation decision, and thus share the same loss in precision.

Coarse-grained LTO. CollAFL [67] (and fuzzers based on it; e.g., GREYONE [68], PathAFL [227]) uses Eq. (3.1) by selecting x , y , and z at link-time so that edges have unique hashes. Importantly, this approach only ensures the elimination of collisions for known edges; run-time collisions may still exist.

Fine-grained CFG transformation. Hash collisions can be eliminated by breaking *critical edges* in the CFG. Breaking critical edges (i.e., edges whose start/end blocks have multiple outgoing/incoming edges, respectively) allows block coverage to *precisely* imply edge coverage (assuming the number of blocks is fewer than the size of \mathbb{C}) at the cost of increased compile-time complexity. AFL++ [63], honggfuzz [207], and libFuzzer [188] use LLVM’s SanitizerCoverage (SANCOV) instrumentation [208] to achieve this.

3.2.3 Path

While edge coverage improves on basic-block coverage by differentiating between forward and backward edges in a loop, it is unable to reason about the complete ordering of *all* blocks in an execution trace. Path coverage resolves this issue by tracking a complete execution trace—consisting of the sequence of executed blocks—through the target. However, using this definition of path coverage is infeasible: the number of paths (which can be infinite) and the memory overhead associated with maintaining this coverage information is prohibitive [67].

AFL-Sensitive [214] introduced *n-gram coverage* (later adopted by AFL++), where a history of n edges is used to approximate a path. These n edges are xor-ed together to generate an index into C . In contrast, INSTRIM applies “CFG-aware” instrumentation to reduce the number of instrumented blocks (improving run-time performance) while ensuring different execution paths remain distinguishable [97]. Finally, PTRIX [40] uses Intel Processor Tracing (PT)—a low-overhead execution tracing feature available in modern Intel CPUs—to record program paths while fuzzing. PTRIX proposes an *elastic* approach to measure code coverage: encoded PT packets are concatenated and then hashed to achieve low-overhead path coverage. Like the coarse-grained probabilistic edge coverage metric, these approaches are lossy and collisions may occur.

3.2.4 Calling Context

Function call information allows a fuzzer to track context-sensitive coverage information (per Section 1.2, this was how Ormandy [162] discovered CVE-2021-43527). However, efficiently computing accurate calling context is difficult [25, 206]. For example, each stack frame requires tracking: (i) the current function; (ii) the current block; (iii) the set of local variables and their values; and (iv) allocated memory. This is (a) computationally expensive to convert into a single index into C , and (b) memory-intensive for efficient access (i.e., without walking the stack) [25]. Unsurprisingly, this leads fuzzers to approximate calling context in practice.

AFL-Sensitive [214], Angora [37], ParmeSan [163], and AFL++ [63] approximate calling context by maintaining a rolling xor hash of call site identifiers (similar to the n -gram approach described in Section 3.2.3, but limited to function call/return sites). AFL++ also supports *k-call-site-sensitivity*, where the last k call sites are recorded in the hash. This hash is combined with edge coverage to provide *context-sensitive edge coverage* (again, with the possibility of hash collisions). While xor has minimal overhead, it reduces precision because it does not differentiate between recursive calls of arbitrary depth. Other approaches (e.g., those proposed by Bond and McKinley [25] and Sumner et al. [206]) are needed if differentiating between recursive function calls of arbitrary depth is required.

3.2.5 Memory Access

Memory-safety vulnerabilities remain one of the most common bug classes [147]. Consequently, some fuzzers incorporate *memory accesses* into their coverage metrics; inputs that lead to new memory access patterns (that may subsequently lead to memory-safety errors) are prioritized. This requires encoding memory reads/writes

in C . When this encoding is code-dependent (e.g., encodes a memory access location or static address), then collected coverage approximates *control flow*. However, when it is data-dependent (e.g., encodes an input-controlled value or address), collected coverage approximates *data flow*. Therefore, C may record information about the location of the memory access (the code location) or the target of the memory address (the data location). Similarly, the value read from/written to a particular memory address may (or may not) be included in the coverage metric.

AFL-Sensitive proposes *memory-access-aware* edge coverage, where both edge coverage and memory reads/writes are combined in the same C (although reads and writes are distinguished by allocating their hash to different regions of C). MEMFUZZ [49] adopts a similar approach, except (a) a lightweight static analysis limits instrumentation to “*interesting memory accesses*”, and (b) edge coverage and memory accesses are disjoint in C (allowing for more fine-grained checks for new program behaviors). Both AFL-Sensitive and MEMFUZZ are prone to hash collisions in C , reducing precision. Finally, CollAFL proposes a seed-selection policy that prioritizes inputs with memory access operations. Unlike the previous approaches, CollAFL counts memory accesses *post hoc*: after executing a testcase s , edges covered by s are retrieved from C_s and the number of memory accesses made per covered block is determined. This count is combined with C_s to weigh s .

3.2.6 Comparison Operator

Instrumented comparison operators offer a form of lightweight data flow. Both honggfuzz and libFuzzer capture comparisons containing constant operands, feeding these constants into `MutateInput` (Algorithm 1). Steelix [123] instruments comparisons to measure “comparison progress”. Comparison progress—which measures how many consecutive bytes are matched—is stored with block coverage in C . Similar to MEMFUZZ’s memory access instrumentation, Steelix limits its comparison instrumentation to “interesting” comparisons; one byte and return value comparisons are ignored.

3.2.7 Program Variable

The previously-discussed coverage metrics are primarily control-flow based. While pervasive, control-flow-based coverage metrics lose important information about program behavior. For example, greybox fuzzers rely on coverage information to decide which input mutations lead to new program behaviors. However, the process for uncovering new behaviors can be highly inefficient, because a fuzzer guided by control-flow coverage alone cannot identify *which* mutated input bytes led to new

program behavior. Differences in data access and manipulation within a single code path are lost. Data-flow-based metrics attempt to solve this problem.

INVSCOV [61] augments coarse-grained probabilistic edge coverage with values and relationships among program variables. However, rather than looking at all program variables (resulting in state explosion), INVSCOV only considers “*likely invariants*”: constraints on the values and relationships of variables learned by dynamically tracing a seed set. Invariant violations are encoded in `C`, signaling to the fuzzer (via `IsInteresting`) that a new program state has been reached.

DDFuzz [141] also augments coarse-grained probabilistic edge coverage with data flows between program variables. Here, data flows are derived from the target’s data dependency graph (DDG). DDGs describe the data flows between instructions in a program and are traditionally used by optimizing compilers [60]. Like INVSCOV, DDFuzz only considers a subset of program variables (again, to prevent state explosion): variable definition sites are restricted to `load` and `alloca` instructions in the LLVM intermediate representation (IR), while variable uses are restricted to `store` and `call` instructions. Importantly, the DDG is not used to expand coverage. Instead, the DDG is used to revisit particular program locations via new paths involving different variable dependencies (thus exposing program states not visible in the CFG). Consequently, further filtering is applied to discard data flows subsumed by edge coverage. Both INVSCOV and DDFuzz use heuristics to decide which program variables to track. This reduces run-time overhead, but may mean important variables are missed.

Directed greybox fuzzers (e.g., AFLGo [24]) aim to reach a target location in a codebase. They typically use a control-flow-based distance metric to determine how close a given test case is to this location. CAFL [121] extends this idea by incorporating “*data conditions*” into the distance metric. Satisfying data conditions requires (a) capturing relevant variables, and (b) computing distances² between these variables’ values. `IsInteresting` uses these distances to determine if the fuzzer is making progress toward the target location.

3.3 Chapter Summary

In this chapter, we analyzed the coverage metrics used by 79 fuzzers released from 2013–2022. In doing so, we identified and surveyed three primary metrics (basic block, edge, and path coverage) and four secondary metrics (calling context, memory access, comparison operator, and program variable). Among these seven metrics,

²This calculation depends on the comparison operator used in the condition.

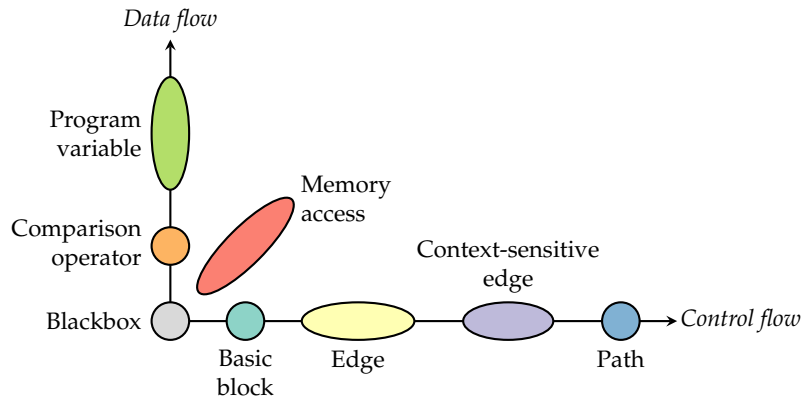


FIGURE 3.2: Qualitative comparison of greybox fuzzers' coverage metrics. Precision increases from left to right (along the x -axis) and bottom to top (along the y -axis). Most fuzzers rely exclusively on control-flow coverage (along the x -axis). However, these metrics can be augmented with data-flow-based metrics (along the y -axis). Some metrics (e.g., edge) vary in implementation, leading to varying levels of precision.

there exist a wide variety of design and implementation choices (edge coverage alone has been implemented in several different ways according to precision/performance tradeoffs).

Inspired by the hierarchy of control- and data-flow coverage metrics proposed by Horgan, London, and Lyu [94], Fig. 3.2 qualitatively summarizes and compares the coverage metrics we discussed in this chapter. Blackbox fuzzers appear in the bottom-left corner: they have no feedback mechanism and hence are the least precise. Most greybox fuzzers lie on the x -axis, as the majority use some form of edge coverage. Implementation decisions affect the precision of edge, memory access, and program variable metrics. Similarly, memory access coverage can be both control- and data-flow-based depending on implementation tradeoffs. A third axis—coverage of *input structure*—could also be included on this graph. This is particularly relevant to grammar-based fuzzers, where (a) the structure of the input format is known, and (b) maximizing the generation of syntactic input features will lead to increased coverage of the target's state space [83]. However, we omit this axis for readability and because our focus is on a target's state space, not its input space.

This chapter sets the scene for the remainder of this dissertation: how do we use these (sometimes incompatible) opinions about coverage to measure and improve a fuzzer's state space exploration. In the next chapter we return to the start of a fuzzing campaign (① in Fig. 1.1, Chapter 1); how to bootstrap the fuzzer to simplify its state space search.

Chapter 4

Seed Selection

Mutational fuzzers require an initial set of non-crashing inputs (a “corpus”) to bootstrap their state space search. These inputs—or *seeds*—should exercise different parts of the target’s state space. By mutating these seeds, the fuzzer aims to reach new, potentially bug-inducing states. In this chapter, we present: (i) an optimal approach for constructing corpora from a large collection of potential inputs; and (ii) a comparative evaluation of how corpus selection affects fuzzing outcomes.

4.1 Introduction

“Garbage in, garbage out” is the idea that the quality of input data determines the quality of outputs. It applies to many fields of computing, including fuzzing. We argue that fuzzers more efficiently explore a target’s state space when bootstrapped with “high quality” inputs. What determines the quality of these inputs? We answer this question here.

Chapter outline. We begin by systematically examining seed selection practices used in prior work (Section 4.2). Then, we focus on the problem of *corpus minimization* (Section 4.3): given a (potentially large) corpus of initial seeds, how to best select a subset of these seeds (e.g., to remove redundancy) to bootstrap a fuzzer? Here we also introduce OPTIMIN, our corpus minimization technique. OPTIMIN is unique in that—unlike other corpus minimization techniques—it produces an *optimal minimum* corpus. Finally, we perform a comprehensive empirical evaluation (33 CPU-yr) to understand how seed choice—including various corpus minimization techniques—affects the fuzzing process (Section 4.4).

4.2 Seed Selection Practices

The process for selecting the initial corpus of seeds varies wildly across fuzzer evaluations and deployments. We systematically review this variation in the following sections, first considering experimental evaluations, followed by industrial-scale deployments.

4.2.1 An Ideal Initial Corpus

Ideally, an initial fuzzing corpus should contain inputs that exercise a diverse set of states in the target’s state space. A diverse set of states means that a fuzzer’s mutations is more likely to “nudge” an input into a new, potentially buggy, state. In contrast, having only a small number of states covered by the initial corpus means the fuzzer must waste time attempting to reach valid states, rather than focusing on exploring “corner cases” in the target’s state space.

Seeds in a fuzzer’s initial corpus should also be fast to execute, so that a fuzzer’s iteration rate is not negatively impacted. Thus, large seeds are discouraged.

Finally, there should be minimal redundancy in the corpus. That is, seeds that cover the same states should be removed from the corpus. This is analogous to *test suite minimization*, where redundant tests should be removed from the test suite to reduce the time required to run the test suite [86, 98, 126, 220].

4.2.2 In Experimental Evaluation

Remarkably, Klees et al. [115] found “most papers treated the choice of seed casually, apparently assuming that any seed would work equally well, without providing particulars”. In particular, of the 32 papers they surveyed (authored between 2012 and 2018): ten used non-empty seed(s), but it was not clear whether these seeds were valid inputs; nine assumed the existence of valid seed(s), but did not report how these were obtained; five used a random sampling of seed(s); four used manually constructed seed(s); and two used empty seed(s). Additionally, six papers used a combination of seed selection techniques. Klees et al. [115] conclude “it is clear that a fuzzer’s performance on the same program can be very different depending on what seed is used” and recommend “papers should be specific about how seeds are collected”.

We examine an additional 74 papers published since 2018¹ to see if these recommendations have been adopted. Table 4.1 summarizes our findings.

¹Published in top security and software testing venues: ACSAC, ASIA CCS, CCS, EURO S&P, NDSS, RAID, S&P, SEC, ASE, ESEC/FSE, ICSE, and ISSTA.

TABLE 4.1: Summary of past fuzzing evaluation, focusing on seed selection. We adopt the categories and notation used by Klees et al. [115]: **R** means randomly sampled seeds; **M** means manually constructed seeds; **G** means automatically generated seed; **N** means non-empty seed(s) with unknown validity; **V** means the paper assumes the existence of valid seed(s), but with unknown provenance; **E** means empty seeds; **/** means different seeds were used in different programs, but only one kind of seeds in one program; and an empty cell means that the paper’s evaluation did not mention seed selection. We introduce an additional category: **V*** means valid seed(s) with known provenance. We also indicate whether the evaluation is reproducible with the same seeds.

Year	Paper	Seed	Reproducible	Year	Paper	Seed	Reproducible	
2018	CollAFL [67]		✗		[20]	V*	✓	
	Hawkeye [34]	E/V*	✗		AFLCHURN [239]	V*	✓	
	QSYM [230]	M, V*	✓		AFL-HIER [213]	M/V*	✓	
2019	AFL-Sensitive [214]	E/V*	✓		CAFL [121]	V*	✗	
	CEREBRO [122]		✗		Controlled Compilation [195]		✗	
	Eclipser [46]	M/V*	✓		Gramatron [199]	E	✓	
	FUZZFACTORY [164]	M, V*	✓		HEXCITE [153]		✗	
	GRIMOIRE [16]	M	✓		INVSov [61]	V	✗	
	MEMFUZZ [49]	E, V	✗		LSym [145]	V	✗	
	MOPT [134]	R	✓		POLYGLOT [43]	V*	✗	
	NAUTILUS [8]	G, M	✓		SIVO [157]	R	✗	
	NEUZZ [193]	G, V*	✓		SNAP [51]	V	✗	
	PFUZZER [142]	E	✓		SoFi [85]	R, V, V*	✗	
	PTRIX [40]	V*	✓		STOCHFuzz [236]	V/V*	✓	
	REDQUEEN [9]	M	✓		SYMQEMU [174]	V*	✓	
	Superion [216]	R, V	✗		Token-Level AFL [184]	V*	✗	
	UnTracer [151]	V	✓		UNIFUZZ [124]	R	✓	
Zest [165]	V	✓		ZAFUZZ [152]	E, V*	✓		
2020	[19]	E	✓		[23]	V*	✓	
	Ankou [137]	V*	✓		[222]	G	✗	
	CRFUZZ [198]		✗		[223]	V*	✗	
	CUPID [78]	E/V*	✓		BEACON [99]	V*	✗	
	DIE [168]	V*	✓		DDFUZZ [141]	R	✗	
	EcoFuzz [229]	V*	✓		EMS [133]	R/V	✗	
	ENTROPIC [21]	E/V*	✓		JIGSAW [36]	V*	✓	
	FuZZan [104]	E/V*, R	✓		K-Scheduler [190]	V*	✓	
	FuzzGen [103]		✗		LTL-Fuzzer [143]	R/V*	✓	
	FuzzGuard [240]	E/V	✗		MobFuzz [232]	V*	✗	
	GREYONE [68]	R	✗		PATA [125]	E/R/V*	✗	
	IJON [7]	M/R/V	✗		TRUZZ [233]	V*	✓	
	Magma [84]	V*	✓					
	MEMLOCK [219]	V	✓					
	MEUZZ [39]	V*	✓					
	MTFUZZ [191]	G, V*	✗					
	MUZZ [35]	V	✗					
	PANGOLIN [100]	V*	✓					
	ParmeSan [163]	E	✗					
	PathAFL [227]	E/R/V	✗					
	SAVIOR [41]	V*	✓					
	SYMCC [173]		✗					
	TortoiseFuzz [217]	R	✗					
UAFUZZ [156]	E/V*	✗						
WEIZZ [62]	M/V/V*	✗						
ZEROR [237]	E/V*	✗						

Unreported seeds. Seven papers make no mention of their seed selection procedure. One, FuzzGen, explicitly mentions “*standardized common seeds*” [103] are key to a valid comparison, yet does not mention the seeds used.

Benchmark and fuzzer-provided seeds. Eight papers (Hawkeye, FUZZFACTORY, CUPID, ENTROPIC, FuZZan, ZEROR, STOCHFUZZ, and PATA) evaluate fuzzers on the Fuzzer Test Suite (FTS) [75], which provides seeds for 14 of its 24 targets (commit 6955fc9). Of these seed sets, eight contain only one or two seeds. When no seed is provided, three papers (CUPID, FuZZan, and PATA) default to the empty seed. It is unclear what seeds the remaining five papers (Hawkeye, FUZZFACTORY, ENTROPIC, ZEROR, and STOCHFUZZ) default to. When evaluating on other fuzzer benchmarks—including FUZZBENCH [144], Magma [84], and LAVA-M [55]—the provided seeds are used. Finally, seven papers (AFL-Sensitive, PTRIX, EcoFuzz, PANGOLIN, SAVIOR, Wu et al. [223], and MobFuzz) used the singleton seed sets provided by American Fuzzy Lop (AFL).

Manually-constructed seeds. Two papers (REDQUEEN and GRIMOIRE) used “*an uninformed, generic seed consisting of different characters from the printable ASCII set*” [9]. However, the authors do not justify (a) why this specific singleton corpus was chosen, and (b) what impact this choice has on the authors’ real-world results, particularly when fuzzing `binutils`, where most of the targets accept non-ASCII, binary file formats (e.g., `readelf`, `objdump`). Similarly, IJON used “*a single uninformative seed containing only the character 'a'*” [7], QSYM used “*a dummy ASCII file containing 256 'A's*” [230], and Eclipser used “*a dummy seed composed of 16 consecutive NULL bytes*” [46]. Finally, AFL-HIER started fuzzing the CGC binaries with a single seed containing the string “`123\n456\n789\n`” [213] (rather than the provided seeds).

Random seeds. Thirteen papers (MOPT, Superior, FuZZan, GREYONE, IJON, PathAFL, TortoiseFuzz, SoFi, UNIFUZZ, DDFuzz, EMS, LTL-Fuzzer, and PATA) randomly select seeds from either (a) a larger corpus of seeds provided by developers of a particular target, or (b) by crawling the Internet. For example, DDFuzz selects a small subset of seeds “*from the test directories... in the repositories of the projects*” [141]: these subsets range from three seeds (`bison`) to 65 seeds (`mir`, which contains over 900 test inputs). Which seeds are selected and why is not explained.² Of these papers, three (Superior, GREYONE, and SoFi) specifically mention using `af1-cmin`, AFL’s corpus minimization tool (discussed further in Section 4.3), to remove duplicate seeds from the random seed set.

²The authors were happy to share their starting seeds when contacted.

Empty seeds. Sixteen papers use an empty seed to bootstrap the fuzzing process. In particular, both Böhme, Manès, and Cha [21] and Böhme and Falk [19] explicitly removed the corpora provided by OSS-Fuzz (discussed further in Section 4.2.3) and used the empty seed because they found “*initial seed corpora... are often for saturation: feature discovery has effectively stopped shortly after the beginning of the campaign*”.

A reproduction experiment: REDQUEEN. To demonstrate the importance of seed selection, we reproduce an experiment from the REDQUEEN evaluation. Aschermann et al. [9] fuzz several `binutils`’ programs, bootstrapping each trial with an “uninformed, generic seed” (discussed previously). Their `readelf` results are particularly striking: notably, AFLFast and honggfuzz cover little code. We repeat this experiment but use a variety of initial seeds, including: (i) the original, uninformed seed; (ii) a single, valid ELF file (from AFL’s seed set); and (iii) a collection of ELF files sourced from the ALLSTAR [200] and Malpedia [172] datasets (reduced from 104,737 to 366 seeds using `af1-cmin`). In place of the original REDQUEEN (which we were unable to build and reproduce) we use AFL++ [63] with “`CmpLog`” instrumentation enabled; this reimplements REDQUEEN’s “input-to-state correspondence”.

Our results appear in Fig. 4.1 and clearly show seed choice’s impact on code coverage. Similar to the results of Aschermann et al. [9], AFLFast bootstrapped with the uninformed seed explores very little of `readelf`’s code: less than 1%. However, this increases to ~38% for AFLFast bootstrapped with the valid ELF file, making it much more competitive against both honggfuzz and AFL++ (although AFL++ still outperforms them both by ~15%). Finally, while the `af1-cmin` corpus has a negligible impact on AFLFast and honggfuzz, it results in a significant improvement when fuzzing with AFL++, increasing coverage to ~60%.

Requirement

At a minimum, fuzzer evaluations must report the seed set used to bootstrap the fuzzing process. To ensure reproducibility, artifacts must provide the initial seed set (since results vary wildly depending on the seeds used). Ideally, fuzzer evaluations should experiment with different initial seed corpora to see how varying initial seeds affect fuzzing outcomes.

4.2.3 In Deployment

In addition to being an active research topic, fuzzers are frequently deployed to find bugs in real-world software [3, 33, 150, 159]. Notably, security professionals




FIGURE 4.1: Code coverage of `readelf` with different initial seed sets. The mean coverage (using Clang’s source-based region coverage metric) and 95% bootstrap confidence interval over five repeated 10h trials is shown. The x -axis uses a log scale.

also recognize the importance of seed selection. For example, the developers of the Mozilla Firefox browser remark [149]:

Mutation-based strategies are typically superior to others if the original samples [i.e., seeds] are of good quality because the originals carry a lot of semantics that the fuzzer does not have to know about or implement. However, success here really stands and falls with the quality of the samples. If the originals don't cover certain parts of the implementation, then the fuzzer will also have to do more work to get there.

In contrast to the evaluation practices described in Section 4.2.2, “industrial fuzzing” eschews small seed sets and the empty seed in favor of large corpora. For example, seed corpora in Google’s continuous fuzzing service for open-source software, OSS-Fuzz (commit 0deef6), ranges in size from a single seed (e.g., OpenThread, ICU) to 62,726 seeds (Suricata). Of the 363 OSSFuzz projects, 135 projects supply an initial corpus (37%) for 706 fuzzable targets. The mean corpus size is 1,083 seeds and the median is 36 seeds. More than half of the 135 projects include more than 100 seeds.

We examined the Suricata corpus in more detail because it provided the largest number of seeds (62,726). We found large redundancy in these 62,726 seeds: we reduced the corpus to 31,234 seeds (~50% decrease) by discarding seeds with an identical MD5 hash. We were able to reduce the size of the corpus further (down to 145 seeds, a 99% reduction) by again applying `af1-cmin` (similarly to the `readelf` reproduction experiment described in Section 4.2.2). This redundancy is wasteful, as it leads to seeds clogging the fuzzing queue, which hinders and delays the mutation of more promising seeds. We discuss corpus minimization in the following section.

Requirement

When deploying fuzzers at an industrial scale, it is imperative that seeds exhibiting redundant behavior be removed from the fuzzing queue, as they will lead to wasted cycles.

4.3 Corpus Minimization

Orthogonal to the seed selection practices examined in Section 4.2, many popular fuzzers (e.g., AFL [231], libFuzzer [188], honggfuzz [207]) provide *corpus minimization* (sometimes called *distillation*) tools. Corpus minimization assumes a large corpus of seeds already exists, and thus a corpus minimization tool *reduces* this large corpus to

a subset of seeds that are then used to bootstrap the fuzzing process. How are these seeds selected?

Abdelnur et al. [1] first formalized this problem as an instance of the *minimum set cover problem* (MSCP). The MSCP states that given a set of elements U (the universe) and a collection of N sets $S = s_1, s_2, \dots, s_N$ whose union equals U , what is the *smallest* subset of S whose union still equals U . This smallest subset $\mathcal{C} \subseteq S$ is known as the *minimum set cover*. Moreover, each $s_i \in S$ may be associated with a weight w_i . In this case, the *weighted MSCP* (WMSCP) attempts to minimize the total cost of elements in \mathcal{C} .

(W)MSCP is NP-complete [111], so Abdelnur et al. [1] used a *greedy* algorithm to solve the unweighted MSCP. Here, U consists of code coverage information for the set of seeds in the original collection corpus. Subsequently, code coverage (e.g., basic blocks, edges) continues to be used to characterize seeds in a fuzzing corpus due to the strong positive correlation between code coverage and bug finding [23, 76, 117]. Computing \mathcal{C} is therefore equivalent to finding the minimum set of seeds that still maintains the code coverage observed in the collection corpus.

Several corpus minimization techniques have been proposed since the work of Abdelnur et al. [1]. We discuss prior work in Section 4.3.1, and introduce our own minimization technique, OPTIMIN, in Section 4.3.2.

4.3.1 Prior Work

Peach Minset. The Peach fuzzer [71] provides a corpus minimization tool, Peach Minset. Despite its name, Rebert et al. [177] found that Peach Minset does not in fact calculate \mathcal{C} , “*nor a proven competitive approximation thereof*”. Peach Minset uses basic block coverage.

MINSET. Rebert et al. [177] extended the work of Abdelnur et al. [1] by also computing \mathcal{C} weighted by execution time or file size. They designed six corpus minimization techniques, simulating and empirically evaluating these techniques over several fuzzing campaigns using the Basic Fuzzing Framework (BFF). These techniques use basic block coverage. Rebert et al. [177] found that UNWEIGHTED MINSET—an unweighted greedy-reduced minimization—performed best in terms of minimization ability, and that the PEACH SET algorithm (based on the previously-discussed Peach Minset) found the highest number of bugs. We extend the work of Rebert et al. [177] in Section 4.4 with an evaluation based on modern coverage-guided greybox fuzzing.

LibFuzzer. LLVM’s libFuzzer [188] provides a corpus “merge” feature for minimizing a large corpus “*while still preserving the full coverage [of the collection corpus]*” [130]. Unlike MINSET and Peach Minset, which use basic block coverage, libFuzzer uses *feature* coverage. Here, a feature combines the edge executed and a hit count [19]. This means two seeds s_i and s_j exercising the same edges will have $C_{s_i} \neq C_{s_j}$ if one seed executes an edge more often [21]. LibFuzzer’s merge algorithm is similar to MINSET: seeds are sorted by size (with smaller seeds preferred) and the number of covered features (with higher hit counts preferred). Once sorted, seeds are greedily added to \mathcal{C} whenever a new feature is found to be covered.

MoonShine. MoonShine [166] is a corpus minimization tool for operating system (OS) fuzzers. OS fuzzers typically test the system-call interface between the OS kernel and user-space applications. As such, the seeds minimized by MoonShine are a list of system calls gathered from program traces.

SmartSeed. SmartSeed [135] takes a different approach to those previously described. Rather than minimizing a corpus of seeds, SmartSeed instead uses a machine learning model to generate “valuable” seeds, where a seed is considered valuable if it uncovers new code or produces a crash.

af1-cmin. Due to AFL’s popularity, af1-cmin [231] is perhaps the most widely-used corpus minimization tool. It implements a greedy minimization algorithm but has a unique approach to coverage. In particular, af1-cmin reuses AFL’s implementation of edge coverage to categorize seeds at minimization time, recording an approximation of edge *frequency count*, not just whether the edge has been taken. Moreover, af1-cmin *bins* edge counts such that changes in edge frequency counts within a single bin are ignored, while transitions from one bin to another are “*flagged as an interesting change in program control flow*” [231]. When minimizing, af1-cmin chooses the smallest seed in the collection corpus that covers a given edge count and then performs a greedy, weighted minimization. We consider af1-cmin and MINSET as representatives of the state-of-the-art in corpus minimization tools and include both in our evaluation.

4.3.2 OPTIMIN

The corpus minimization techniques described in Section 4.3.1 all employ heuristic algorithms to approximate \mathcal{C} . This is because the underlying problem, the (W)MSCP, is NP-complete. However, in the case of corpus minimization, we found exact solutions were nonetheless computable in reasonable time. We achieved this by encoding the problem as a maximum satisfiability (MaxSAT) problem and using




FIGURE 4.2: An example program and set of traces for illustrating corpus minimization. The program in Fig. 4.2a is executed with seeds S_1 , S_2 , and S_3 , producing the traces in Fig. 4.2b .

an off-the-shelf MaxSAT solver. Thus, we implement OPTIMIN, an optimal corpus minimization tool for AFL.

OPTIMIN³ uses the EvalMaxSat solver [10] to pose and solve corpus minimization as a MaxSAT problem. Unlike the boolean satisfiability (SAT) problem—which determines whether the variables in a given Boolean formula can be assigned values to make the formula evaluate to true—the MaxSAT problem divides constraints into hard and soft constraints with the aim to satisfy all hard constraints and maximize the total number (or weighted total) of satisfied soft constraints. OPTIMIN considers edge coverage as a hard constraint, while seed exclusion is considered a soft constraint. This approach ensures the solution \mathcal{C} covers all edges with the minimal number of seeds, and is optimal in the sense that \mathcal{C} is guaranteed to be exact (rather than an approximation).

We describe our approach using the example in Fig. 4.2. The program in Fig. 4.2a is executed with three seeds, producing the traces in Fig. 4.2b . To obtain \mathcal{C} , OPTIMIN:

1. Encodes each edge as a disjunction of seeds covering that edge. For example, the edge (A, B) is encoded as $S_1 \vee S_2 \vee S_3$ because it is covered by all seeds. Similarly, the edge (B, E) is encoded as $S_1 \vee S_2$ because it is only covered by seeds S_1 and S_2 . Intuitively, this encoding tells the solver to select S_1 or S_2 to ensure the edge (B, E) is represented in \mathcal{C} .
2. Produces a conjunction of seed disjunctions. This conjunction is the *hard* constraint (thus assigned the maximum weight \top), guaranteeing all edges are covered in \mathcal{C} (i.e., no coverage is “lost” during minimization).

³The “optimal minimizer”. OPTIMIN is available at <https://github.com/HexHive/fuzzing-seed-selection>.

3. Encodes each seed S_i as a *soft* constraint with a given weight W_i . In a MaxSAT solver, minimization is performed via negation. Intuitively, this tells the solver to exclude a seed from \mathcal{C} . If the solver must choose between a set of seeds (to satisfy the hard constraint), then the one with a lower weight is preferred. For unweighted minimizations (i.e., when only the number of seeds matters, while their features—such as size—do not), $W_i = 1$. In the weighted case, W_i is S_i 's file size (leading OPTIMIN to preference smaller seeds).
4. Encodes the hard and soft constraints in weighted conjunctive normal form (CNF) [12]. These constraints are passed to EvalMaxSat, which produces \mathcal{C} .

We present the complete set of constraints for this example in Table 4.2.

OPTIMIN is not the first tool to generate optimal solutions to minimization problems in software testing. For example, MINTS [98] and Nemo [126] both use integer linear programming (ILP) solvers to perform *test-suite minimization*; i.e., eliminate redundant test cases from a test suite “based on any number of criteria [e.g., statement coverage, time-to-run, setup effort]” [98]. When developing OPTIMIN, we explored the use of mixed integer programming solvers (which are more general than ILP solvers), but found them to be *orders-of-magnitude slower* than EvalMaxSat.

4.4 Evaluation

We perform a large-scale evaluation (33 CPU-yr of fuzzing) to understand the impact of seed selection on fuzzing’s ultimate goal: finding bugs in real-world software. We aim to answer the following research questions:

- RQ 1** How effective are corpus *minimization* tools at producing a minimal corpus? (Section 4.4.2)
- RQ 2** What effect does seed selection have on a fuzzer’s *bug finding* ability? Do fuzzers perform better when bootstrapped with (a) a small seed set (e.g., empty or singleton set), or (b) a large corpus of seeds, derived from an even larger collection corpus after applying a corpus minimization tool? (Section 4.4.3)
- RQ 3** How does seed selection affect *state space exploration*? Does starting from a corpus that executes more instrumentation data points result in greater code coverage, or does a fuzzer’s mutation engine naturally achieve the same coverage (e.g., when starting from an empty seed)? (Section 4.4.4)

We find that while corpus minimization greatly impacts fuzzing campaigns, the underlying minimization tool is less important, as long as *some* form of minimization

TABLE 4.2: OPTIMIN constraints for the example in Fig. 4.2. Edges are encoded as a disjunction of seeds covering that edge. The hard constraint—a conjunction of seed disjunctions—ensures all edges are covered in \mathcal{C} . Soft constraints—weighted by W_i —minimize the number of seeds selected.

(A) Weighted minimization (WOPT in Section 4.4.1).

Coverage.

Edges	Constraint
(A, B)	$S_1 \vee S_2 \vee S_3$
(B, E)	$S_1 \vee S_2$
(B, D)	S_3
$(E, F), (F, H), (F, I), (H, J), (I, D),$ and (J, E)	S_2
$(E, G),$ and (G, D)	S_1

Hard constraints.

Constraint	Weight
$(S_1 \vee S_2 \vee S_3) \wedge (S_1 \vee S_2) \wedge S_3 \wedge S_2 \wedge S_1$	\top

Soft constraints.

Constraint	Weight
$\neg S_1$	W_1
$\neg S_2$	W_2
$\neg S_3$	W_3

(B) Weighted minimization accounting for edge frequencies (WMOPT in Section 4.4.1). The loop backedge (J, E) is split into two edges (the number of times the back edge was executed).

Coverage.

Edges	Constraint
(A, B)	$S_1 \vee S_2 \vee S_3$
(B, E)	$S_1 \vee S_2$
(B, D)	S_3
$(E, F), (F, H), (F, I), (H, J), (I, D), (J, E)_1,$ and $(J, E)_2$	S_2
$(E, G),$ and (G, D)	S_1

Hard constraints.

Constraint	Weight
$(S_1 \vee S_2 \vee S_3) \wedge (S_1 \vee S_2) \wedge S_3 \wedge S_2 \wedge S_1$	\top

Soft constraints.

Constraint	Weight
$\neg S_1$	W_1
$\neg S_2$	W_2
$\neg S_3$	W_3

occurs. All experimental data is available at <https://datacommons.anu.edu.au/DataCommons/rest/display/anudc:6106>.

4.4.1 Methodology

Fuzzer Selection

Our real-world fuzzing campaigns use AFL (v2.52b) in greybox mode, while our Magma and FTS campaigns also include AFL++ with (a) CmpLog instrumentation enabled, and (b) a 250 KiB coverage map. We configure both fuzzers for single-system parallel execution with one main and one secondary node; the primary node focuses on deterministic checks, while the secondary node proceeds skips these checks and proceeds directly to random mutation.

For the FTS and the real-world targets we compile using AFL’s LLVM [120] (v8) instrumentation for 32-bit x86 and AddressSanitizer (ASAN) [187]. We chose LLVM instrumentation over AFL’s assembler-based instrumentation because LLVM’s offers the best level of interoperability with ASAN . We compile Magma targets using their default build configuration (i.e., 64-bit x64 without ASAN).

We tune AFL’s timeout and memory parameters for each target to enable effective fuzzing.⁴ When fuzzing the FTS we configure the target process to respawn after *every* iteration (due to stability issues encountered when fuzzing in parallel execution mode). All other parameters are left at their default values.

Target Selection

We use targets from the Magma [84] and FTS [75] benchmarks, and six popular open-source programs (spanning 14 different file formats) to test different seed selection strategies. Table 4.3 summarizes these targets.

We exclude 14 of the 24 FTS targets, because: (i) they contain only memory leaks (e.g., *proj4-2017-08-14*), which are not detected by AFL by default; or (ii) we were unable to find a suitably-large collection corpus for a particular file type (e.g., ICC files for *lcms-2017-03-21*). Similarly, two Magma targets were excluded (*openssl* and *sqlite3*) because we were unable to find a suitably large corpus, leaving us with five targets. The six real-world targets are popular programs that are both (a) commonly fuzzed, and (b) operate on a diverse range of file formats (e.g., images, audio, documents).

⁴Per-target settings are available at <https://datacommons.anu.edu.au/DataCommons/rest/display/anudc:6106>.

TABLE 4.3: Fuzzing targets. The FTS provides drivers.

	Target	Driver	Command-line	Version	File type
Magma	<i>libpng</i>	<code>libpng_read_fuzzer</code>			PNG
	<i>libtiff</i>	<code>tiff_read_rgba_fuzzer</code>			TIFF
	<i>libxml2</i>	<code>libxml2_xml_reader_for_file_fuzzer</code>			XML
	<i>php-exif</i>	<code>exif</code>			JPEG
	<i>php-json</i>	<code>json</code>			JSON
	<i>php-parser</i>	<code>parser</code>			PHP
	<i>poppler</i>	<code>pdf_fuzzer</code>			PDF
FTS	<i>freetype2</i>			2017	TTF
	<i>guetzli</i>			2017-3-30	JPEG
	<i>json</i>			2017-02-12	JSON
	<i>libarchive</i>			2017-01-04	GZIP
	<i>libjpeg-turbo</i>			07-2017	JPEG
	<i>libpng</i>			1.2.56	PNG
	<i>libxml2</i>			2.9.2	XML
	<i>pcr2</i>			10.00	Regex
	<i>re2</i>			2014-12-09	Regex
	<i>vorbis</i>			2017-12-11	OGG
Real-world	<i>freetype2</i>	<code>char2svg</code>	<code>@@ @</code>	2.5.3	TTF
	<i>librsvg</i>	<code>rsvg-convert</code>	<code>-o /dev/null @@</code>	2.40.20	SVG
	<i>libtiff</i>	<code>tiff2pdf</code>	<code>-o /dev/null @@</code>	4.0.9	TIFF
	<i>libxml2</i>	<code>xmllint</code>	<code>-o /dev/null @@</code>	2.9.0	XML
	<i>poppler</i>	<code>pdftotext</code>	<code>@@ /dev/null</code>	0.64.0	PDF
	<i>sox-mp3</i>	<code>sox</code>	<code>-single-threaded @@ -b 16 -t aiff /dev/null channels 1 rate 16k fade 3 norm</code>	14.4.2	MP3
	<i>sox-wav</i>	<code>sox</code>	<code>-single-threaded @@ -b 16 -t aiff /dev/null channels 1 rate 16k fade 3 norm</code>	14.4.2	WAV

Sample Collection

For each file type in Section 4.4.1, we built a Web crawler using Scrapy [241] to crawl the Internet for 72 h to create the collection corpus. For image files, crawling started with Google search results and the Wikimedia Commons repository. For media and document files, crawling started from the Internet Archive and Creative Commons collections. We used the regular expressions from RegExLib [178], and sourced OGG files from old video games [112, 148, 161] (in addition to the Internet Archive). We sourced PHP files from test suites for popular PHP interpreters (e.g., Facebook’s HipHop Virtual Machine [58]) and from popular GitHub repositories (e.g., WordPress [221]). Finally, we found TIFF files to be relatively rare, so we generated 40 % of the TIFF seeds by converting other image types, such as JPEG and BMP, using ImageMagick (v6.9.7).

We preprocessed each collection corpus to remove (a) duplicates identified by MD5 checksum, and (b) files larger than 300 KiB. The cutoff file size of 300 KiB is our best effort to conform to the AFL authors’ suggestions regarding seed size, while still having enough eligible seeds in the preprocessed corpora. We split audio files larger than 1 MiB into smaller files using FFmpeg (v3.4.8). In total, we collected 2,899,208 seeds across 14 different file formats. After preprocessing our collection corpus, we were left with a total of 1,019,688 seeds.

Experimental Setup

We run the Magma experiments on a cluster of AWS EC2 machines with 36-core Intel® Xeon® E5-2666 v3 2.9 GHz CPUs and 60 GiB of RAM. We conduct the FTS and real-world experiments on a Dell PowerEdge server with a 48-core Intel® Xeon® Gold 5118 2.30 GHz CPU, 512 GiB of RAM, and running Ubuntu 18.04.

Following Section 2.6, we run one fuzzing campaign per {program × fuzzer × initial corpus} combination (where a *program* is an executable driver for a particular *target*). Each fuzzing campaign consists of thirty independent 18 h trials.

Experiment

We evaluate the following six seed selection strategies against the previously-described targets and fuzzers:

FULL The collection corpus without minimization, preprocessed to remove duplicates and filtered for size (per Section 4.4.1).

EMPTY A per-target corpus comprising a single “empty” seed. Per Klees et al. [115], “the empty seed should be considered [when evaluating fuzzers], despite its use contravening conventional wisdom.” This seed is a zero-length file for six file types (JSON, MP3, REGEX, TIFF, TTF, and XML). For the remaining eight file types, the seed is a small file handcrafted to contain the bytes necessary to satisfy file header checks. This follows from the readelf experiments in Section 4.2.2, which demonstrates how poorly AFL performs when these header checks are not satisfied by the initial corpus. These files range in size from 11 B (SVG) to 13 KiB (OGG), with a median size of 51 B.

PROV The corpus provided with the benchmark (if any). This strategy is only applicable for the two fuzzer benchmarks (Magma and FTS).

MSET The corpus obtained using the UNWEIGHTED MINSET tool. We present UNWEIGHTED MINSET (rather than TIME or SIZE MINSET) because it finds more bugs than other MINSET configurations [177].

CMIN The corpus produced using AFL’s `af1-cmin` tool.

WOPT The *optimal minimum* corpus weighted by *file size*.

WMOPT The *weighted optimal minimum* corpus that takes into account edge frequencies. WMOPT attempts to *minimize* file sizes while *maximizing* an edge’s frequency count. We tried to implement an “optimal `af1-cmin`” (i.e., minimizing file size while treating the same edge with different hit counts as distinct

constraints), but EvalMaxSat was unable to find a solution (after 6 h) for many targets. Maximizing the total hit count for a given edge is a compromise and one that we hypothesize results in *deeper* state space exploration.

We exclude REDQUEEN’s uninformed, generic seed due to its poor performance in Section 4.2.2. We also explored an unweighted optimal minimum corpus, but found that EvalMaxSat produced the same corpora as WOPT for all but three targets (Magma’s *libpng*, the FTS’ *libarchive*, and the real-world *poppler* target). Thus, we also exclude unweighted minimal corpora from our results.

We compare the performance of each seed selection strategy across three measures:

Bug count. The ultimate aim of fuzzing is to find bugs in software. Thus, we use a direct bug count to answer RQ 2 and compare fuzzer effectiveness. We perform a manual triage of *all* crashes in the real-world target set. This is in contrast to many prior works [115, 122, 177, 214], which uses stack-hash deduplication to determine unique bugs from crashes—a technique known to both over- and under-count bugs [84, 108, 115].

Bug survival. We use survival analysis and the log-rank test as described in Section 2.6.2. When computing the restricted mean survival time (RMST) we use $N = 30$ repeated trials and an upper bound $T = 18$ h. The log-rank test is computed under the null hypothesis that two corpora share the same survival function. Thus, we consider two corpora to have statistically-equivalent bug survival times if the log-rank test’s p -value > 0.05 .

Code coverage. Coverage is often used to measure fuzzing effectiveness [23, 115, 144]. Like FUZZBENCH [144], we use Clang’s source-based region coverage metric to answer RQ 3 and compare coverage across fuzzers using different coverage maps. This is achieved by replaying each trial’s fuzzing queue through a version of the target compiled with Clang’s source-based coverage [209].

We report both the mean and 95 % bootstrap confidence interval (CI) of the percentage of regions executed across each campaign. Per Section 2.6.3, region coverage is compared across corpora using the Mann-Whitney U -test; a p -value > 0.05 implies a {program \times fuzzer \times corpus} combination yields a statistically-equivalent result compared to another.

TABLE 4.4: Corpora sizes. Each corpus is summarized by the number of files contained within (“#”) and total size (“S”, in MiB unless stated otherwise). The smallest *minimized* corpus for both “#” and “S” are highlighted in green and blue, respectively.

Target	FULL		PROV		MSET		CMIN		WOPT		WMOPT		
	#	S	#	S	#	S	#	S	#	S	#	S	
Magma	libpng	66,512	7,773.60	4	1.22 KiB	36	2.69	172	9.61	33	2.01	51	5.37
	libtiff	99,955	446.63	21	0.20	35	0.14	115	0.39	33	0.13	55	0.23
	libxml2	79,032	205.64	1,268	3.90	42	0.38	132	0.82	42	0.40	56	0.51
	php-exif	120,000	222.86	60	0.26	2	0.01	2	0.01	2	0.01	2	0.01
	php-json	19,978	76.46	55	4.23 KiB	17	0.72	212	4.29	17	0.85	30	1.91
	php-parser	75,777	224.51	2,934	0.92	606	1.94	2,229	11.43	569	1.71	1,187	7.86
	poppler	99,986	6,085.07	392	18.39	237	28.11	2,273	119.27	222	26.70	510	66.89
Google FTS	freetype2	466	35.50	2	2.83 KiB	43	5.40	246	20.92	37	5.04	53	6.95
	guetzli	120,000	222.86	2	544 B	17	0.04	463	0.60	13	0.03	51	0.10
	json	19,978	76.46	1	14 B	17	0.95	149	2.56	16	1.21	27	1.77
	libarchive	108,558	850.64	1	500 B	41	1.05	180	2.80	40	0.94	57	1.52
	libjpeg-turbo	120,000	222.86	1	413 B	3	0.01	93	0.11	3	0.01	13	0.02
	libpng	66,512	7,773.60	1	1.23 KiB	22	1.91	107	4.05	19	1.71	28	2.85
	libxml2	79,032	205.64	0	-	97	2.23	440	7.71	89	1.40	175	4.50
	pcre2	4,520	0.46	0	-	183	0.04	691	0.13	175	0.03	321	0.09
	re2	4,520	0.46	0	-	56	0.01	155	0.01	55	0.01	84	0.01
	vorbis	99,450	8,902.70	1	2.54 KiB	8	0.33	237	12.06	8	0.27	20	1.88
Real-world	freetype2	466	35.50	-	-	23	3.04	73	8.68	23	3.02	33	4.70
	libsvg	71,763	744.59	-	-	173	4.34	881	17.05	159	3.80	333	10.47
	libtiff	99,955	446.63	-	-	23	0.10	67	0.27	23	0.10	33	0.14
	libxml2	79,032	205.64	-	-	103	1.67	505	9.04	95	1.60	196	6.70
	poppler	99,986	6,085.07	-	-	189	22.70	1,318	121.90	177	22.04	381	50.30
	sox-mp3	99,691	4,094.22	-	-	9	0.17	137	3.75	6	0.30	15	0.64
sox-wav	74,000	2,490.61	-	-	10	0.39	68	1.65	9	0.27	14	0.49	

4.4.2 Minimization (RQ 1)

Table 4.4 shows the sizes of 14 collection corpora minimized across 21 target programs based on the code coverage measured by AFL. We reapplied the four corpus minimizers when fuzzing Magma with AFL++, as AFL++’s larger coverage map (250 KiB, compared to AFL’s 64 KiB) theoretically results in a more fine-grained coverage view. Indeed, we saw small variations (up to 10%) between the AFL and AFL++ minimized corpora, due to both the different-sized coverage maps and hash collisions inherent to AFL’s method for computing edges.

Across both fuzzers, corpora produced by CMIN are significantly larger than that produced by MSET (mean $8\times$ larger), WOPT (mean $9\times$ larger), and WMOPT (mean $4\times$ larger). We attribute this to CMIN distinguishing seeds with different edge frequency counts: MSET and WOPT only look at edges executed, ignoring the number of times these edges are executed, while WMOPT maximizes an edge’s frequency count. In comparison, MSET was only at most 37 seeds larger than WOPT (*php*’s parser), and on average only five seeds larger than WOPT. WMOPT corpora were (on average) twice as large as WOPT corpora. Finally, the minimized *php*’s exif corpora are notable because they discard 99% of the full JPEG corpus, due to a lack of diverse EXIF data. The small minimized exif corpora demonstrate the importance of selecting a

diverse range of initial inputs and minimizing large corpora.

In addition to generating smaller corpora, W[M]OPT incur lower run-time costs during minimization (compared to CMIN). We exclude the time required to trace the target and collect coverage data for each seed. WOPT’s minimization times range from 12 ms (*freetype2*) to 23 min (*libjpeg-turbo*), with a mean time of 141 s. WMOPT takes a similar amount of time: between 31 ms (*freetype2*) and 24 min (*php’s exif*), with a mean time of 143 s. In comparison, CMIN’s minimization times range from 12 s to 130 min, with a mean time of 25 min. Despite CMIN’s significantly slower minimization times, recall that (a) `af1-cmin` is a BASH script, while we wrote our optimal solver in C++, and (b) corpus minimization is a one-time upfront cost.

What ultimately matters is if the minimized corpora lead to better fuzzing outcomes. To this end, the following section discusses the bug-finding ability of the different corpus minimization techniques across our three benchmark suites (Magma, FTS, and a set of real-world targets) and two fuzzers (AFL and AFL++). We analyze these results with respect to the performance measures outlined in Section 4.4.1.

Finding

OPTIMIN produces significantly smaller corpora compared to existing state-of-the-art corpus minimization tools, while also incurring lower run-time costs.

4.4.3 Bug Finding (RQ 2)

Table 4.5 summarizes the bugs found in our Magma and FTS campaigns. For our real-world campaigns, we summarize the seven Common Vulnerabilities and Exposures (CVEs) assigned to us (for bugs found in these campaigns) in Table 4.6. In total, 77 (25 Magma, 15 FTS, and 33 real-world) bugs were found.

FTS Coverage Benchmarks

While this section’s main focus is on the bug-finding ability of each corpus, we first discuss six FTS “bugs” that are not actual bugs, but are instead code locations that the fuzzer must reach (marked with [†] in Table 4.5b). Notably, two of these locations are reached instantaneously (i.e., seeds in the fuzzing corpora reach the particular line of code without requiring any fuzzing) by most corpora *except* EMPTY. Naturally, EMPTY takes some time to reach the target locations, as AFL must construct valid inputs from “nothing”. Nevertheless, EMPTY reaches four of the six target locations within two hours (on average). A *libpng* location and the *freetype2* locations are never reached by EMPTY, because: (i) *freetype2* requires a valid composite glyph, which

TABLE 4.5: Bug-finding results, presented as the RMST in hours with 95 % bootstrap CI. Bugs never found have an RMST of \top (to distinguish bugs with an RMST of 18 h). The best performing corpus (corpora if the bug survival times are statistically equivalent per the log-rank test) for each target is highlighted in green (smaller is better).

(A) Magma bugs found by AFL and AFL++. We only report the RMST for bugs *triggered*. Bugs not triggered by any corpus are omitted (irrespective of whether the bug was reached or not). The *php json* and *parser* targets are omitted because no bugs were found.

Target	Bug	FULL		EMPTY		PROV		MSET		CMIN		WOPT		WMOPT	
		AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++
libpng	PNG001	\top	\top	16.83 ± 4.30	\top	16.96 ± 3.82	\top	17.95 ± 0.33	\top	\top	\top	17.54 ± 2.83	\top	17.47 ± 3.25	\top
	PNG003	1.93 ± 0.07	8.67 ± 3.45	\top	0.12 ± 0.05	0.0042 ± 0.01	0.0050 ± 0.01	0.01 ± 0.01	0.01 ± 0.01	0.02 ± 0.01	0.10 ± 0.08	0.01 ± 0.01	0.04 ± 0.01	0.01 ± 0.01	0.34 ± 0.13
	PNG006	\top	\top	\top	17.96 ± 0.26	\top	\top	\top	\top	\top	\top	\top	\top	\top	\top
	PNG007	17.04 ± 3.53	\top	\top	17.41 ± 3.61	6.85 ± 2.53	17.56 ± 2.69	11.03 ± 3.46	\top	12.26 ± 3.68	\top	10.65 ± 3.26	\top	10.04 ± 3.54	\top
libtiff	TIF001	17.77 ± 0.96	\top	\top	\top	\top	\top	15.98 ± 3.23	\top	15.81 ± 2.85	\top	13.93 ± 3.64	\top	15.20 ± 2.55	\top
	TIF002	16.85 ± 2.97	16.48 ± 4.03	17.65 ± 1.17	\top	15.29 ± 2.77	17.21 ± 2.93	16.39 ± 2.73	17.20 ± 3.10	12.81 ± 2.73	17.46 ± 3.27	14.09 ± 3.32	16.40 ± 4.25	14.50 ± 2.84	\top
	TIF007	0.34 ± 0.07	7.33 ± 4.14	1.27 ± 0.66	0.42 ± 0.25	0.05 ± 0.02	0.05 ± 0.03	0.03 ± 0.01	0.01 ± 0.01	0.03 ± 0.01	0.13 ± 0.01	0.03 ± 0.01	0.01 ± 0.01	0.04 ± 0.01	0.02 ± 0.01
	TIF008	16.96 ± 1.07	\top	17.02 ± 2.94	\top	16.23 ± 2.81	\top	12.53 ± 2.23	13.36 ± 3.85	7.91 ± 2.24	17.43 ± 3.50	8.68 ± 2.29	\top	9.39 ± 2.18	\top
	TIF012	3.64 ± 1.81	12.24 ± 3.66	2.45 ± 1.24	9.04 ± 3.93	0.47 ± 0.17	7.65 ± 2.79	0.65 ± 0.21	7.84 ± 3.03	0.53 ± 0.12	9.85 ± 3.31	0.68 ± 0.18	9.33 ± 3.04	0.70 ± 0.14	9.83 ± 3.11
TIF014	0.76 ± 0.15	11.34 ± 3.97	1.87 ± 1.19	0.73 ± 0.55	0.88 ± 0.36	6.46 ± 3.23	1.38 ± 0.78	3.08 ± 1.68	1.56 ± 0.72	5.16 ± 2.31	1.45 ± 0.35	2.60 ± 1.16	1.66 ± 0.51	1.39 ± 1.13	
libxml2	XML001	17.87 ± 0.82	\top	\top	\top	17.42 ± 2.49	\top	\top	\top	17.25 ± 2.10	\top	\top	\top	\top	
	XML003	8.69 ± 2.00	14.94 ± 2.89	\top	\top	14.93 ± 4.32	14.74 ± 3.69	\top	\top	\top	\top	\top	\top	\top	
	XML004	0.69 ± 0.12	0.85 ± 0.08	0.01 ± 0.01	0.02 ± 0.01	0.05 ± 0.01	0.02 ± 0.01	0.77 ± 0.30	0.08 ± 0.05	0.61 ± 0.29	0.15 ± 0.08	1.21 ± 0.41	0.77 ± 0.26	0.63 ± 0.27	0.80 ± 0.34
	XML009	15.99 ± 3.58	16.42 ± 3.26	17.94 ± 0.36	15.00 ± 2.75	16.44 ± 3.64	17.09 ± 2.93	17.55 ± 1.69	17.50 ± 2.19	15.33 ± 3.22	17.07 ± 2.58	15.51 ± 3.55	16.91 ± 2.39	14.31 ± 3.16	16.89 ± 2.01
	XML017	0.63 ± 0.12	1.32 ± 0.11	16.88 ± 2.66	2.60 ± 1.22	0.06 ± 0.02	0.38 ± 0.15	0.09 ± 0.01	0.09 ± 0.01	0.11 ± 0.03	0.07 ± 0.01	0.11 ± 0.01	0.07 ± 0.01	0.09 ± 0.01	0.08 ± 0.01
php-exif	PHP004	\top	\top	\top	\top	0.00 ± 0.01	16.20 ± 5.28	\top	\top	\top	\top	\top	\top	\top	
	PHP009	\top	\top	\top	\top	0.00 ± 0.01	0.00 ± 0.01	\top	\top	\top	\top	\top	\top	\top	
	PHP011	11.79 ± 2.09	7.76 ± 3.78	8.68 ± 2.32	0.05 ± 0.01	0.00 ± 0.01	13.74 ± 2.40	9.66 ± 2.67	0.05 ± 0.01	7.40 ± 2.16	0.06 ± 0.01	8.43 ± 1.97	0.06 ± 0.02	9.76 ± 2.77	0.05 ± 0.01
poppler	PDF002	\top	\top	\top	\top	17.43 ± 3.45	\top	17.86 ± 0.82	\top	\top	\top	\top	\top	\top	
	PDF005	\top	\top	\top	\top	\top	\top	17.46 ± 3.28	\top	17.18 ± 3.01	\top	16.93 ± 2.92	\top	15.40 ± 3.38	
	PDF010	\top	\top	\top	\top	0.83 ± 0.64	7.66 ± 4.15	4.06 ± 0.74	17.42 ± 3.54	5.54 ± 0.62	17.65 ± 1.13	4.51 ± 0.48	\top	8.07 ± 1.49	
	PDF016	\top	\top	0.0042 ± 0.01	0.0066 ± 0.01	0.08 ± 0.01	0.09 ± 0.02	0.14 ± 0.01	2.19 ± 0.14	1.11 ± 0.04	3.13 ± 0.17	0.14 ± 0.01	2.48 ± 0.17	0.30 ± 0.01	2.51 ± 0.16
	PDF018	\top	\top	\top	\top	6.17 ± 2.62	\top	\top	\top	\top	\top	\top	\top	\top	
	PDF021	\top	\top	\top	\top	\top	\top	17.61 ± 2.37	\top	17.67 ± 2.00	\top	\top	\top	\top	
	PDF022	\top	\top	\top	\top	\top	\top	17.79 ± 0.57	\top	\top	\top	17.90 ± 0.48	\top	\top	

(B) Bug-finding results (cont.). FTS bugs found by AFL and AFL++. IDs are derived from the order in which the bugs are presented in the target’s README (from the FTS repo). Bugs marked with [†] denote benchmarks that attempt to verify the fuzzer can reach a known location. Results with “-” indicate the FTS does not contain seeds for that target (see Table 4.4), and so we ignore it. The *vorbis* target is omitted because none of its three bugs were found.

Target	Bug	FULL		EMPTY		PROV		MSET		CMIN		WOPT		WMOPT	
		AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++
freetype2	A [†]	0.00 ± 0.00	0.00 ± 0.00	T	T	6.78 ± 2.12	6.10 ± 2.73	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
guetzli	A	T	T	T	T	10.25 ± 2.33	17.48 ± 1.38	15.88 ± 2.09	17.45 ± 3.37	17.78 ± 1.07	T	15.45 ± 2.82	T	16.92 ± 2.83	T
json	A	T	T	T	16.41 ± 4.24	0.09 ± 0.08	0.28 ± 0.27	17.64 ± 2.17	T	16.93 ± 3.92	17.57 ± 2.60	T	17.90 ± 1.43	T	17.90 ± 0.62
libarchive	A	T	T	T	15.31 ± 1.76	T	9.08 ± 0.90	4.44 ± 0.21	11.86 ± 1.74	9.39 ± 1.42	11.97 ± 1.90	12.50 ± 0.50	6.78 ± 0.44	7.38 ± 0.05	13.50 ± 1.75
libjpeg-turbo	A [†]	17.19 ± 2.37	T	1.92 ± 0.45	10.43 ± 2.73	3.00 ± 0.95	14.14 ± 2.47	3.36 ± 0.98	16.57 ± 2.70	3.82 ± 1.19	15.62 ± 2.75	4.71 ± 1.52	16.91 ± 2.70	3.68 ± 0.93	15.70 ± 2.11
libpng	A [†]	2.41 ± 0.01	0.0043 ± 0.0020	0.03 ± 0.01	0.11 ± 0.02	0.08 ± 0.08	0.21 ± 0.03	0.0051 ± 0.0029	0.09 ± 0.01	0.0072 ± 0.0036	0.0041 ± 0.0027	0.0025 ± 0.0017	0.0049 ± 0.0032	0.0038 ± 0.0025	0.08 ± 0.01
	B [†]	0.00 ± 0.00	0.00 ± 0.00	T	0.33 ± 0.04	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00
	C [†]	2.42 ± 0.02	0.0089 ± 0.0092	0.0003 ± 0.0004	0.07 ± 0.14	0.0008 ± 0.0007	2.47 ± 0.85	0.0046 ± 0.0029	0.0049 ± 0.0025	0.0070 ± 0.0035	0.0054 ± 0.0033	0.0023 ± 0.0017	0.0062 ± 0.0034	0.0036 ± 0.0022	0.0040 ± 0.0032
libxml2	A	4.46 ± 0.32	17.84 ± 0.58	T	16.46 ± 2.48	-	-	0.54 ± 0.06	3.10 ± 0.09	0.79 ± 0.12	4.72 ± 0.96	0.66 ± 0.07	4.70 ± 0.51	0.85 ± 0.23	7.81 ± 2.23
	B	16.46 ± 1.27	T	5.32 ± 1.22	13.43 ± 2.26	-	-	13.35 ± 1.95	T	10.57 ± 2.20	17.85 ± 0.91	8.93 ± 1.81	17.95 ± 0.28	14.46 ± 2.09	T
	C	T	T	T	T	-	-	T	T	17.53 ± 2.89	T	T	T	T	T
pcr2	A	2.57 ± 0.39	4.56 ± 0.36	1.88 ± 0.25	2.70 ± 0.44	-	-	2.15 ± 0.24	5.43 ± 0.74	2.07 ± 0.23	3.10 ± 0.46	1.59 ± 0.16	4.83 ± 0.47	1.39 ± 0.18	3.34 ± 0.44
	B	2.04 ± 0.73	5.34 ± 1.58	3.25 ± 0.70	5.71 ± 0.88	-	-	2.01 ± 0.68	3.83 ± 0.73	2.29 ± 0.65	3.98 ± 1.36	1.42 ± 0.42	3.47 ± 0.96	1.61 ± 0.72	3.39 ± 1.22
re2	A [†]	0.82 ± 0.16	2.91 ± 0.52	1.72 ± 0.31	9.72 ± 1.61	-	-	0.74 ± 0.52	3.30 ± 0.53	0.52 ± 0.12	3.85 ± 0.70	0.50 ± 0.16	7.52 ± 2.88	0.42 ± 0.09	3.35 ± 1.14
	B	16.20 ± 3.83	16.41 ± 2.49	17.66 ± 1.48	17.69 ± 1.87	-	-	11.52 ± 2.70	15.69 ± 2.53	11.85 ± 3.22	16.97 ± 2.35	17.08 ± 3.07	12.36 ± 2.62	15.91 ± 3.18	15.91 ± 2.55

TABLE 4.6: Real-world targets assigned a CVE.

Target	CVE	Description
libtiff	2019-14973	Elision of integer overflow check by compiler
poppler	2019-12293	Heap buffer overread
sox	2019-8354	Integer overflow causes improper heap allocation
	2019-8355	Integer overflow causes improper heap allocation
	2019-8356	Stack buffer bounds violation
	2019-8357	Integer overflow causes failed memory allocation
	2019-13590	Integer overflow causes failed memory allocation

EMPTY never synthesizes in the given timeframe, and (ii) *libpng* requires a specific chunk type (sRGB), which is difficult to synthesize without any knowledge of the PNG file format.

The only coverage benchmark that is not reached within minutes, *libjpeg-turbo*, is reliably reached by all corpora *except* FULL within the first five hours (on average) of each trial. The FULL corpus is highly unreliable on this target: it only reaches the target location in 10% of trials, and when it does, it takes double the time of the other corpora. This results in a high survival time of 17.19 h.

The EMPTY Seed

Of the 41 (14 Magma, 8 FTS, and 19 real-world) bugs EMPTY was able to find, it was the (equal) fastest to do so for 21 of these (5 Magma, 3 FTS, and 13 real-world). This result is particularly striking on *sox* (both MP3 and WAV), where EMPTY found the most bugs with the lowest RMSTs (including three of the CVEs in Table 4.6). However, EMPTY also suffers from the highest “false negative” rate: it is the most likely corpus to miss a bug when one exists (as evident from the number of \top entries in Table 4.5; the most of any corpus).

We hypothesize this low RMST is due to the reduced search space when mutating the empty seed, but that the mutation engine is less likely to “get lucky” in generating a bug-inducing input when starting from nothing. Indeed, for the three FTS bugs where EMPTY statistically outperforms the other corpora (*libjpeg-turbo*, *libpng*’s bug C, and *libxml2*’s bug B), EMPTY finds the bug with the lowest number of mutations (on average, half the number of mutations compared to the other corpora on these three bugs) while also achieving a comparable (and sometimes slightly slower) iteration rate than the other corpora (in particular, WOPT achieves a higher iteration rate than EMPTY on these three targets).

PROVided Seeds

The FTS PROV seeds—half of which are *singleton seeds*—are selected (by the FTS developers) based on their ability to trigger the target bug(s) within a few hours. For example, the FTS *json* bug is “usually found in about 5 minutes using the provided seed” [75] (confirmed by our results). While serving as a useful regression test (i.e., ensuring a new fuzzer does not perform worse than an existing one), this is not indicative of typical fuzzing practices, as (a) the location of bugs is unknown *a priori*, and (b) large seed sets are used in practice (per Section 4.2.3). Given the former, it is notable that the minimized corpora (CMIN, MSET, WOPT, and WMOPT) also

successfully found the same FTS bugs found by PROV, and even outperform the PROV corpus in half of these targets (*freetype2*, *libarchive*, and *libpng*).

The PROV corpus is the best performer at finding bugs in Magma: it triggers the most bugs—21 of the 25 bugs found by all fuzzers—and achieves the (equal) lowest RMST for 15 of these bugs. Similarly to FTS’s *freetype2* and *libpng*, all three *php* *exif* bugs were found without any mutation of the PROV seeds. Closer inspection of this corpus reveals why: PROV contains images that serve as regression tests for each of these three bugs (*bug77753.tif*, *bug77563.jpg*, *bug77950.jpg*, corresponding to bugs PHP004, PHP009, and PHP011, respectively). These regression tests immediately trigger their respective bugs, but with Magma’s *ideal sanitization*—“*in which triggering a bug immediately results in a crash*” [84]—disabled by default, these seeds do not cause a crash and hence are not excluded by AFL.⁵

Iteration Rates

Low iteration rates coupled with large corpora have a detrimental effect on a fuzzer’s ability to find bugs. For example, with FULL, *guetzli* achieves mean iteration rates of 0.84 and 0.74 execs/s for AFL and AFL++, respectively. At the other end of the spectrum, EMPTY achieves mean iteration rates of 229.23 and 167 execs/s (for AFL and AFL++, respectively), while the minimized corpora achieve iteration rates between 2 and 5 execs/s. FULL’s low iteration rate has a severe impact: both AFL and AFL++ fail to complete an initial pass over the 120,000 seeds in this corpus (in an 18 h trial), let alone perform any mutations and discover the bug. In comparison, the *guetzli* bug is found by all minimized corpora (CMIN, MSET, WOPT, and WMOPT) and PROV. We encounter similar results with *poppler*, where again neither AFL nor AFL++ complete a full pass over the collection corpus (resulting in no bugs triggered).

We find iteration rates vary significantly between fuzzers. For example, fuzzing Magma’s *libpng* with AFL and EMPTY achieves a mean iteration rate of 693 execs/s, while AFL++ achieves a mean iteration rate of 2,508 execs/s. Conversely, fuzzing the same target with AFL and WOPT achieves a mean iteration rate of 4,575 execs/s, compared to AFL++ at 351 execs/s. These results correlate with the bug survival times in Table 4.5a (where AFL++ outperforms AFL on the EMPTY seed, while AFL outperforms AFL++ with the WOPT corpus), suggesting higher iteration rates contribute to a fuzzer’s bug-finding ability (and at the very least, allow a fuzzer to more-quickly discard inputs that are not worth exploring).

⁵At the time of writing, this is a known issue flagged by the Magma developers, per <https://github.com/HexHive/magma/issues/54>.

When fuzzing with non-empty corpora, AFL achieves a higher iteration rate than AFL++. This is likely due to a combination of (a) more complex target instrumentation (where more of this instrumentation is being exercised with valid inputs), and (b) a coverage map $\sim 4\times$ larger than AFL's (which has L2 cache implications). These results further reinforce the need for corpus minimization when starting with a large collection corpus, particularly as fuzzer complexity increases.

Comparison to Previous Magma Evaluations

Our results improve on those originally presented by Hazimeh, Herrera, and Payer [84]. Nine of the bugs triggered during our 18 h trials (PNG006, TIF001, XML001, XML003, PDF002, PDF005, PDF018, PDF021, and PDF022) were *never* triggered by AFL/AFL++ in the original 24 h campaigns. Furthermore, two of these bugs (*poppler's* PDF005 and PDF022) were never found by *any* of the seven fuzzers originally evaluated by Hazimeh, Herrera, and Payer [84]. These two bugs were only found by the distilled corpora fuzzed with AFL and never with FULL, EMPTY, PROV, or AFL++. This is significant because over 200,000 CPU-hr were spent fuzzing Magma targets (across many 24 h and 7 d trials).

CVEs

Our real-world fuzzing campaigns led to the assignment of seven CVEs across three targets, summarized in Table 4.6 (our campaigns uncovered another 26 bugs, but these were already under investigation). Of these bugs, *libtiff's* CVE-2019-14973 is particularly interesting: discovered by all corpora, but found the fastest and most reliably by EMPTY (with an RMST of 6.02 h), this bug is only evident because we build our real-world targets for 32-bit x86. The *libtiff* maintainers report that the undefined behavior at the root of this bug does not manifest on 64-bit targets.

Our campaigns also reproduced an uncontrolled resource consumption bug in *libtiff* (CVE-2018-5784). This bug is caused by an infinite loop in the TIFF image directory linked list and is again found most frequently by EMPTY (11 out of 30 trials, resulting in an RMST of 12.37 h). In comparison, this bug is never found by FULL and MSET and only once by the other corpora. Notably, the initial EMPTY seed does not contain any image file directories, while all of the TIFF files in our minimized corpora do. We suspect AFL's mutations break existing directory structures (leading to parser failures), whereas EMPTY can construct a (malformed) directory list from scratch. These mutations eventually lead to a loop in the list, causing uncontrolled resource consumption.

Finding

Seed selection significantly impacts a fuzzer’s ability to find bugs. Both AFL and AFL++ perform better when bootstrapped with a minimized corpus, although the exact minimization tool is inconsequential. While both AFL and AFL++ find a similar number of bugs, AFL is generally faster to do so (and with less variance in bug-finding times).

4.4.4 Code Coverage (RQ 3)

Table 4.7 summarizes the coverage achieved overall Magma and FTS campaigns. For most targets, EMPTY explores less code than the other corpora: on average, half as much code (although this difference slightly decreases when fuzzing with AFL++). In particular, on targets accepting highly-structured input formats (e.g., *libxml2*, *poppler*), EMPTY explores less than half as much of the program’s code compared to the four minimized corpora. EMPTY’s results slightly improve when fuzzing with AFL++, likely due to the additional CmpLog instrumentation (reflecting our `readelf` results in Section 4.2.2). However, despite this improvement, EMPTY’s performance remains inferior to any other seed selection strategies.

After an 18 h trial, little distinguishes the code coverage achieved by the non-EMPTY corpora, and once again the four minimized corpora with AFL produced the best results. Curiously, the coverage gains AFL++ saw when fuzzing `readelf` with CMIN (Section 4.2.2) did not manifest in any of the five Magma targets: both AFL and AFL++ achieved similar levels of code coverage.

Finding

Seed selection has a significant impact on a fuzzer’s ability to expand code coverage. When fuzzing with the empty seed, more advanced fuzzers (such as AFL++) can cover more code. However, this advantage all but disappears when bootstrapping the fuzzer with a minimized corpus, as faster iteration rates become more critical. The exact minimization tool remains inconsequential.

4.4.5 Discussion

Selecting a corpus minimization tool. We evaluated three corpus minimization tools: MINSET, `af1-cmin`, and our OPTIMIN. Our results do not reveal a “best” minimization tool; while minimized corpora sizes varied markedly between tools, stochastic fuzzing variability means this ultimately has no statistically significant

TABLE 4.7: Code coverage, expressed as the mean region coverage (as a percentage of total program regions) with 95 % bootstrap CI. The best performing corpus (corpora if the code coverages are statistically equivalent per the Mann-Whitney U -test) for each target is highlighted in green (larger is better).

(A) Magma code coverage with AFL and AFL++.

Target	FULL		EMPTY		PROV		MSET		CMIN		WOPT		WMOPT	
	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++
libpng	28.71	25.56	15.26	19.10	25.19	26.55	30.00	29.53	29.98	29.60	30.03	29.39	29.08	28.81
	± 0.09	± 1.51	± 0.06	± 0.97	± 0.29	± 0.74	± 0.05	± 0.06	± 0.05	± 0.06	± 0.04	± 0.12	± 0.05	± 0.14
libtiff	44.37	44.76	35.10	27.41	44.77	41.72	45.40	44.45	46.09	44.38	45.70	44.38	45.86	43.90
	± 0.39	± 0.30	± 2.23	± 2.05	± 0.25	± 0.43	± 0.23	± 0.22	± 0.23	± 0.39	± 0.26	± 0.26	± 0.29	± 0.32
libxml2	21.10	20.39	10.60	14.99	22.47	22.74	19.47	19.84	20.75	20.64	19.53	19.14	20.06	19.07
	± 0.29	± 0.31	± 0.29	± 0.44	± 0.22	± 0.37	± 0.14	± 0.20	± 0.21	± 0.32	± 0.22	± 0.10	± 0.14	± 0.07
php-exif	2.24	2.34	2.28	2.36	2.37	2.37	2.25	2.36	2.30	2.36	2.29	2.36	2.26	2.36
	± 0.04	± 0.01	± 0.04	± 0.00	± 0.00	± 0.00	± 0.04	± 0.00	± 0.03	± 0.00	± 0.03	± 0.00	± 0.04	± 0.00
poppler	35.96	35.95	1.49	1.76	41.12	38.35	41.40	36.74	41.13	37.92	41.30	36.69	41.44	36.85
	± 0.00	± 0.00	± 0.00	± 0.01	± 0.05	± 0.06	± 0.06	± 0.30	± 0.06	± 0.32	± 0.07	± 0.27	± 0.07	± 0.33

(B) FTS code coverage with AFL and AFL++.

Target	FULL		EMPTY		PROV		MSET		CMIN		WOPT		WMOPT	
	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++	AFL	AFL++
freetype2	48.01	44.65	17.31	33.85	34.99	45.41	47.58	44.84	47.84	44.11	47.82	43.24	47.73	45.12
	± 0.12	± 0.41	± 0.16	± 0.78	± 0.17	± 0.59	± 0.14	± 0.34	± 0.15	± 0.42	± 0.20	± 0.24	± 0.16	± 0.21
guetzli	67.34	67.34	31.05	30.31	75.55	73.61	72.95	69.84	73.12	70.76	72.93	70.16	72.79	69.92
	± 0.00	± 0.00	± 0.14	± 0.14	± 0.12	± 0.08	± 0.14	± 0.10	± 0.11	± 0.07	± 0.10	± 0.07	± 0.11	± 0.08
json	90.29	90.87	90.00	90.18	90.32	90.74	90.60	90.96	90.62	91.14	90.82	90.96	90.79	91.05
	± 0.08	± 0.11	± 0.34	± 0.31	± 0.26	± 0.22	± 0.06	± 0.06	± 0.07	± 0.08	± 0.10	± 0.07	± 0.08	± 0.06
libarchive	17.14	16.67	17.84	23.09	18.04	24.63	18.51	20.14	18.54	22.18	18.65	21.55	18.36	21.61
	± 0.12	± 0.47	± 0.46	± 0.77	± 0.63	± 0.68	± 0.19	± 0.39	± 0.22	± 0.24	± 0.10	± 0.21	± 0.20	± 0.27
libjpeg-turbo	16.31	15.38	18.56	18.01	18.84	18.13	20.53	19.24	20.54	19.50	20.47	19.25	20.41	19.01
	± 0.30	± 0.08	± 0.09	± 0.07	± 0.24	± 0.07	± 0.11	± 0.32	± 0.11	± 0.24	± 0.12	± 0.30	± 0.11	± 0.38
libpng	32.81	34.87	19.30	29.66	25.40	33.83	34.92	36.62	35.09	36.82	34.97	36.95	34.94	36.63
	± 0.12	± 0.10	± 0.00	± 0.23	± 0.00	± 0.21	± 0.09	± 0.18	± 0.06	± 0.13	± 0.06	± 0.15	± 0.07	± 0.12
libxml2	14.90	14.49	6.77	8.12	-	-	15.81	15.08	15.91	14.95	16.01	15.06	15.76	15.29
	± 0.09	± 0.03	± 0.09	± 0.49	-	-	± 0.15	± 0.12	± 0.13	± 0.03	± 0.02	± 0.15	± 0.15	± 0.18
pcr2	60.81	63.97	59.65	63.13	-	-	60.94	63.21	61.04	63.64	61.09	63.03	61.00	62.92
	± 0.16	± 0.14	± 0.15	± 0.20	-	-	± 0.16	± 0.22	± 0.22	± 0.24	± 0.23	± 0.16	± 0.16	± 0.22
re2	59.00	58.96	59.26	58.13	-	-	59.05	59.03	59.08	59.00	59.10	57.40	59.05	59.00
	± 0.07	± 0.06	± 0.16	± 0.53	-	-	± 0.06	± 0.05	± 0.04	± 0.06	± 0.05	± 0.83	± 0.08	± 0.07

impact on fuzzing outcomes (with respect to both bug finding and code coverage). However, a minimized corpus is always better due to the faster iteration rate. While our results show that this may not necessarily find more bugs in a given trial, the fuzzer can more quickly discard inputs not worth exploring. We, therefore, recommend the adoption of OPTIMIN, given the considerably smaller corpora it produces.

When to use the empty seed. While our results demonstrate that corpus minimization achieves the best results, there were nine occasions (three in each of the Magma, FTS, and real-world benchmarks) where EMPTY performed as well as or better than the minimized corpora. These occasions correspond to when coverage is at its lowest, suggesting these are shallow bugs. Thus, we recommend an additional campaign with the empty seed to quickly weed out shallow bugs where possible.

Corpus minimization as lossy compression. Prior work [183, 214] shows how different coverage metrics can affect fuzzing results in practice. Similarly, corpus minimization can also use coverage metrics not solely based on code coverage (or approximate edge coverage in AFL’s case). Corpus minimization based solely on code coverage is effectively a form of *lossy compression* [59]: program states may be discarded if they do not expand code coverage. Indeed, we saw this in Section 4.4.2, where the AFL/AFL++ corpus sizes differed due to different-sized coverage maps. We leave it to future work to explore how corpus minimization generalizes to other coverage metrics.

Generalizing to other fuzzers. We limit our experiments to two coverage-guided, mutational greybox fuzzers: AFL and AFL++. We selected AFL because it is widely evaluated and deployed, while AFL++ is an updated and maintained version of AFL incorporating broad improvements from recent fuzzing research and regularly outperforms *all* other fuzzers on Google’s FUZZBENCH [144]. While it is unclear how our results might generalize to other fuzzers (e.g., honggfuzz [207] and libFuzzer [188], both of which provide corpus minimization capabilities), we believe our AFL++ results—which demonstrate how seed selection practices impact a range of recent advances in fuzzing research—are generalizable to other mutation-based greybox fuzzers. We leave it to future work to confirm this.

4.5 Chapter Summary

In this chapter we presented, to the best of our knowledge, the first in-depth analysis of the impact seed selection has on mutation-based coverage-guided greybox fuzzing.

We argue that the choice of fuzzing corpus is a critical decision—often overlooked—made before a fuzzing campaign begins. Our results provide ample confirmation of this criticality. In particular, we demonstrate how fuzzing outcomes vary significantly depending on the initial seeds used to bootstrap the fuzzer; especially with respect to state space exploration.

Now that the fuzzer is bootstrapped with an appropriate seed corpus, we turn our attention to running the fuzzer and the instrumented target (② and ③ in Fig. 1.1, Chapter 1). In particular, in the next chapter we introduce a new coverage metric: *def-use* chains.

Chapter 5

Data-Flow-Guided Fuzzing

In Chapter 3 we surveyed the coverage metrics fuzzers use to abstract a target’s state space and measure state space search. We found that most fuzzers use coverage metrics derived from *control-flow* features (e.g., basic block, edge). However, in some targets, control flow offers only a coarse-grained approximation of program behavior; instead, *data flow* provides a more fine-grained view.

5.1 Introduction

Exploiting dynamic information drives fuzzer efficiency. Tracking code executed in a target allows a fuzzer to focus its mutations on inputs that reach new code. However, per Section 1.2, tracking code coverage (i.e., control-flow coverage) only covers one dimension of a program’s state space. Is this sufficient? We answer this question here.

Chapter outline. We begin by motivating the need for data-flow-guided fuzzers (Section 5.2). Then, we introduce our approach for tracking run-time data flows with low overhead: the DATAFLOW fuzzer (Section 5.3). Inspired by AFL-Sensitive [214], DATAFLOW provides a tunable sensitivity range, balancing the computational cost of exploration with precision. We discuss this in Section 5.3.1, followed by a discussion of DATAFLOW’s implementation in Section 5.4. Finally, we evaluate DATAFLOW in Section 5.5. This evaluation compares DATAFLOW to state-of-the-art fuzzers driven by control flow, taint analysis, and data flow.

5.2 Motivation

In some targets, control flow offers only a coarse-grained approximation of program behavior. This includes targets whose control structure is decoupled from its semantics (e.g., LR parsers generated by yacc [224]). Such targets require *data-flow*

coverage [65, 88, 101, 175, 202, 224] to accurately capture program behavior. Whereas control flow focuses on the order of operations in a program (i.e., branch and loop structures), data flow instead focuses on how variables (i.e., data) are defined and used [175]; indeed, there may be no control dependence between variable definition and use sites.

5.3 Design of a Data-Flow-Guided Fuzzer

A greybox fuzzer should maintain *accurate* coverage information without negatively impacting *performance*. These requirements exist irrespective of the coverage metric used. With this in mind, we describe: (i) a theoretical foundation for constructing data-flow-based coverage metrics; (ii) how DATAFLOW incorporates these observations; and (iii) the implementation of a DATAFLOW prototype.

5.3.1 Coverage Sensitivity

We define data-flow coverage as follows:

Definition

Data-flow coverage is the tracking of *def-use* chains executed at run time.

This definition allows us to explore data-flow-based coverage metrics with different *sensitivities* [183, 214]. We follow the program analysis literature and define sensitivity as a coverage metric's ability to discriminate between a set of program behaviors [114]. In fuzzing, a coverage metric's sensitivity is its ability to preserve a chain of mutated test cases until they trigger a bug [214]. Different sensitivities allow us to balance efficacy and performance: more sensitive metrics incur higher performance penalties. For example, edge coverage sensitivity is increased by incorporating function call context [37]. However, this requires additional instrumentation, increasing run-time overhead [183].

Like traditional data-flow analysis, our data-flow coverage metric requires identifying variable *def* and *use* sites. Following Horgan and London [93], we define a data-flow variable *def* site as a name referring to storage allocated statically (e.g., storage class `static`, `global`) or automatically (i.e., local to a procedure). We deviate from this definition by: (i) including calls to dynamic memory allocation routines (e.g., `malloc`); and (ii) excluding reallocations/reassignments that would traditionally *kill* a definition. Instead, *defs* are only killed when they (a) go out of scope (e.g., a local variable in a returning procedure), or (b) are explicitly deallocated (e.g., via `free`).

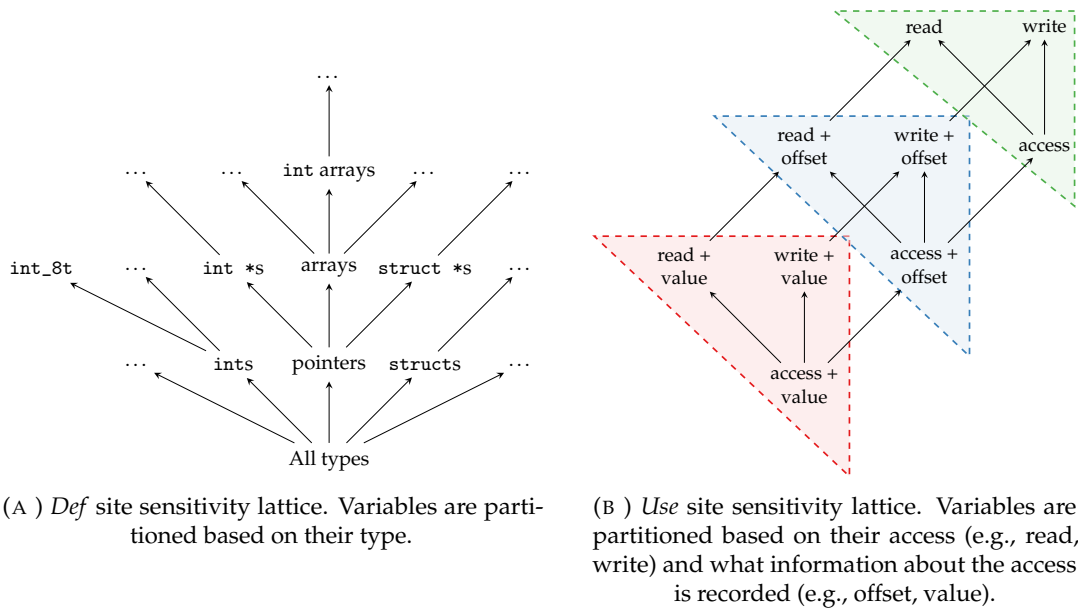


FIGURE 5.1: *Def* and *use* site sensitivity lattices. The sensitivity of coverage metrics increases towards the bottom.

Consequently, a *use* site includes both reads/writes from/to a *def* site. We deviate from the classic definition to ensure scalability: the difficulties of scaling data-flow analyses on real-world programs are well known [29, 88, 202]. We believe reducing precision by not killing definitions (when assigning a new value to a variable) is a suitable trade-off to maintain scalability.

Once we identify *def* and *use* sites, DATAFLOW instruments these sites (using compiler-based instrumentation, discussed in Section 5.4) so *def-use* chains can be tracked at run time. However, exactly which *def-use* sites are instrumented (and hence which are tracked) depends on the required sensitivity. Inspired by Wang et al. [214], this leads us to define a pair of sensitivity lattices—one for *def* sites and another for *use* sites, in Fig. 5.1—that can be composed to achieve the desired overall sensitivity (we discuss related limitations in Section 5.4.4). The remainder of this section reuses the code from Fig. 1.2a, which we reproduce in Fig. 5.2 for convenience.

Def Site Sensitivity

Complete data-flow coverage requires identifying and instrumenting *all* variable *def* sites. Unfortunately, the overhead to achieve this level of sensitivity is prohibitively expensive [28]. Therefore, a method for identifying (and hence instrumenting) a subset of important program variables is required. Ideally, this would be an (almost entirely) automated process, reducing the developer burden on the user.

```
1 static char buf[128];
2
3 void foo(int a, size_t b) { memset(buf, a, b); }
4 void bar(int a, size_t b) { foo('A', sizeof(buf)); }
5 void baz(int a, size_t b) { bar(a, sizeof(buf)); }
6
7 const struct handler_t {
8     char code;
9     void (*handler)(int, size_t);
10 } handlers[] = {{0x66, foo}, {0x61, bar}, {0x7a, baz}};
11
12 int sig_cmp(void *data, size_t start, size_t len) {
13     if (start + 3 > len)
14         return 1;
15
16     static const char sig[3] = {0xa, 0xb, 0xc};
17     return memcmp(&data[start], sig, 3);
18 }
19
20 int main(int argc, char *argv[]) {
21     struct stat st;
22
23     int fd = open(argv[1], O_RDONLY);
24     fstat(fd, &st);
25     size_t size = st.st_size;
26     char *data = malloc(size);
27     read(fd, data, size);
28
29     if (sig_cmp(data, 0, size))
30         return 1;
31
32     for (unsigned i = 4; (i + 3) <= size; i += 3)
33         for (unsigned j = 0; j < 3; ++j)
34             if (data[i] == handlers[j].code)
35                 handlers[j].handler(data[i + 1], data[i + 2]);
36
37     free(data);
38     close(fd);
39
40     return 0;
41 }
```

FIGURE 5.2: Source code for the simple dispatch table from Fig. 1.2a .

One approach is to partition *def* sites by *type* and restrict instrumentation to *def* sites of a given type (or type set). Figure 5.1a shows the sensitivity lattice for this type-based partitioning.

Partitioning *def* sites by type has several advantages. For example, instrumenting array variables focuses the fuzzer on memory-safety vulnerabilities. Similarly, tracking the data flow of structs may allow for the discovery of type confusion vulnerabilities [105, 197]. Type-based partitioning requires some upfront knowledge of the target to ensure meaningful variables are tracked at run time. For example, the fuzzer may miss important program behaviors (and hence bugs) if “uninteresting” variables are tracked (e.g., *st* at line 21 in Fig. 5.2).

Requirement

Tracking *all* data flows is prohibitively expensive. Identification (and instrumentation) of only important variables is required.

Use Site Sensitivity

Figure 5.1b shows the *use* site sensitivity lattice. Variables are either read from or written to (i.e., “accessed”). Variable accesses are strictly more sensitive than writes or reads on their own. The simplest and least sensitive metrics only track when a variable is accessed (shown at the top of the lattice).

Conversely, the most sensitive data-flow coverage metrics are ones that track not only *when* a particular variable is accessed, but the *value* of that variable when accessed. For example, considering line 17 in Fig. 5.2, this is the difference between reading from data and reading the value 0xa from data. The latter is akin to traditional data-flow testing, which focuses on the values that variables take at run time [175, 202], and is similar to GREYONE, which monitors (a subset of) program variables and their values to infer taint [68]. Depending on the *def* site sensitivity, this approach will quickly saturate the fuzzer’s coverage map (due to the path collision problem [67]); a middle ground between this overly sensitive approach and simple accesses is required.

We achieve this middle ground by incorporating more fine-grained spatial information into a variable’s *use*. This is particularly useful when *def* sites include arrays and/or structs (e.g., *handlers* in Fig. 5.2), as *def-use* chains are now differentiated by the *offset* at which an array/struct is accessed (analogous to a field-sensitive static analysis).

Requirement

Information at different granularities is recorded at *use* sites. Care is required when recording more precise information to ensure the coverage map does not saturate, clogging the fuzzing queue.

Composing Sensitivity Lattices

Different *def-use* sensitivities can be composed to track data flow at different granularities. We reuse the code in Fig. 1.2a —reproducing it in Fig. 5.2 for convenience—to illustrate this. For example, given the *def* sensitivity lattice in Fig. 5.1a, the following may be tracked: (i) all variables; (ii) integer variables (e.g., `fd`, `i`); (iii) arrays (e.g., `handlers`); or (iv) each `struct` in `handlers`. To simplify our presentation, we restrict *def* site instrumentation (and hence *def-use* chain tracking) to the `handlers` array. This leads to varying *def-use* chains depending on the *use* site sensitivity.

Simple access. The **green** region in Fig. 5.1b. Tracks when `handlers` is accessed (lines 34 and 35 in Fig. 5.2). This results in two *def-use* chains: line 10 \rightsquigarrow line 34, and line 10 \rightsquigarrow line 35. However, this provides a poor approximation of program behavior: information about the specific `handler` executed is lost.

Access with offset. The **blue** region in Fig. 5.1b. Tracks when `handlers` is accessed along with the offsets where `handlers` is accessed (at index `j`). This provides a more complete view of how `handlers` is used with negligible overhead. This is similar to MEMFUZZ’s approach, which incorporates memory accesses into code coverage [49]. This results in 2×3 *def-use* chains: one for each read at each index `j`.

Access with value. The **red** region in Fig. 5.1b. Tracks when `handlers` is accessed along with the values (being read) during these accesses. This is the most sensitive *use* site coverage metric and achieves the goal of traditional data-flow coverage: associate values with variables, and how these associations can affect the execution of the target [175]. This is similar to GREYONE’s “taint inference”, which looks at the value of variables used in path constraints [68].

Again, this level of sensitivity results in six *def-use* chains. Here, `handlers`’ value range is fully deterministic. However, in general these values will depend on user input, resulting in rapid saturation of the fuzzer’s coverage map.




FIGURE 5.3: High-level overview of DATAFLOW.

Requirement

Sensitivity lattice composition must balance efficacy and performance: too precise, and the fuzzer’s coverage map will saturate, reducing throughput.

By composing *def* and *use* sensitivity lattices, we realize a variety of data-flow-based coverage metrics. We do this in our fuzzer, DATAFLOW, described in the following sections.

5.4 DATAFLOW Implementation

Figure 5.3 depicts DATAFLOW’s high-level architecture, including: (i) compiler instrumentation (built on LLVM v12) for capturing *def-use* sites at the desired sensitivity (Sections 5.4.1 and 5.4.2); and (ii) a runtime library for feeding data-flow information to the fuzzing engine (Section 5.4.3).

Our architecture is agnostic to the underlying fuzzer; the instrumented target produced by the compiler (and linked with the `fuzzalloc` runtime library) can be executed by any American Fuzzy Lop (AFL)-based fuzzer (i.e., any fuzzer using an AFL-style coverage map). However, instead of recording and tracking control-flow coverage, the fuzzer’s coverage map tracks data-flow coverage. DATAFLOW is available at <https://github.com/HexHive/datAFLOW>.

5.4.1 *Def-use* Site Identification

We must first identify *def* and *use* sites so that data flows between these sites can be tracked. Per Section 5.3.1, *def* site selection impacts coverage sensitivity: more instrumented *def* sites leads to more complete data-flow coverage. We implement several *def* site instrumentation schemes based on the type-based partitioning described in Section 5.3.1.

We make the following assumptions during *def-use* site identification. First, we assume debug metadata is available in the LLVM intermediate representation (IR).

We use this metadata to identify and limit variable *def* sites to source-level variables. Second, we assume tracked variables are accessed via memory references (i.e., load/store instructions), rather than registers. This is automatic for most composite types (e.g., arrays). For primitive types (e.g., integers), this requires demoting registers to memory references (via LLVM’s `reg2mem` pass).

The first assumption reduces the number of potential data flows and is adopted from prior work [61, 141]. The second assumption limits *use* sites to memory access instructions, simplifying instrumentation. We apply existing LLVM transforms to limit *use* sites to two instructions: loads and stores.¹ Exactly which instructions we instrument depends on the *use* sensitivity required (configured at compile time). We describe our instrumentation in Section 5.4.2.

5.4.2 Def-use Tracking

We reduce the run-time tracking of *def-use* chains to a metadata management problem. Here, *def* site identifiers are the metadata requiring efficient retrieval at *use* sites. Inspired by AFL’s approach for tracking edge coverage—where basic blocks (in the LLVM IR) are statically assigned a random 16-bit integer—we statically “tag” *def* sites (again, in the LLVM IR) with a random 16-bit integer (Section 5.4.2). This tag is then propagated to *use* sites, where it is retrieved and used to construct a *def-use* chain (Section 5.4.3).

Def Site Instrumentation

We adopt Padding Area MetaData (PAMD) [128] for tracking *def-use* chains. PAMD extends *baggy bounds checking*, a technique proposed by Ding et al. [50] for protecting C and C++ code against buffer overruns. PAMD attaches inline metadata to memory objects (hence our assumption that tracked variables are accessed via memory references; Section 5.4.1) and provides constant-time lookup of this metadata. This lookup occurs via the “baggy bounds table”, which stores the binary logarithm of an object’s size and alignment (denoted e). Once e is retrieved from the baggy bounds table, the base and size of an object pointed to by p is computed using:

$$\text{base} = p \ \& \ \sim(2^e - 1) \tag{5.1}$$

$$\text{size} = 2^e \tag{5.2}$$

Equations (5.1) and (5.2) require an object’s size and alignment to be a power-of-two. To meet this requirement, PAMD pads static objects (i.e., stack and global variables)

¹We lower atomic memory intrinsics and expand `llvm.mem*` intrinsics so we can focus on load/store instructions (both of which are trivial to identify and hence instrument).




FIGURE 5.4: PAMD’s approach for inline metadata. “Object” is aligned to a power-of-two boundary and “padding” is inserted to ensure *size* is a power-of-two.

before attaching the *def* site tag. Figure 5.4 illustrates this process. For example, given a 4-byte object, then $size = 8$, $e = 3$ (the binary logarithm of *size*), and two bytes of padding is inserted before the tag.

Objects whose padding or overall size becomes too large for static allocation are “*heapified*” (i.e., move to the heap). We adopt CCured’s [154] approach to heapify objects. For heap-allocated objects (including heapified objects), calls to `malloc`, `calloc`, and `realloc` are replaced with tagged versions (e.g., `__bb_malloc`) accepting the 16-bit tag as an additional argument. Figure 5.5 demonstrates our *def* site instrumentation on a static global variable.

Figure 5.5 shows a snippet of the (un)instrumented LLVM IR for the simple dispatch table in Fig. 1.2 (Chapter 1). We focus our *def* site instrumentation on the global `handlers` array. During compilation, DATAFLOW resizes `handlers` to meet PAMD’s object size requirement. Here, 14 B of padding are inserted before the two-byte tag (line 8 in Fig. 5.5b). DATAFLOW then tags this *def* site with the identifier 1337 (line 9). `Handlers` is a global variable, so registration in the baggy bounds table occurs via a constructor (line 13) inserted by the compiler.

Use Site Instrumentation

Per Section 5.4.1, *use* sites are limited to `load` and `store` instructions in the LLVM IR (e.g., line 12 in Fig. 5.6a). We instrument these instructions with a call to `__hash_def_use`, which retrieves the object’s size from the baggy bounds table and uses this size to retrieve the *def* tag. The size is also used to determine the offset at which an object is accessed (enabling the *access with offset* sensitivity described in Section 5.3.1). Like *def* sites, *use* sites are tagged at compile time with a randomly-generated identifier (e.g., 4242 at line 12 in Fig. 5.6b). Finally, we leverage several techniques from AddressSanitizer (ASAN) [187] to limit the number of *use* instrumentation sites, thereby reducing overhead without sacrificing precision. We describe the

```

1 ; const struct handler_t handlers[] =
2 ;   {'f', foo}, {'a', bar}, {'z', {baz}}
3 @handlers = [3 x %struct.handler_t] [
4   %struct.handler_t {i8 102, void (i32, i64)* @foo},
5   %struct.handler_t {i8 97, void (i32, i64)* @bar},
6   %struct.handler_t {i8 122, void (i32, i64)* @baz}
7 ], align 16

```

(A) Original code.



```

1 ; Assign handlers the tag "1337"
2 @handlers = <{ [3 x %struct.handler_t], [14 x i8], i16 }> <{
3   [3 x %struct.handler_t] [
4     %struct.handler_t {i8 102, void (i32, i64)* @foo},
5     %struct.handler_t {i8 97, void (i32, i64)* @bar},
6     %struct.handler_t {i8 122, void (i32, i64)* @baz}
7   ], ; The original array
8   [14 x i8] zeroinitializer, ; Padding
9   i16 1337 ; Tag
10 }>, align 64
11
12 ; Register the allocation in the baggy bounds table
13 define void @fuzzalloc.ctor() {
14 entry:
15   %cast = bitcast @handlers to i8*
16   call void @_bb_register(i8* %cast, i64 64)
17   ret void
18 }

```

(B) Instrumented code.

FIGURE 5.5: Example *def* site instrumentation of the simple dispatch table in Fig. 1.2a . DATAFLOW replaces the handlers global array with a tagged version (tag = 1337) and registers this allocation in the baggy bounds table (in the `fuzzalloc.ctor`).

```

1 ; data[i] == handlers[j].code
2
3 ; Load data[i]
4 %data_idx = getelementptr inbounds i8*, i8* @data, i64 %i
5 %0 = load i8, i8* %data_idx
6
7 ; Load handlers[j].code
8 %handlers_idx = getelementptr inbounds [3 x %struct.handler_t],
9   [3 x %struct.handler_t]* @handlers, i64 0, i64 %j
10 %code = getelementptr inbounds %struct.handler_t,
11   %struct.handler_t* %handlers_idx, i32 0, i32 0
12 %1 = load i8, i8* %code
13
14 ; Check equality
15 %cmp = icmp i8 %0, %1

```

(A) Original code.



```

1 ; Instrument handlers' use with a call to __hash_def_use
2
3 %data_idx = getelementptr inbounds i8*, i8* @data, i64 %i
4 %0 = load i8, i8* %data_idx
5
6 ; Load handlers[j].code and instrument use
7 %handlers_idx = getelementptr inbounds [3 x %struct.handler_t],
8   [3 x %struct.handler_t]* @handlers, i64 0, i64 %j
9 %code = getelementptr inbounds %struct.handler_t,
10   %struct.handler_t* %handlers_idx, i32 0, i32 0
11 %1 = load i8, i8* %code
12 call void @__hash_def_use(i16 4242i8* %code, i64 1)
13
14 %cmp = icmp i8 %0, %1

```

(B) Instrumented code.

FIGURE 5.6: Example *use* site instrumentation of the simple dispatch table in Fig. 1.2a. DATAFLOW tags (tag = 4242) and instruments the *use* site with a call to `__hash_def_use`.

internals of `__hash_def_use`, and how it integrates with the fuzzer, in the following section.

5.4.3 Fuzzer Integration

The `__hash_def_use` function constructs a *def-use* chain by hashing together the *def* and *use* sites. This hash is used as a lookup into the fuzzer’s coverage map to guide the fuzzer toward discovering new data flows. This is analogous to AFL tracing edges to discover new control flow paths. Consequently, we leverage techniques used by traditional greybox fuzzers (e.g., compact bitmaps) to efficiently record data-flow coverage [138].

In particular, we use *coarse* data-flow coverage metrics—*def-use* chain hit counts stored in a compact bitmap—to achieve efficient fuzzing. While these techniques result in path collisions [67], we are willing to tolerate such imprecision to limit overhead costs. Coarse coverage metrics also lower implementation costs, enabling the reuse of existing fuzzing engines (here, AFL++ [63]).

We adopt AFL’s hashing process (Eq. (3.1) in Section 3.2.2) for looking up data flows in the fuzzer’s coverage map. *What* is hashed and *how* it is hashed varies according to the desired sensitivity (Section 5.3.1):

Simple access. Xor of the *def* and *use* site tags:

$$p \leftarrow \text{def} \oplus \text{use} \quad (5.3)$$

Access with offset. The *def* site tag, *use* site tag, and the offset being accessed. The offset (e.g., array index, struct offset) is computed by subtracting the base address—found using Eq. (5.1)—from pointer *p*. We compute the hash as:

$$p \leftarrow \text{def} \oplus (\text{use} + \text{offset}) \quad (5.4)$$

Access with value. The *def* site tag, *use* site tag, the offset, and the value accessed. The *def/use* tags and offset are left-shifted to allow room for the value hash and reduce collisions. The accessed value is divided into single-byte chunks $\{v_0, v_1, \dots\}$ that are hashed into the *def-use* chain:

$$p \leftarrow (\text{def} \oplus (\text{use} + \text{offset}) \ll 2) \oplus (v_0 \oplus v_1 \oplus \dots) \quad (5.5)$$

This is implemented as a loop, resulting in a double-load of the accessed object (in addition to the load in the original code). We implement the `__hash_def_use`

function so that uninstrumented data flows (i.e., those without an entry in the baggy bounds table) are bucketed in their own coverage map entry.

5.4.4 Limitations

Def Site Selection

Our *def* site selection approach (Section 5.3.1) is incomplete: important data flows may be missed if the appropriate *def* sites are not instrumented. Per our *def* site sensitivity lattice, our prototype focuses on composite types (i.e., arrays and structs) and eschews instrumenting primitive types (e.g., integers). While this approach may miss important data flows, we accept this trade-off, given (a) memory safety remains a key concern [147], and (b) the prohibitive run-time overheads when tracking all *def* sites.

Custom Memory Allocators

Identifying *def* sites is complicated because many applications do not directly call the standard allocation routines (e.g., `malloc`), but indirectly through a custom memory allocator. For example, standard memory allocation routines may be wrapped in other functions. These functions may then be indirectly called via global variables/aliases, stored and passed around in structs, or used as function arguments.

To address the challenge imposed by custom memory allocators and memory allocation patterns, DATAFLOW allows the user to specify wrapper functions to tag (in addition to the standard allocation routines). While DATAFLOW requires the user to find these wrappers manually, existing techniques [38] could assist in this process. We wrap these memory allocation routines within *trampoline* functions when their address is taken (e.g., stored in a global variable). Rather than a compile-time *def* site tag (which may not be statically computable), these trampolines revert to using the lower 16-bits of the PC as the *def* site tag. This approach avoids the need for expensive and imprecise static analysis (e.g., to track the access of memory allocators through global variables).

C++ Dynamic Memory Allocation

To simplify our instrumentation, we rewrite C++ `new` calls as `malloc` calls. However, this prevents us from handling any `std::bad_alloc` exceptions, meaning any failed allocations will cause a program crash (irrespective of any exception handlers in place). Such false negatives are removed by replaying crashing inputs through the original target.

Coverage Imprecision

Storing coarse coverage information in a compact bitmap is inherently inaccurate and incomplete [67]. While this may limit DATAFLOW’s ability to discover and explore data flows, this limitation is not unique to DATAFLOW, and affects many greybox fuzzers [7, 8, 37, 49, 61, 68, 97, 134, 141, 214, 217, 231].

5.5 Evaluation

We perform an extensive evaluation (over 3 CPU-yr of fuzzing) to answer the following research questions:

- RQ 1** Is data-flow-guided fuzzing viable with minimal run-time overheads? (Section 5.5.2)
- RQ 2** Does data-flow-guided fuzzing find more or different bugs? (Section 5.5.3)
- RQ 3** Does data-flow-guided fuzzing expand more coverage? (Section 5.5.4)
- RQ 4** Can we predict *a priori* the targets most amenable to data-flow-guided fuzzing? (Section 5.5.5)

5.5.1 Methodology

Fuzzer Selection

Our evaluation compares the performance of fuzzers using: (i) pure control-flow coverage; (ii) pure data-flow coverage; and (iii) exact and approximate dynamic taint analysis (DTA), combining control-flow coverage with data-flow tracking.

We select AFL++ (commit 3e2986d) as the pure control-flow-guided fuzzer because it is the current state-of-the-art coverage-guided greybox fuzzer. We configure AFL++ with: (i) link-time optimization (LTO) instrumentation, eliminating hash collisions; and (ii) with and without “CmpLog” instrumentation. CmpLog—inspired by REDQUEEN’s input-to-state correspondence [9]—approximates DTA by capturing comparison operands. Similarly, we select Angora as an alternative control-flow-guided fuzzer (using context-sensitive edge coverage) that also incorporates exact DTA. Finally, we select DDFuzz as an alternative data-flow-guided fuzzer.²

We configure DATAFLOW with: (i) two *def* site sensitivities: arrays only (“A”) and arrays + structs (“A+S”); and (ii) three *use* site sensitivities: simple access (“A”), accessed offset (“O”), and accessed value (“V”). We use the notation “X/Y” to refer to

²DDFuzz was published concurrently with DATAFLOW.

TABLE 5.1: Evaluated fuzzer configurations. Angora and DDFuzz use their default map sizes. AFL++’s LTO instrumentation does not require a fixed-size map.

Name	Map size (KiB)	Description
A _{LTO}	–	AFL++ with LTO instrumentation
A _{CL}	–	AFL++ with LTO and CmpLog instrumentation
An	1,024	Angora
DD	64	DDFuzz
D _{A/A}	1,024	DATAFLOW with array <i>defs</i> with accessed <i>uses</i>
D _{A/A+O}	1,024	DATAFLOW with array <i>defs</i> with accessed offset <i>uses</i>
D _{A/A+V}	1,024	DATAFLOW with array <i>defs</i> with accessed value <i>uses</i>
D _{A+S/A}	1,024	DATAFLOW with array & struct <i>defs</i> with accessed <i>uses</i>
D _{A+S/A+O}	1,024	DATAFLOW with array & struct <i>defs</i> with accessed offset <i>uses</i>
D _{A+S/A+V}	1,024	DATAFLOW with array & struct <i>defs</i> with accessed value <i>uses</i>

the composition of X *def* and Y *use* site sensitivities; e.g., “A/A” refers to array *def* and access *use* sites; “A+S/O” refers to arrays + structs *def* and accessed offset *use* sites. The evaluated fuzzers are summarized in Table 5.1.

Target Selection

We evaluate the ten fuzzers in Table 5.1 on the following targets. We fuzz 20 target programs in total.

SPEC CPU2006. The SPEC CPU benchmark suite [89] is an industry-standardized, CPU-intensive benchmark suite for stress-testing a system’s processor, memory subsystem, and compiler. We use SPEC CPU2006 to answer RQ 1.

Magma. Unlike other fuzzing benchmarks (e.g., UNIFUZZ [124]), Magma [84] contains ground-truth bug knowledge. We exclude the *php* target because it failed to build with AFL++’s CmpLog instrumentation (failing with a segmentation fault). We use 15 Magma targets to answer RQ 2.

DDFuzz target dataset. Mantovani, Fioraldi, and Balzarotti [141] select five targets—*bison*, *pcr2*, *mir*, *qbe*, and *faust*—they believe to contain a large number of data dependencies, and hence are amenable to data-flow-guided fuzzing. We use newer versions of these targets (because some did not compile on Ubuntu 20.04), shown in Table 5.2. We use these targets to answer RQ 3.

Experimental Setup

We conduct all experiments on an Ubuntu 20.04 AWS EC2 instance with a 48-core Intel[®] Xeon[®] Platinum 8275CL 3.0GHz CPU and 92GiB of RAM. Each fuzz run

TABLE 5.2: DDFuzz target dataset.

Target	Driver	Command line	Commit
<i>bison</i>	bison	@@ -o /dev/null	5555f4d
<i>pcre2</i>	pcre2test	@@ /dev/null	db53e40
<i>mir</i>	c2m	@@	852b1f2
<i>qbe</i>	qbe	@@	c8cd282
<i>faust</i>	faust	@@	13def69

was conducted for 24 h and repeated five times (ensuring statistically sound results). All targets were bootstrapped with their provided seeds.³ Finally, we (a) manually located and specified memory allocation functions for DATAFLOW to tag, and (b) used Angora’s default behavior to discard taint when calling an external library.

5.5.2 Run-time Overheads (RQ 1)

Conventional wisdom assumes data-flow-based coverage metrics are too heavy-weight, adversely affecting a fuzzer’s performance by reducing its iteration rate. We investigate the extent to which this assumption is true by isolating the effects of instrumentation overhead *outside* of a fuzzing environment. Per Section 5.5.1, we measure performance overheads on SPEC CPU2006.

Table 5.3 shows the overhead of all ten evaluated fuzzers on all 19 C and C++ targets in the SPEC CPU2006 v1.0 benchmark suite. We compare these measurements against a baseline without instrumentation (clang v12), calculating the geometric mean (“geomean”) and 95 % bootstrap confidence interval (CI) over three repeated iterations. The following results are omitted because they failed to build or run: AFL++ (LTO) 445.gobmk triggered a run-time assertion; DATAFLOW (all configurations) 429.mcf crashed with a run-time segmentation fault; and Angora 447.dealll, 471.omentpp, 473.astar, and 483.xalancbmk failed to link with DFSAN’s runtime library.

Angora has a geomean overhead of $32.79\times$. This is particularly notable because previous work has found DFSAN—the framework upon which Angora’s taint tracking mode is built—to be one of the more performant DTA frameworks [192]. However, while this overhead is significantly higher than AFL++ (LTO) and AFL++ (CmpLog)—which have geomean overheads of $1.19\times$ and $2.80\times$, respectively—it is important to recall Angora amortizes this cost over the lifetime of a fuzzing campaign by only tracking taint once on a given input over many mutations.

Of the six DATAFLOW configurations, A/A has the lowest overhead ($10.69\times$), while A + S/V has the highest ($15.01\times$). This is unsurprising, given the rolling hash

³We contacted Mantovani, Fioraldi, and Balzarotti [141] to obtain their initial seed sets.

TABLE 5.3: SPEC CPU2006 overheads. Computed as the geomean (over three repeated iterations) relative to an uninstrumented benchmark (compiled with `clang v12`). The 95% bootstrap CI is reported for the geomean across all targets (for a given fuzzer). The bootstrap CI is zero for individual targets and hence is omitted. A \times indicates the target failed to compile or run.

Target	Fuzzer (\times)									
	A _{LTO}	A _{CL}	An	DD	DF _{A/A}	DF _{A/O}	DF _{A/V}	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}
<i>400.perlbench</i>	1.27	3.86	141.85	21.81	12.04	12.75	16.34	12.51	13.21	16.79
<i>401.bzip2</i>	1.26	2.17	25.83	2.54	7.75	8.53	11.12	7.69	8.49	11.09
<i>403.gcc</i>	1.30	3.40	21.19	3.45	19.53	21.22	26.07	19.58	21.18	26.20
<i>429.mcf</i>	1.12	2.46	12.08	1.52	\times	\times	\times	\times	\times	\times
<i>445.gobmk</i>	\times	2.48	23.41	5.26	6.99	7.51	9.48	6.92	7.44	9.73
<i>456.hammer</i>	1.12	3.08	60.41	1.56	13.47	15.07	21.61	13.60	14.95	21.62
<i>458.sjeng</i>	1.21	4.36	29.69	4.44	7.57	8.13	10.05	7.54	8.01	10.32
<i>462.libquantum</i>	1.20	2.40	27.09	1.61	3.81	4.04	6.97	3.70	4.14	6.97
<i>464.h264ref</i>	1.19	1.82	41.01	1.88	100.63	109.40	134.10	100.96	109.68	134.58
<i>471.omnetpp</i>	1.06	2.02	\times	2.05	6.58	6.34	6.82	6.15	6.34	7.56
<i>473.astar</i>	1.13	2.19	\times	1.59	5.53	5.83	6.80	5.60	5.96	7.28
<i>483.xalancbmk</i>	1.29	5.04	\times	3.48	11.44	12.25	15.67	11.66	12.43	15.93
Geomean	1.19 ± 0.00	2.80 ± 0.01	32.79 ± 0.34	2.91 ± 0.03	10.69 ± 0.13	11.41 ± 0.18	14.64 ± 0.22	10.65 ± 0.16	11.47 ± 0.21	15.01 ± 0.21

approach used for the “access with value” *use* sensitivity (Section 5.4.3). Performance improvements are possible by specializing the hash function based on the *type* of value accessed (e.g., hashing a `uint64_t` or `float` value directly, rather than dividing it into single-byte chunks). Increasing the *def* site sensitivity to include structs added minimal overhead. However, this is target specific: the median number of tracked arrays (across the 12 SPEC CPU2006 targets) is 51, compared to 33 structs. This result may not generalize across targets where structs outnumber arrays.

Our results reflect those presented by Liu and Criswell [128] (e.g., *464.h264ref* has the highest run-time overhead in both the original and our work). However, there is a significant increase in our run-time overheads compared to the original PAMD implementation [128]. To validate our PAMD (re)implementation we evaluated a version of DATAFLOW that *only* performed metadata lookup in the baggy bounds table (i.e., it did not construct *def-use* chains nor update the fuzzer’s coverage map). This version of DATAFLOW has a geomean overhead of $3.97\times$. *Def-use* chain construction is a simple xor operation (Section 5.4.3), so we attribute this dramatic increase in run-time overhead to the interaction of the baggy bounds table and coverage map. In particular, cache effects associated with reading from/writing to these two tables.

Finding

Despite building on PAMD—an efficient metadata encoding scheme—DATAFLOW remains impaired by high run-time overheads. Maximizing fuzzer iteration rates (e.g., by lowering run-time instrumentation costs) is crucial to maximizing fuzzing outcomes.

5.5.3 Bug Finding (RQ 2)

Following Sections 2.6.2 and 4.4.3, we again use survival analysis to summarize our bug-finding results. When computing the restricted mean survival time (RMST) we use $N = 5$ repeated trials and an upper bound $T = 24$ h. The log-rank test is again used to compute the statistical significance of our bug-finding results; two fuzzers have statistically-equivalent bug survival times if the log-rank test's p -value > 0.05 .

We present our bug-finding results in Table 5.4. Based on raw bug counts, AFL++ was the best performing fuzzer, triggering 60 bugs. The two data-flow-driven fuzzers followed this; DDFuzz (44 bugs) and DATAFLOW (41 bugs). Angora was the worst-performing fuzzer, triggering only 24 bugs.

DATAFLOW with “simple access” *use* sensitivity ($DF_{A/A}$ and $DF_{A+S/A}$) was the best-performing version of DATAFLOW (39 bugs). This was followed by $DF_{A+S/O}$ (31 bugs). DATAFLOW with “accessed value” *use* sensitivity was the worst performer. This suggests incorporating variable values at *use* sites is not worth the increased run-time cost; simply tracking the existence of *def-use* chains is “good enough” (for discovering bugs).

AFL++ remains the best-performing fuzzer when accounting for RMSTs (i.e., it triggers bugs fastest), outperforming the data-flow-guided fuzzers for the majority of bugs triggered (60%). However, this result is reversed (i.e., the data-flow-guided fuzzers outperform AFL++) for 14% of the triggered bugs. Notably, DATAFLOW was the only fuzzer to trigger LUA003 (not previously triggered by any fuzzer in any prior Magma evaluation), while DATAFLOW and DDFuzz triggered XML001 (`xml1.int`) and LUA004 orders-of-magnitude faster than AFL++. DDFuzz was the only fuzzer to trigger PDF008. However, this bug was only triggered once (over five trials) and towards the end of the trial (after 20 h). This suggests that the bug is difficult to find and DDFuzz may have just “got lucky”. Finally, AFL++ either failed to trigger or was orders-of-magnitude slower at triggering SSL009 (`x509`) and PDF003 (`pdf.images`). *Do these bugs share properties that make them amenable to discovery via data-flow-guided fuzzing?* To answer this question, we examine the two lua bugs in greater depth.

TABLE 5.4: Magma bugs, presented as the RMST (in hours) with 95 % bootstrap CI. Bugs never found by a particular fuzzer have an RMST of \top (to distinguish bugs with a 24 h RMST). We only report the RMST for bugs triggered; bugs not triggered by any fuzzer are omitted. The best performing fuzzer (fuzzers if the bug survival times are statistically equivalent per the log-rank test) for each bug is highlighted in green (smaller is better).

Target	Driver	Bug	Fuzzer									
			ALTO	ACL	An	DD	DF _{A/A}	DF _{A/O}	DF _{A/V}	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}
libpng	read_fuzzer	PNG003	0.01 ± 0.01	0.01 ± 0.01	0.01 ± 0.01	0.05 ± 0.01	0.03 ± 0.02	0.06 ± 0.01	0.24 ± 0.15	0.04 ± 0.02	0.03 ± 0.03	0.99 ± 0.44
		PNG006	\top	0.02 ± 0.02	0.07 ± 0.03	\top	\top	\top	\top	\top	\top	\top
		PNG007	7.47 ± 6.93	17.54 ± 4.15	\top	19.49 ± 7.37	23.38 ± 1.39	\top	\top	19.49 ± 10.21	19.94 ± 9.19	\top
libsndfile	sndfile_fuzzer	SND001	0.43 ± 0.27	0.72 ± 0.29	-	1.75 ± 0.55	0.78 ± 0.20	19.97 ± 6.74	\top	0.77 ± 0.32	20.44 ± 8.07	\top
		SND005	0.55 ± 0.21	0.62 ± 0.43	-	2.05 ± 0.79	5.63 ± 2.25	23.25 ± 1.69	\top	5.63 ± 1.55	15.32 ± 7.81	20.49 ± 7.96
		SND006	0.36 ± 0.23	0.26 ± 0.16	-	1.06 ± 0.21	6.54 ± 5.97	\top	19.52 ± 10.15	3.05 ± 2.43	\top	18.96 ± 8.45
		SND007	0.64 ± 0.27	0.27 ± 0.16	-	2.53 ± 0.54	1.39 ± 0.73	\top	\top	3.33 ± 2.21	21.12 ± 6.52	22.57 ± 3.24
		SND017	0.43 ± 0.28	0.06 ± 0.02	-	0.00 ± 0.01	0.01 ± 0.02	0.74 ± 1.21	0.04 ± 0.04	0.02 ± 0.02	0.01 ± 0.01	14.42 ± 11.50
		SND020	0.71 ± 0.30	0.49 ± 0.14	-	0.60 ± 0.11	0.30 ± 0.20	4.76 ± 2.72	1.71 ± 0.57	0.61 ± 0.36	5.53 ± 7.39	15.15 ± 10.63
		SND024	0.31 ± 0.24	0.26 ± 0.16	-	1.03 ± 0.21	0.50 ± 0.30	21.61 ± 3.20	19.38 ± 10.45	0.60 ± 0.26	21.03 ± 6.73	15.07 ± 10.72
		libtiff	read_rgba_fuzzer	TIF002	18.44 ± 6.92	\top	\top	\top	\top	\top	\top	\top
TIF007	0.01 ± 0.01			0.02 ± 0.01	0.81 ± 0.40	0.27 ± 0.11	0.17 ± 0.07	0.95 ± 0.63	19.62 ± 9.92	0.31 ± 0.18	2.20 ± 1.85	14.93 ± 10.89
TIF008	19.97 ± 5.16			\top	\top	\top	\top	\top	\top	\top	\top	\top
TIF012	0.16 ± 0.05			0.99 ± 0.47	6.09 ± 7.19	10.32 ± 6.45	3.71 ± 1.45	14.83 ± 11.01	\top	4.56 ± 3.05	19.50 ± 10.17	\top
tiffcp	TIF014		0.60 ± 0.20	1.26 ± 0.37	\top	15.46 ± 10.33	14.15 ± 7.31	20.12 ± 8.79	\top	20.07 ± 8.90	19.45 ± 10.29	\top
	TIF005		\top	\top	14.91 ± 10.91	\top	\top	\top	\top	\top	\top	\top
	TIF006		7.52 ± 4.11	15.40 ± 8.28	10.14 ± 6.14	\top	22.86 ± 2.59	\top	\top	18.12 ± 7.14	\top	\top
	TIF007		0.03 ± 0.02	0.02 ± 0.02	0.73 ± 0.66	0.29 ± 0.09	0.26 ± 0.08	6.26 ± 7.22	1.68 ± 0.45	0.38 ± 0.20	0.66 ± 0.16	10.57 ± 9.61
	TIF008		\top	22.17 ± 4.13	\top	\top	\top	\top	\top	\top	\top	\top
	TIF009		10.33 ± 6.86	13.30 ± 7.97	\top	\top	17.70 ± 7.79	\top	\top	\top	\top	\top
	TIF012		0.26 ± 0.13	0.71 ± 0.27	11.57 ± 9.09	11.14 ± 7.90	4.99 ± 1.97	20.42 ± 8.11	\top	7.02 ± 6.80	\top	\top
	TIF014		0.50 ± 0.21	1.60 ± 0.70	14.23 ± 7.56	8.79 ± 6.65	12.87 ± 8.05	20.31 ± 8.35	\top	13.78 ± 7.91	14.22 ± 6.59	\top

TABLE 5.4: Magma bugs (continued).

Target	Driver	Bug	Fuzzer										
			ALTO	ACL	An	DD	DF _{A/A}	DF _{A/O}	DF _{A/V}	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	
libxml2	read_memory_fuzzer	XML001	T	T	10.16 ± 3.03	13.11 ± 6.48	23.23 ± 1.74	T	T	19.99 ± 4.84	T	T	
		XML002	19.51 ± 10.15	T	T	T	T	T	T	T	T	T	
		XML003	4.00 ± 3.45	0.78 ± 0.61	0.01 ± 0.01	0.03 ± 0.03	0.04 ± 0.05	0.05 ± 0.00	0.05 ± 0.04	0.05 ± 0.05	0.04 ± 0.05	0.05 ± 0.06	
		XML009	1.07 ± 0.35	0.95 ± 0.29	T	14.59 ± 6.94	3.38 ± 3.11	15.31 ± 10.48	T	6.70 ± 6.33	15.72 ± 10.08	T	
		XML017	0.01 ± 0.01	0.01 ± 0.01	0.01 ± 0.01	0.01 ± 0.01	0.07 ± 0.00	0.08 ± 0.00	0.07 ± 0.00	0.07 ± 0.00	0.07 ± 0.00	0.07 ± 0.00	
	xmllint	XML001	22.88 ± 2.53	T	T	0.01 ± 0.01	0.06 ± 0.00	0.08 ± 0.00	0.07 ± 0.00	0.06 ± 0.00	0.07 ± 0.00	0.07 ± 0.00	
		XML002	22.83 ± 2.64	T	T	T	T	T	T	T	T	T	
		XML009	0.59 ± 0.24	1.14 ± 0.61	4.29 ± 0.12	7.16 ± 3.70	1.57 ± 0.74	18.53 ± 8.70	T	1.18 ± 0.65	15.65 ± 10.03	T	
		XML012	20.61 ± 4.99	22.30 ± 3.84	T	T	T	T	T	T	T	T	
		XML017	0.02 ± 0.02	0.02 ± 0.02	0.01 ± 0.01	0.02 ± 0.02	0.05 ± 0.04	0.06 ± 0.04	0.05 ± 0.00	0.05 ± 0.05	0.05 ± 0.00	0.05 ± 0.00	
	lua	lua	LUA003	T	T	T	T	T	T	19.81 ± 9.47	T	T	
			LUA004	6.06 ± 3.88	10.38 ± 5.67	T	0.22 ± 0.37	0.02 ± 0.02	0.01 ± 0.02	0.02 ± 0.01	0.01 ± 0.02	0.21 ± 0.37	0.01 ± 0.02
	openssl	asn1	SSL001	3.53 ± 3.38	8.34 ± 4.86	T	21.83 ± 3.63	T	T	T	T	T	T
			SSL003	0.05 ± 0.07	0.03 ± 0.03	0.00 ± 0.01	0.01 ± 0.01	0.01 ± 0.02	0.01 ± 0.02	0.02 ± 0.03	0.01 ± 0.02	0.01 ± 0.01	0.01 ± 0.02
		client	SSL002	0.06 ± 0.00	0.06 ± 0.00	0.01 ± 0.01	0.05 ± 0.05	0.02 ± 0.02	0.03 ± 0.03	0.03 ± 0.03	0.02 ± 0.03	4.82 ± 10.85	0.03 ± 0.02
			server	SSL002	0.08 ± 0.00	0.09 ± 0.01	0.18 ± 0.02	5.00 ± 7.60	0.52 ± 0.00	0.69 ± 0.16	0.69 ± 0.27	5.22 ± 7.51	0.39 ± 0.04
		SSL020		20.78 ± 6.00	T	6.13 ± 0.14	21.90 ± 4.76	T	T	T	T	T	T
x509		SSL009	T	T	0.01 ± 0.02	0.11 ± 0.00	0.02 ± 0.02	0.02 ± 0.02	0.02 ± 0.02	0.02 ± 0.02	0.01 ± 0.02	0.02 ± 0.02	
poppler	pdf_fuzzer	PDF010	1.08 ± 0.54	1.47 ± 0.96	T	2.23 ± 1.33	11.01 ± 4.15	T	16.14 ± 9.47	7.91 ± 6.84	16.90 ± 5.17	22.22 ± 4.03	
		PDF016	0.03 ± 0.01	0.02 ± 0.02	0.23 ± 0.00	0.16 ± 0.03	0.37 ± 0.09	0.31 ± 0.01	0.27 ± 0.01	0.31 ± 0.02	0.25 ± 0.01	0.49 ± 0.03	
		PDF018	17.60 ± 7.22	15.86 ± 4.14	T	17.35 ± 8.27	20.09 ± 5.16	T	T	T	T	T	
		PDF019	23.14 ± 1.94	T	T	T	T	T	T	T	T	T	
		PDF021	23.52 ± 1.09	T	T	22.30 ± 3.86	T	T	T	T	T	T	

```

1 #define l_checkmodep(m) \
2 ((m[0] == 'r' || m[0] == 'w') && m[1] == '\0')

```

FIGURE 5.7: LUA003 missing popen check.

LUA003. This bug is caused by a missing check of the “mode” argument to popen. The check is shown in Fig. 5.7. While the check is quickly reached by DDFuzz (after ~4 h) and all six DATAFLOW variations (on average, after ~60 s), the exact trigger conditions were only met once by DF_{A+S/A}. Upon examining the compiled binary, we found the second check (`m[1] == '\0'`) was optimized to a *branchless* operation (i.e., it did not contain conditional control flow). This effectively makes the program state where `m[1] != '\0'` invisible to a control-flow-guided fuzzer (in particular, there is no explicit edge for AFL++ to instrument). This state is explicitly visible to DATAFLOW, which reaches it after ~19 h of fuzzing.

LUA004. This is a logic bug, caused by a missing update to the interpreter’s “old” program counter (occurring under particular conditions when tracing the execution of a Lua function). Again, there is no explicit “state” in the target’s control-flow graph (CFG) for the fuzzer to reach. Rather, the bug is triggered when the `oldpc` field in the `lua_State` struct is not updated. This only happens under particular conditions, again depending on specific data values.

Finding

The control-flow-guided fuzzers (AFL++ and Angora) outperform the data-flow-guided fuzzers (DDFuzz and DATAFLOW) on 60 % of the triggered Magma bugs. However, the data-flow-guided fuzzers significantly outperform the control-flow-guided fuzzers (by orders-of-magnitude) on 11 % of the triggered bugs. These results suggest that fuzzers guided by control flow and data flow should be combined to maximize bug-finding potential.

5.5.4 Coverage Expansion (RQ 3)

Control-flow coverage is typically quantified by reasoning over the target’s CFG (e.g., basic blocks, edges, lines of code). For example, FUZZBENCH replays the fuzzer’s queue through an independent and precise (i.e., collision-free) coverage metric; specifically, Clang’s source-based coverage [144, 209]. However, the equivalent process for quantifying *data-flow* coverage does not exist.

We quantify coverage expansion using both control-flow and data-flow metrics, using (a) static analyses to approximate an upper bound, and (b) dynamic analyses to quantify coverage expansion against this upper bound. The usual limitations of static analysis (e.g., undecidability) mean this upper bound may be larger than the set of executable coverage elements (e.g., a code region may not be reachable from the target’s driver, or a pointer’s points-to set may be over-approximated). We accept this imprecision for both metrics (because it is a comparative metric that does not rely on an absolute value). Per Section 2.6.4, we use the Mann-Whitney U -test [139] to statistically compare dynamic coverage across fuzzers: two fuzzers cover the same number of coverage elements if the Mann-Whitney U -test’s p -value > 0.05 .

Control-flow coverage. We use Clang’s existing source-based coverage metric [209]. Specifically, we use *region coverage* (as used by FUZZBENCH), Clang’s version of statement coverage. Like classic statement coverage, region coverage is more granular than function and line coverage [94]. Region information is embedded into the target during compilation and can be statically extracted using existing LLVM tooling (to obtain the upper bound).

Data-flow coverage. We build a static analysis on SVF [204] (commit d6fe474), a state-of-the-art value flow and pointer analysis framework. This static analysis computes the set of *def-use* chains in a target (for the set of tracked variables, as determined by the chosen *def* site sensitivity). This analysis leverages a flow- and context-insensitive interprocedural pointer analysis based on the Andersen algorithm [4].⁴ For the dynamic analysis, we modify the PAMD metadata stored at each *def* site (Section 5.4.2) to store a tuple of $\langle \text{variable name}, \text{location} \rangle$, where *location* is another tuple $\langle \text{source filename}, \text{function name}, \text{line}, \text{column} \rangle$. Both tuples are constructed by extracting source-level information from the target’s debug information. A *use* site (Section 5.4.2) is similarly labeled with a *location* tuple. Unlike the 16-bit tags used by DATAFLOW, this approach does not result in hash collisions and is precise (albeit with a higher run-time cost). Importantly, neither the static nor dynamic analysis takes into account *def-use* chain *values*. We also exclude dynamic memory allocations from these analyses (to simplify run-time *def-use* tracking when faced with custom memory allocators, per Section 5.4.4).

Table 5.5 and Figs. 5.8 and 5.9 summarize our coverage expansion results. Two targets, bison and faust, failed to build with AFL++’s CmpLog (again, due to a segmentation fault) and are excluded from our results.

⁴We experimented with SVF’s flow-sensitive interprocedural analysis but found the run-time overheads prohibitively large.

TABLE 5.5: Coverage expansion (control and data flow) on DDFuzz targets, reported as the mean over five repeated trials with 95 % bootstrap CI. The “static” results give an approximate upper bound, while the “dynamic” results give the percentage of coverage elements covered at run time. The best performing fuzzer (fuzzers if the coverages are statistically equivalent per the Mann-Whitney U -test) for each target is highlighted in green (larger is better). A \times indicates the target failed to compile or run.

(A) Control-flow coverage. Quantified in terms of Clang’s code region coverage.

Target	Static (#)	Dynamic (%)									
		A _{LTO}	A _{CL}	An	DD	DF _{A/A}	DF _{A/O}	DF _{A/V}	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}
bison	35,476	33.47 ± 0.08	\times	30.84 ± 0.31	29.95 ± 3.21	32.47 ± 0.39	30.83 ± 0.11	27.95 ± 0.56	31.31 ± 0.12	29.59 ± 0.23	28.28 ± 0.27
pcre2test	36,973	56.35 ± 1.78	62.22 ± 2.25	23.01 ± 1.68	37.56 ± 0.30	36.21 ± 0.00	22.63 ± 1.31	21.51 ± 1.25	36.13 ± 0.32	21.13 ± 0.66	23.01 ± 1.98
c2m	43,765	45.35 ± 0.07	45.86 ± 0.08	43.81 ± 0.22	45.45 ± 0.07	44.22 ± 0.09	41.99 ± 0.15	41.95 ± 0.22	44.18 ± 0.08	42.02 ± 0.21	41.78 ± 0.17
qbe	5,400	76.15 ± 0.19	76.34 ± 0.21	73.94 ± 0.03	76.03 ± 0.15	74.17 ± 0.09	73.03 ± 0.03	73.84 ± 0.12	74.54 ± 0.11	73.20 ± 0.14	74.11 ± 0.07
faust	26,872	32.11 ± 0.32	\times	29.02 ± 0.09	32.07 ± 0.20	31.02 ± 0.08	30.17 ± 0.04	30.57 ± 0.04	31.13 ± 0.17	30.16 ± 0.03	30.31 ± 0.12

(B) Data-flow coverage. Quantified in terms of interprocedural *def-use* chains (scaled by $\times 10^{-3}$, due to the small percentage of *def-use* chains covered across all targets). The c2m target is excluded because it unexpectedly crashed with our precise data-flow tracking instrumentation.

Target	Static (#)	Dynamic ($\times 10^{-3}$ %)									
		A _{LTO}	A _{CL}	An	DD	DF _{A/A}	DF _{A/O}	DF _{A/V}	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}
bison	9,923,196	11.60 ± 0.10	\times	9.94 ± 0.17	10.40 ± 0.17	11.27 ± 0.26	10.71 ± 0.10	9.57 ± 0.33	11.26 ± 0.16	10.60 ± 0.03	9.84 ± 0.28
pcre2test	1,401,778	181.03 ± 14.78	205.78 ± 16.14	78.24 ± 7.23	125.53 ± 1.53	122.72 ± 1.30	76.12 ± 4.87	69.51 ± 5.71	124.83 ± 1.65	69.73 ± 2.86	73.82 ± 5.99
c2m	25,462,192	-	-	-	-	-	-	-	-	-	-
qbe	450,407	381.97 ± 3.06	378.68 ± 4.46	178.28 ± 0.47	376.64 ± 3.62	363.01 ± 1.71	352.61 ± 0.36	356.21 ± 1.38	366.20 ± 0.40	353.24 ± 0.62	355.77 ± 4.22
faust	159,515,187	5.46 ± 0.06	\times	4.88 ± 0.07	5.55 ± 0.06	5.19 ± 0.01	4.81 ± 0.03	5.01 ± 0.05	5.34 ± 0.03	4.78 ± 0.02	4.89 ± 0.07

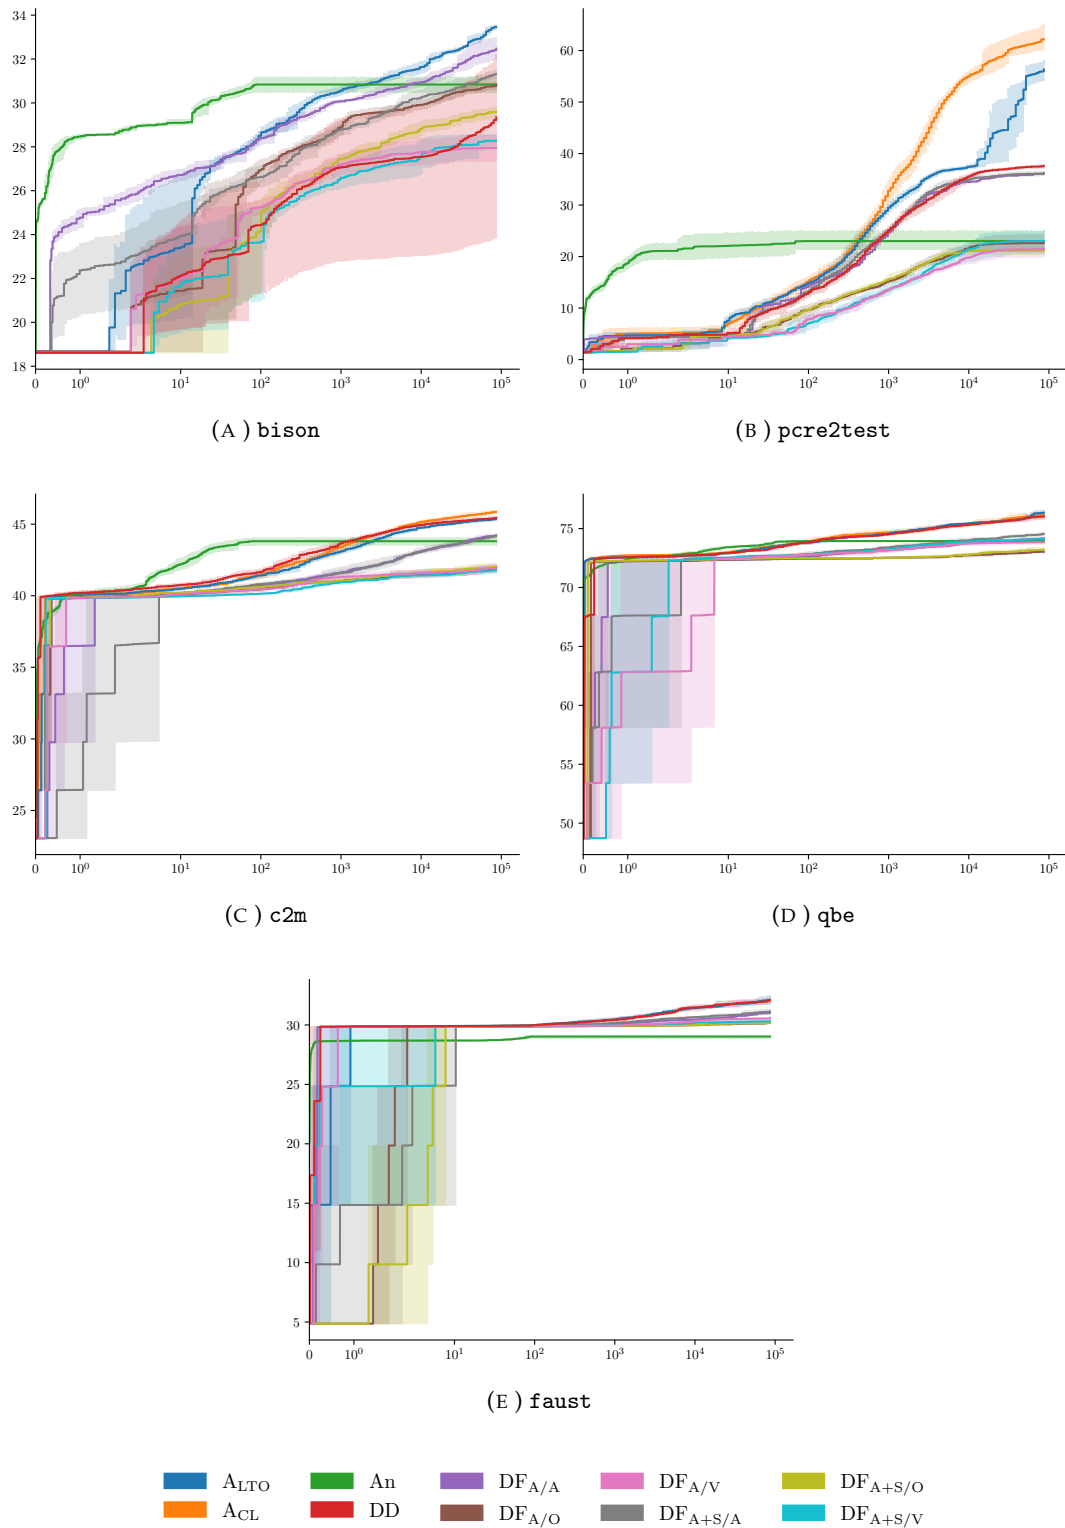


FIGURE 5.8: Control-flow coverage expansion over time. The x -axis is time in seconds (log scale), and the y -axis is the percentage of code regions expanded (against the static upper bound in Table 5.5a). The mean coverage (over five repeated trials) and 95% bootstrap CI is shown.

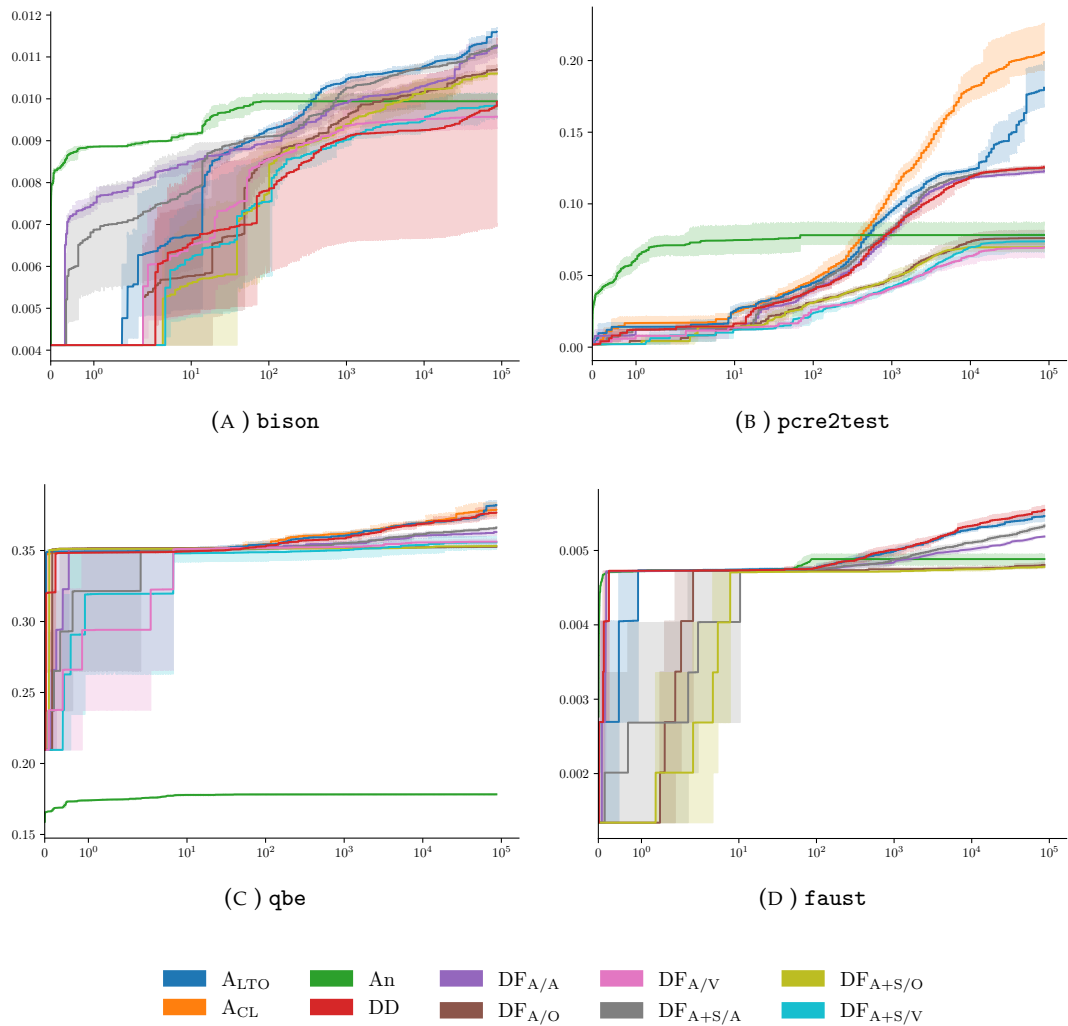


FIGURE 5.9: Data-flow coverage expansion over time. The x -axis is time in seconds (log scale), and the y -axis is the percentage of *def-use* chains expanded (against the static upper bound in Table 5.5b). The mean coverage (over five repeated trials) and 95% bootstrap CI is shown.

AFL++ is again the best-performing fuzzer, achieving the highest control-flow (i.e., code region) coverage. CmpLog improves AFL++’s already-strong coverage expansion capabilities. These results are unsurprising, given that control-flow coverage (specifically, edge coverage) guides AFL++. Similarly, Angora again performs poorly, outperformed by both DDFuzz and DATAFLOW in maximizing both control- and data-flow coverage. Curiously, however, AFL++ also achieves the highest *data-flow* (i.e., *def-use* chain) coverage. This is despite DATAFLOW’s data-flow guidance. We attribute this (surprising) result to the differences in *fuzzer iteration rates* (i.e., the number of inputs executed by the fuzzer per unit of time).

Accounting for Iteration Rates

AFL++ (LTO) achieves a mean iteration rate of 1,172 execs/s (median 347 execs/s). In contrast, DDFuzz, Angora, and DATAFLOW achieve mean iteration rates of 974, 616 and 270 execs/s, respectively (median 442, 249 and 144 execs/s). This dramatic decrease in iteration rates reflects our overhead results in Section 5.5.2.

To account for differences in iteration rates, rather than comparing coverage at the *end* of each fuzz run (i.e., after 24 h of fuzzing), we compare coverage at a given *execution* (“exec”). Specifically, we compare coverage at the last exec of the slowest fuzzer (i.e., with the lowest iteration rate). Intuitively, this places a “ceiling” on the coverage achieved by faster fuzzers (i.e., those able to execute more inputs within a single 24 h fuzz run). For example, $DF_{A+S/V}$ is the slowest fuzzer on *bison* (85 execs/s). Thus, we compare the coverage achieved at the last execution of $DF_{A+S/V}$ (exec = 7,358,231), implicitly ignoring any additional coverage expanded after this exec. Unfortunately, Angora does not provide the necessary information to map coverage to a particular exec, so we exclude it from our analysis (despite it being the slowest fuzzer on two targets: *bison* and *faust*).

We present coverage “normalized” against iteration rates in Table 5.6. DATAFLOW is now more competitive (against AFL++) on expanding data-flow coverage. It achieves the highest *def-use* chain coverage on *bison* and *faust*, and is only ~4% behind the number of *def-use* chains expanded by AFL++ on *qbe*. Again, increasing DATAFLOW’s *use* sensitivity to include variable values fails to improve fuzzing outcomes. These results reinforce our belief that fuzzer iteration rates have a significant impact on fuzzing outcomes.

TABLE 5.6: Coverage expansion—control and data flow (“CF” and “DF”, respectively)—on DDFuzz targets. In contrast to Table 5.5, which reports the (mean) coverage achieved at the *end of a 24 h fuzz run*, here we report the (mean) coverage at the *last exec of the slowest fuzzer* (“Exec” occurring at the given “Time”). Dynamic coverage is quantified in terms of the static analysis results in Table 5.5 (the DF results are again scaled by $\times 10^{-3}$). The best performing fuzzer (fuzzers if the coverage is statistically equivalent per the Mann-Whitney U -test) for each target is highlighted in **green** (larger is better).

Target	Exec (#)	Time (hr)	Metric	Dynamic (% , $\times 10^{-3}$ %)								
				A _{LTO}	A _{CL}	DD	DF _{A/A}	DF _{A/O}	DF _{A/V}	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}
bison	7,358,231	13.74	CF	32.35 ± 0.09	×	28.39 ± 3.71	32.17 ± 0.47	30.64 ± 0.11	27.89 ± 0.53	30.90 ± 0.20	29.54 ± 0.23	28.11 ± 0.51
			DF	10.95 ± 0.09	×	9.56 ± 2.08	11.06 ± 0.24	10.59 ± 0.06	9.55 ± 0.34	11.03 ± 0.08	10.56 ± 0.08	9.79 ± 0.35
pcre2test	37,992,897	8.22	CF	37.51 ± 0.30	55.79 ± 1.09	36.66 ± 0.16	35.85 ± 0.00	22.56 ± 1.28	21.38 ± 1.15	35.82 ± 0.29	21.12 ± 0.66	23.01 ± 1.98
			DF	124.40 ± 1.69	182.84 ± 11.61	122.42 ± 0.77	121.02 ± 0.68	75.96 ± 4.72	69.08 ± 5.31	123.47 ± 1.68	69.73 ± 2.83	73.82 ± 5.99
c2m	516,654	5.27	CF	43.20 ± 0.13	43.44 ± 0.08	43.41 ± 0.27	42.84 ± 0.11	41.44 ± 0.24	41.67 ± 0.25	42.77 ± 0.22	41.39 ± 0.11	41.61 ± 0.19
			DF	-	-	-	-	-	-	-	-	-
qbe	12,589,430	15.78	CF	75.39 ± 0.07	75.27 ± 0.12	75.28 ± 0.10	74.11 ± 0.12	72.93 ± 0.06	73.82 ± 0.12	74.41 ± 0.10	73.02 ± 0.13	74.05 ± 0.09
			DF	368.69 ± 1.33	368.38 ± 1.73	368.33 ± 1.20	362.03 ± 1.89	352.53 ± 0.38	356.17 ± 1.33	364.83 ± 0.75	352.93 ± 0.67	355.68 ± 4.26
faust	1,033,846	23.97	CF	31.26 ± 0.04	×	31.05 ± 0.08	30.80 ± 0.06	30.12 ± 0.04	30.49 ± 0.06	30.86 ± 0.08	30.15 ± 0.04	30.25 ± 0.12
			DF	5.25 ± 0.02	×	5.22 ± 0.01	5.13 ± 0.01	4.79 ± 0.03	4.99 ± 0.05	5.22 ± 0.00	4.77 ± 0.01	4.87 ± 0.08

Finding

The control-flow-guided fuzzers (specifically, AFL++) achieve the highest control- and data-flow coverage. Data-flow-guided fuzzers require more complex instrumentation (compared to control-flow-guided fuzzers), impairing the fuzzers’ iteration rates.

5.5.5 Characterizing Data-Flow (RQ 4)

The fuzzing community has largely settled on control-flow-based coverage metrics—in particular, edge coverage—to drive a fuzzer’s exploration. While prior successes have largely validated this approach [52, 181, 189, 207, 231], we wish to understand what (if any) program characteristics lend themselves to data-flow-based coverage.

Mantovani, Fioraldi, and Balzarotti [141] propose the *DD ratio*—defined as the ratio between the number of basic blocks instrumented with data-dependency information over the total number of basic blocks in the target—to determine whether data-flow-based coverage—derived from the target’s data dependency graph (DDG)—adds value (e.g., over edge coverage). A higher DD ratio suggests the target is

TABLE 5.7: Characterizing data flow using the data dependency ratio (“DD ratio”) introduced by Mantovani, Fioraldi, and Balzarotti [141]. *Strongly* data-dependent targets (i.e., those with a DD ratio $\geq 10\%$) are highlighted in green.

(A) Magma.		(B) DDFuzz target dataset.	
Target	DD ratio (%)	Target	DD ratio (%)
png_read_fuzzer	13.40	bison	6.59
sndfile_fuzzer	12.14	pcr2test	22.60
tiff_read_rgba_fuzzer	12.01	c2m	21.82
tiffcp	11.37	qbe	12.45
xml_read_memory_fuzzer	12.86	faust	7.29
xmllint	13.03		
lua	12.73		
asn1	9.89		
client	9.99		
server	9.98		
x509	9.98		
pdf_fuzzer	11.62		
pdfimages	9.33		
pdftoppm	11.77		
sqlite3_fuzz	12.80		

more amenable to data-flow-guided fuzzing; a target with a DD ratio above 10% is considered *strongly* data dependent.

Table 5.7 summarizes the DD ratio of our 20 target programs.⁵ Thirteen of these targets (65%) have DD ratios $\geq 10\%$, indicating their suitability for data-flow-guided fuzzing. However, we found little correlation between a target’s DD ratio and fuzzing outcomes (both bug finding and coverage expansion). For example, `png_read_fuzzer` had the highest DD ratio among the Magma targets (13.40%), closely followed by `xmllint` (13.03%). However, AFL++ outperformed the data-flow-guided fuzzers (DDFuzz and DATAFLOW) on both targets (across bug counts and survival times). Similarly, `pcr2test` and `c2m` had the highest DD ratios among the DDFuzz targets (22.60 and 21.82%, respectively). Again, AFL++ outperformed the two data-flow-guided fuzzers (across both control- and data-flow coverage expansion).

Based on these results, we conclude that the DD ratio is not suitable for determining a target’s suitability for data-flow-guided fuzzing. Instead, we propose using *subsumption*.

⁵These values differ from the original DDFuzz evaluation [141] because we use newer versions of the targets (per Section 5.5.1).

Finding

The “DD ratio” is not suitable for determining whether a target is amenable to data-flow-guided fuzzing. Alternative techniques—such as those based on subsumption—may yield more accurate results.

5.5.6 Discussion

Coverage sensitivity. In Section 5.3.1, we introduced a framework for reasoning about and constructing data-flow coverage metrics for greybox fuzzing. This framework allows the user to balance precision with performance. Our results suggest that fuzzing outcomes (i.e., bug finding and coverage expansion) fail to improve as precision increases. Notably, this finding also applies to Angora; Angora’s exact DTA provided little benefit over the approximate DTA used by AFL++’s CmpLog mode. Our results reflect prior findings that demonstrate the importance of maximizing fuzzer iteration rates [9, 68, 92, 173, 225].

Bugs vs. coverage. Böhme, Szekeres, and Metzman [23] found the fuzzer best at maximizing coverage expansion may not be best at finding bugs. Our results reflect this finding; despite AFL++ outperforming DATAFLOW on coverage expansion (Section 5.5.4), DATAFLOW triggered bugs AFL++ failed to find (Section 5.5.3). Ultimately, fuzzers are deployed to find bugs and vulnerabilities; our findings reinforce the need for bug-based fuzzer evaluation [84, 115, 234] (not only a comparison of coverage profiles).

Computing coverage upper bounds with static analysis. In Section 5.5.4 we used static analysis to approximate a coverage upper bound (for both control- and data-flow coverage). In theory, this upper bound is useful for estimating the *residual risk* of ending a fuzz run before maximizing coverage (analogous to the residual risk of missing a bug [20]). In practice, static analysis of “real-world” programs is fraught; dynamically loaded, just in time (JIT), and inline assembly code all impact precision. Even specific command-line arguments influence the reachability of particular code regions. Thus, it is difficult to determine how realistic the upper bounds in Section 5.5.4 are. We return to this issue in Chapter 6.

Control- or data-flow? We hypothesized that data-flow-guided fuzzing offers superior performance on targets where control flow is decoupled from semantics. Our results lead us to reject this hypothesis. In most cases, control-flow-guided fuzzers

outperformed data-flow-guided fuzzers (across both bug-finding and coverage-expansion metrics, and on targets identified as being amenable to data-flow-guided fuzzing). However, we are not prepared to give up on data-flow-guided fuzzing; despite lower run-time costs than DTA, DATAFLOW’s run-time costs remain high, negatively impacting coverage expansion. Despite this impediment, DATAFLOW discovers bugs control-flow-guided fuzzers do not. We believe reducing the run-time costs of data-flow-guided fuzzers will improve fuzzing outcomes.

5.6 Future Work

Our results highlight the opportunities for further work on data-flow-based coverage metrics. We propose ideas for further exploration here.

Reducing overheads. The significant run-time overheads remain the primary impediment to the adoption of data-flow-guided fuzzing (see Section 5.5.2). Liu and Criswell [128] propose using interprocedural optimizations to eliminate unnecessary object (de)allocation in the baggy bounds table, improving performance. Similarly, more sophisticated pointer analyses (e.g., those provided by SVF) could be used to eliminate unnecessary *def/use* site instrumentation (e.g., removing redundant instrumentation when *def-use* chains can be statically identified).

Reducing hash collisions. Per Section 5.4.3, DATAFLOW is prone to hash collisions. It is well known that hash collisions cause fuzzers to miss program behaviors [67]. While AFL++’s LTO mode solves the hash collision problem for edge coverage, we did not investigate a similar technique for *def-use* chain coverage. A hash-collision-free DATAFLOW may lead to improved coverage expansion.

Combining control- and data-flow coverage. Finally, DATAFLOW exclusively uses *def-use* chain coverage to drive exploration. In contrast, other data-flow-guided fuzzers (e.g., INVS COV [61], DDFuzz [141]) combine data flow with control flow. Given our bug-finding results—i.e., those where DATAFLOW significantly outperformed AFL++ (e.g., LUA003, LUA004, SSL009, and PDF003)—combining DATAFLOW with edge coverage may provide a “best of both worlds” solution (echoing the conclusions reached by Salls et al. [183]). This combination of coverage metrics could be realized by combining control- and data-flow coverage in a single coverage map, maintaining separate coverage maps, or by dynamically switching between different instrumented targets.

Coverage metric subsumption. Our results in Section 5.5.5 led us to conclude that the DD ratio was not suitable for determining a target’s suitability for data-flow-guided fuzzing. Prior work on characterizing programs for automated test suite generation is also unsuitable; e.g., the approaches proposed by Neelofar et al. [155] and Oliveira et al. [160] are specific to object-oriented software and focus on control-flow features. Instead, we propose *subsumption*.

We say that coverage metric \mathcal{M}_1 strictly subsumes metric \mathcal{M}_2 if covering all coverage elements in \mathcal{M}_1 also covers all elements in \mathcal{M}_2 . For example, edge coverage strictly subsumes basic block coverage. Relaxing this definition of strict subsumption allows us to quantify the number of coverage elements in \mathcal{M}_2 not subsumed by \mathcal{M}_1 . Intuitively, more elements in \mathcal{M}_2 *not* subsumed by \mathcal{M}_1 implies fuzzing with \mathcal{M}_2 will lead to behaviors not detectable by \mathcal{M}_1 . Static data-flow analysis frameworks such as those proposed by Chaim et al. [29] can be used to perform this subsumption analysis. We leave the investigation of such techniques for future work.

5.7 Chapter Summary

Observing fuzzers that introduce taint tracking along with control flow, we investigate data flow as an alternate coverage metric, making *data-flow coverage* a first-class citizen. Driven by empirical results and the conventional wisdom gathered over years of software-testing research, we hypothesized data-flow-guided fuzzing to offer superior outcomes (over control-flow-guided fuzzing) in targets where control flow is decoupled from semantics.

Our results show that control-flow-guided fuzzing produces better outcomes (bug finding and coverage expansion) in most cases. The high run-time costs associated with data-flow tracking impaired the fuzzer’s ability to explore a target’s behavior efficiently. Despite these costs, our data-flow-guided fuzzer discovered bugs control-flow-guided fuzzers did not. These results suggest that data-flow-guided fuzzers discover *different*, not *more*, bugs. Specifically, bugs existing in program states not explicitly visible in the target’s CFG.

In the next chapter we conclude the fuzzing campaign (④ in Fig. 1.1, Chapter 1). Specifically, we expand on our use of static analyses to quantify control- and data-flow coverage (Section 5.5.4) and investigate their use more broadly in measuring a fuzzer’s state space search.

Chapter 6

Quantifying State Space Search

In Chapter 5 we introduced a pair of static analyses for quantifying coverage expansion using control-flow and data-flow metrics. Here, we expand on this idea and investigate the use of several state-of-the-art static analyses for quantifying the upper-bound of a fuzzer’s state space search.

6.1 Introduction

Measuring how much of a target’s state space has been explored *per fuzzer* is challenging. Fuzzer evaluations typically avoid measuring this by only reporting a count of covered coverage elements (e.g., the number of lines of code executed). However, a raw count is insufficient in many cases. For example, determining the *residual risk* [20] of stopping the fuzzer requires deriving an accurate upper bound of target \mathcal{P} ’s state space, and from this determining how much of \mathcal{P} ’s state space has been explored.

One approach for deriving an upper bound of \mathcal{P} ’s state space is to count the number of edges in \mathcal{P} ’s control-flow graph (CFG). However, this approach is fraught: while the number of *direct* edges (both intra- and inter-procedural) is computable during instrumentation, *indirect* edges (e.g., function calls made by dynamic dispatch, such as the calls to `bar`, `baz`, and `foo` in Fig. 1.2) and *exceptional* control flow (e.g., `try/catch` blocks) make counting edges difficult (and in some cases, impossible). Moreover, the set of *reachable* basic blocks in \mathcal{P} ’s CFG depends on the driver (i.e., the program that interfaces with the codebase—typically a library—under test) and even the command-line arguments used to run the driver [235]! The former (basic block reachability) requires accurate control flow recovery, while the latter (command-line arguments) requires accurate data flow recovery. While perfect control- and data-flow recovery is undecidable, all is not lost: modern static analyzers [79, 80, 119, 185, 204] have demonstrated precise and scalable control- and data-flow analyses on real-world

programs [54, 119, 203]. *Can we use these techniques to make an accurate per-fuzzer evaluation?*

Despite their long history, the application of static program analyses to fuzzing—in particular, to reason about state space exploration—is under explored. This chapter aims to rectify this, and improve understanding of fuzzers’ state space search through program analysis.

Chapter outline. We investigate the use of several static analyses for quantifying control-flow coverage. We first propose our analysis technique (Section 6.2), and then apply our technique to a large-scale fuzzing campaign (Section 6.3).

6.2 Program Analysis for State Space Search

Traditionally, a fuzzer’s state space search is quantified using control-flow-based metrics (e.g., the number of CFG edges or lines of code executed). However, these control-flow-based metrics are often *context insensitive*, despite the fact that different invocations of a function may lead to different program behaviors (as in Figs. 1.2 and 7.1c). Moreover, accurate comparisons of fuzzers’ state space search require the ability to unambiguously identify and correlate points across executions [205, 224]. Notably, context-insensitive control-flow coverage (such as the approach used by FUZZBENCH) does not meet this requirement: using {source file \times line number} to uniquely identify an execution point introduces ambiguities; e.g., in the presence of loops (iteration count), function calls (context sensitivity), and data values.

We overcome this issue by introducing a technique that quantifies control-flow coverage using a *context sensitive interprocedural control-flow graph (ICFG)* analysis. This technique can be used (a) for determining how much of a target’s state space has been explored by a given fuzzer, and (b) to uniquely identify and compare execution points covered by a fuzzer. The former improves the accuracy of *per-fuzzer* evaluation, while the latter improves *cross-fuzzer* evaluation. We develop (a) a static analysis to quantify an upper bound of reachable context-sensitive edges (Section 6.2.1), and (b) a dynamic analysis to trace context-sensitive edges (Section 6.2.2).

6.2.1 Static Analysis

Our static analysis first constructs an intraprocedural CFG for each function f in \mathcal{P} . These CFGs are then combined to create the ICFG by inserting edges from a `call` instruction to the entry block of the callee’s CFG. Indirect calls are resolved via an off-the-shelf static pointer analysis (discussed further in Section 6.2.3).

We augment edges in the ICFG with calling context information. The number of calling contexts for f can be statically computed by [206, 238]:

$$\mathcal{C}(f) = \sum_{i=1\dots n} \mathcal{C}(g_i) \quad (6.1)$$

Where $\{g_1, \dots, g_n\}$ are the caller functions of f . The total number of context-sensitive edges is obtained by propagating $\mathcal{C}(f)$ to each intraprocedural edge in f . Our ICFG analysis is flow-insensitive, so the number of context-sensitive edges may be infinite in the presence of recursion. In this case, we apply the approach proposed by Sumner et al. [206]: (i) a dummy entry node is inserted into \mathcal{P} 's call graph; (ii) edges are added from this dummy node to (a) the original entry node, and (b) any node that is a target of a backward edge (i.e., recursive call); and (iii) backward edges are removed, making the call graph acyclic. From here, Eq. (6.1) is applied. This process implicitly collapses recursive calls and thus mimics the approach taken by fuzzers using context-sensitive coverage (Section 3.2.4).

6.2.2 Dynamic Analysis

Our dynamic analysis inserts instrumentation (at compile time) at the entry of each basic block and function in \mathcal{P} . At run time this instrumentation logs (a) hit counts for intraprocedural branches and function calls/returns, and (b) calling contexts. Recursive function calls are collapsed into a single call so the dynamic calling context is aligned with the static calling context. An offline, *post hoc* analysis reconstructs the context-sensitive ICFG from these traces.

6.2.3 Implementation

Our static and dynamic analyses are built on LLVM v14 [120]. We use the following pointer analysis frameworks and configurations for resolving indirect calls (thus improving the accuracy of our static analysis).

Fuzz Introspector. Fuzz Introspector—part of Google’s OSS-Fuzz service [189]—is a static analysis for computing the set of reachable functions in a given target [32]. However, Fuzz Introspector “*makes few efforts into resolving indirect calls*”, limiting itself to C++ v-table entries and function pointers assigned to variables and passed as function arguments.¹

¹Per <https://github.com/ossf/fuzz-introspector/blob/f01aeee5/frontends/llvm/lib/Transforms/FuzzIntrospector/FuzzIntrospector.cpp#L967-L969>.

SVF. SVF [204] is a value flow and pointer analysis framework. Our SVF analysis uses a flow- and context-insensitive interprocedural pointer analysis based on the Andersen algorithm [4], similar to that used in Section 5.5.4. Also similar to Section 5.5.4, we experimented with SVF’s flow-sensitive interprocedural analysis but found that the analysis failed to complete within 24 h on a majority of targets.

SEADSA. SEADSA is a context-, field-, and array-sensitive unification-based pointer analysis. We use the same configuration as Kuderski, Navas, and Gurfinkel [119]: “bottom-up + top-down context sensitive” analysis with “type awareness”.

6.3 Evaluation

Our evaluation aims to answer the following research question:

RQ 1 How effective are state-of-the-art static analyses at estimating the upper bound of a target’s state space? (Section 6.3.2)

6.3.1 Methodology

Target Selection

We evaluate our static analyses on the FUZZBENCH [144] benchmark. We chose FUZZBENCH over other fuzzer benchmark suites (e.g., Magma [84]) because it is widely used and “uses code coverage as its primary evaluation metric” [144]. We exclude the “bug-based” targets and focus on the 23 “coverage” targets. We also exclude *libjpeg-turbo*, *libpcap*, *libxml2*, *proj4*, and *systemd*, which failed to build/run. The remaining 23 targets are listed in Table 6.1.

Experimental Setup

We run a number of fuzzing campaigns using a range of fuzzers.² Each campaign consists of ten independent 24 h *trials* (consistent with the recommendations by Klees et al. [115]). All experiments were conducted on a Dell PowerEdge server with a 48-core Intel® Xeon® Gold 5118 2.30 GHz CPU, 512 GiB of RAM, and running Ubuntu 18.04. We use FUZZBENCH’s default starting seeds to bootstrap all trials.

²These are the same campaigns as those in Chapter 7. Because we are only interested in the static analysis results here, we defer the details of these campaigns until Section 7.3.

TABLE 6.1: FUZZBENCH target statistics, including: the programming language the target is written in (“Lang”), the number of functions (“F”), the number of basic blocks (“BB”), and the number of static indirect call sites (“Ind.”).

Target	Lang.	F	BB	Ind.
<i>bloaty</i>	C++	5,626	99,259	1,413
<i>curl</i>	C	7,040	69,772	1,677
<i>freetype2</i>	C	2,279	21,007	672
<i>harfbuzz</i>	C++	17,235	42,748	300
<i>jsoncpp</i>	C++	455	6,922	23
<i>lcms</i>	C	1,032	6,528	202
<i>libpng</i>	C	459	4,788	11
<i>libxslt</i>	C	2,227	36,379	2,163
<i>mbedtls</i>	C	1,851	14,311	99
<i>openh264</i>	C++	664	10,608	96
<i>openssl</i>	C	9,381	48,406	1,244
<i>openthread</i>	C++	7,497	25,302	158
<i>re2</i>	C++	4,011	10,841	49
<i>sqlite3</i>	C	2,265	33,105	247
<i>stb</i>	C	220	2,455	11
<i>vorbis</i>	C	339	3,802	33
<i>woff2</i>	C++	3,380	12,274	66
<i>zlib</i>	C	51	688	9

6.3.2 Static Analysis Performance (RQ 1)

As discussed in Chapter 1, static analyses have garnered a reputation for being more harmful than helpful, due to the many false positives/negatives they emit. Moreover, they often fail to scale to “real-world” programs. In Chapter 5 we developed a suite of static analyses to quantify control- and data-flow-coverage achieved by our DATAFLOW fuzzer. However, we did not investigate the accuracy of these analyses. Here, we empirically investigate the accuracy of modern static analyses in the context of fuzzing, focusing on the effectiveness of a static analysis in quantifying the upper bound of a fuzzer’s state space search.

Table 6.1 summarizes the control-flow elements of the 23 FUZZBENCH coverage targets, while Table 6.2 summarizes the control-flow elements *reachable* from the target’s entry point (i.e., elements that may be covered during fuzzing). We use the static analyses discussed in Section 6.2.3 to determine reachability, namely: Fuzz Introspector (commit `f9e9c68b`), SVF (commit `a3fc9cc`), and SEADSA (commit `fda066e`). We also considered the `cclyzer++` pointer analysis framework [66] (commit `1a64af9`), but found that its analysis failed to complete within 24 h on most targets.

Table 6.2 shows wildly varying results across the three evaluated static analyses. In particular, SVF and SEADSA often report a greater number of reachable functions compared to Fuzz Introspector (e.g., up to $\sim 13\times$ on the *woff2* target). We attribute

TABLE 6.2: Static analysis results, showing the number of functions (“F”), basic blocks (“BB”), edges (“E”), and context-sensitive edges (“E_{ctx}”) reachable from the driver program. A ✗ indicates the analysis failed or did not complete (within 24 h).

Target	Fuzz Introspector				SVF				SEADSA			
	F	BB	E	E _{ctx}	F	BB	E	E _{ctx}	F	BB	E	E _{ctx}
<i>bloaty</i>	278	4,734	8,529	34,979	✗	✗	✗	✗	1,817	37,204	137,497	7,629,437
<i>curl</i>	1,327	14,694	33,323	133,601	✗	✗	✗	✗	✗	✗	✗	✗
<i>freetype2</i>	327	3,066	6,400	15,550	2,010	19,192	130,177	11,265,197	686	4,329	9,654	39,754
<i>harfbuzz</i>	10,521	28,851	87,695	259,811	12,132	32,833	101,924	4,186,525	10,250	27,897	82,532	353,143
<i>jsoncpp</i>	37	204	623	1,352	81	391	1,339	3,319	54	638	1,243	4,460
<i>lcms</i>	253	1,708	4,027	23,535	✗	✗	✗	✗	847	5,424	21,450	106,226
<i>libpng</i>	216	3,106	6,297	12,503	297	3,211	6,913	21,954	275	3,242	6,723	20,480
<i>libxslt</i>	477	7,453	15,095	112,247	1,669	28,732	102,445	8,690,806	1,538	28,389	71,749	4,575,223
<i>mbedtls</i>	816	6,790	17,783	70,045	1,135	8,462	27,626	168,129	1,094	8,338	21,937	102,192
<i>openh264</i>	21	64	141	191	761	10,364	24,342	109,202	611	10,366	19,971	101,145
<i>openssl</i>	1,535	7,494	20,635	166,769	4,698	29,835	176,491	26,363,048	✗	✗	✗	✗
<i>openpthread</i>	5,917	20,622	67,193	261,847	3,239	8,805	26,401	186,350	5,680	20,489	63,866	981,359
<i>re2</i>	545	2,762	8,346	20,659	3,411	7,284	24,093	83,466	725	1,841	5,156	19,899
<i>sqlite3</i>	1,689	27,626	65,628	470,460	2,285	32,647	132,147	21,591,726	2,059	31,730	72,517	4,617,764
<i>stb</i>	158	2,186	4,891	15,636	190	2,262	5,615	24,659	161	2,255	5,115	17,992
<i>vorbis</i>	121	1,302	2,604	8,299	187	1,783	4,007	15,855	157	1,780	3,388	13,321
<i>woff2</i>	95	2,774	4,603	10,802	1,259	4,431	10,495	20,191	1,224	4,685	10,241	25,300
<i>zlib</i>	22	567	936	1,005	23	567	986	1,068	24	569	984	1,062

this increase to the more sophisticated pointer analysis techniques used by SVF and SEADSA, allowing these analyses to resolve more indirect calls.

Notably, there is no significant correlation between the number of indirect calls (listed in Table 6.1) and this increase in reachable functions. For example, *woff2* has only 66 indirect calls and a $\sim 13\times$ increase in reachable functions. Instead, the complexity of the target’s data flow determines how significantly the Fuzz Introspector and SVF/SEADSA results differ.

The number of reachable control-flow elements also varies between SVF and SEADSA. We attribute this to algorithmic and implementation differences. Notably, SVF failed to analyze more targets than SEADSA. While both failed to analyze *curl*, SEADSA failed to analyze *openssl* (while SVF’s analysis completed) and SVF failed to analyze *bloaty* and *lcms* (which SEADSA’s analysis completed). This—combined with *cclyzer++*’s inability to complete its analysis on *any* target within 24 h—suggests that scalability issues remain in even state-of-the-art static analysis frameworks.

Table 6.3 summarizes the number of control-flow elements covered by at least one fuzzer. Comparing these results to those in Table 6.2, we see that Fuzz Introspector often undercounts the number of reachable functions (e.g., reporting 278 reachable functions in *bloaty*, while 1,101 were covered at run time). Differences between the static and dynamic analysis results become more pronounced when comparing more sensitive coverage metrics; in particular, context-sensitive edges (E_{ctx} in Tables 6.2 and 6.3). This is unsurprising, as static analysis errors (false positives and false negatives) are compounded as the analysis becomes more sensitive. This is evident

TABLE 6.3: The number of control-flow elements covered. The control-flow elements are the same as those in Table 6.2: the number of functions (“F”), basic blocks (“BB”), edges (“E”), and context-sensitive edges (“E_{ctx}”).

Target	F	BB	E	E _{ctx}
<i>bloaty</i>	1,101	5,571	8,762	25,944
<i>curl</i>	1,546	10,737	15,208	27,872
<i>freetype2</i>	1,016	9,199	13,588	37,820
<i>harfbuzz</i>	11,223	28,515	51,215	154,310
<i>jsoncpp</i>	113	664	903	2,173
<i>lcms</i>	283	1,484	2,122	5,610
<i>libpng</i>	162	1,789	2,441	4,686
<i>libxslt</i>	798	10,471	15,240	65,560
<i>mbedtls</i>	623	3,885	5,106	10,982
<i>openh264</i>	431	8,575	11,862	15,800
<i>openssl</i>	1,269	6,608	10,169	76,694
<i>openthread</i>	2,477	5,338	11,169	19,499
<i>re2</i>	580	3,049	4,441	7,381
<i>sqlite3</i>	1,762	21,755	32,931	374,597
<i>stb</i>	142	1,765	2,293	3,700
<i>vorbis</i>	131	1,280	1,728	2,753
<i>woff2</i>	101	1,798	2,449	2,934
<i>zlib</i>	20	388	547	554

in *bloaty*, where SEADSA reported 7,629,437 reachable context-sensitive edges, while only 25,944 were covered at run time. These differences (in context-sensitive edges) continue in *curl* (Fuzz Introspector), *harfbuzz* (SVF), *lcms* (SEADSA), *libxslt* (SEADSA), *mbedtls* (SVF), *openssl* (SVF), *openthread* (SEADSA), and *sqlite3* (SVF and SEADSA).

Finding

The number of reachable control-flow elements varies wildly (a) across the evaluated static analyses, and (b) compared to the number of covered control-flow elements. This suggests the static analysis results are unrealistic.

False Negatives

In Table 6.4 we summarize the accuracy of these static analysis results by examining *false negatives*: the number of control-flow elements the static analysis labeled unreachable but were reachable by a fuzzer. Quantifying false positives is undecidable; there is no way of knowing if a control-flow element (e.g., CFG edge) is unreachable or if the fuzzer failed to generate a valid input for covering that element. Because false negatives (and false positives) can only occur due to indirect function calls³, we focus on the function call graph.

³Or bugs in the static analyzer, which we ignore.

TABLE 6.4: The number of static analysis *false negatives*; i.e., the number of indirect function calls that the static analysis (one of “Fuzz Introspector”, “SVF”, and “SEADSA”) labeled unreachable but were covered by a fuzzer. The best performing static analyzer for each target is highlighted in green (lower is better). A ✗ indicates the static analysis failed.

Target	Fuzz Introspector	SVF	SEADSA
<i>bloaty</i>	2,155	✗	2,160
<i>curl</i>	451	✗	✗
<i>freetype2</i>	572	163	407
<i>harfbuzz</i>	285	62	94
<i>jsoncpp</i>	137	137	162
<i>lcms</i>	80	✗	2
<i>libpng</i>	14	6	6
<i>libxslt</i>	177	4	20
<i>mbedtls</i>	45	0	1
<i>openh264</i>	82	0	7
<i>openssl</i>	414	87	✗
<i>openthread</i>	29	20	125
<i>re2</i>	176	165	211
<i>sqlite3</i>	235	22	26
<i>stb</i>	8	0	0
<i>vorbis</i>	23	5	5
<i>woff2</i>	36	20	29
<i>zlib</i>	3	3	0

Despite failing on the most targets, SVF was the best performer with the lowest number of false negatives on the most targets. In comparison, Fuzz Introspector was the worst performer. Despite this—and its relatively straightforward approach to resolving indirect function calls (Section 6.2.3)—Fuzz Introspector outperforms SEADSA on four targets (*bloaty*, *jsoncpp*, *openthread*, and *re2*).

We analyzed these false negatives and attribute most of them to *custom memory allocators*. In Section 5.4.4 we discussed the prevalence and impact of custom memory allocators on DATAFLOW. In particular, custom memory allocators required user intervention to ensure *def-use* sites were appropriately instrumented. Here, custom memory (de)allocation routines are passed to functions inside structs. Dynamically (de)allocating memory thus requires making indirect function calls via struct offsets, which SVF and SEADSA are unable to reason about. We hypothesize that a “structure-sensitive” analysis—such as that used by *cclizer++* [13, 66]—would perform better in these circumstances. Unfortunately, this increased accuracy incurs a significant performance penalty, often resulting in the analysis failing to complete (within 24 h).

Finding

The high variance in results across the three static analyzers, together with the number of false negatives, indicates that even state-of-the-art static analyzers are inadequate for quantifying the upper-bound of a fuzzer’s state space search.

6.4 Discussion

Liyanage et al. [129] also investigated the use of static analysis to quantify *fuzzer effectiveness*; computing an upper bound on the number of reachable coverage elements.⁴ Liyanage et al. [129] also found that static analysis significantly over-counted reachable coverage, and our findings reinforce their results. Consequently, we encourage static analysis developers to further improve the accuracy and scalability of their tools, particularly on “real-world” targets. In particular, this could include developing new benchmark suites that contain larger, more complex codebases more indicative of “real-world” software.

6.5 Future Work

Based on our results, we propose the following ideas for future work.

Root cause quantification for missing control-flow edges. In Section 6.3.2 we analyzed the accuracy of three static analyzers by quantifying the number of unlabeled call graph edges. While this analysis was manual, automatic techniques would help in improving the accuracy of these static analyzers. For example, Chakraborty et al. [30] develop a technique to “*automatically quantify the relative importance of different root causes of call graph unsoundness for a set of target [JavaScript] applications*”. While their study focused on JavaScript call graphs, their techniques are generalizable to other languages (here, we are interested in LLVM intermediate representation (IR)), and could be used to improve the accuracy of LLVM-based static analyzers (which we found to be severely lacking).

Quantifying data flow. In Section 5.5.4 we quantified data-flow coverage by computing the set of *def-use* chains in a target. An alternative approach is to quantify data flow over *program variable values*. This approach is akin to the data abstractions used in abstract interpretation and model checking (discussed in Chapter 1). We

⁴Our work was carried out concurrently. Liyanage et al. [129] evaluated SVF on five targets from the Software-artifact Infrastructure Repository, whereas we evaluated a larger number of targets from FUZZBENCH and also evaluate Fuzz Introspector and SEADSA.

propose using a static analyzer (e.g., CLAM, a static analyzer for LLVM built on the CRAB library [79]) to compute inductive invariants—computed as basic block pre- and post-conditions—for each variable in \mathcal{P} . These invariants provide the range of values a variable can take. By instrumenting these variables, we can track their value at run time, providing another measure of “state space coverage”. We leave the implementation and evaluation of such analyses for future work.

6.6 Chapter Summary

This chapter discussed the difficulties in quantifying a fuzzer’s state space search. To address this, we proposed using context-sensitive static analyses to quantify the control-flow elements of a fuzzer’s state space search. Unfortunately, our results also show that even state-of-the-art static analyses are far from accurate on typical fuzzer targets; they either failed to complete their analysis or produced unrealistic results. We hope these results and ideas for future work inspire others to improve the accuracy and scalability of static analyses for real-world codebases.

In the final chapter, we provide a comprehensive evaluation of the full range of coverage metrics discussed throughout this dissertation.

Chapter 7

Comparison of Coverage Metrics

This final chapter provides a comprehensive evaluation of the full range of coverage metrics discussed throughout this dissertation. To do this, we propose a methodology for fairly comparing coverage metrics.

7.1 Introduction

As described in Chapter 3, fuzzers use a coverage metric to measure their progress in exploring a target’s state space. This coverage metric abstracts the target’s state space, leading the fuzzer to uncover new behaviors in the target (by revealing new states in the state space). Moreover, fuzzers are often evaluated based on their coverage of a target’s state space: fuzzer f_1 is “better” than fuzzer f_2 if f_1 explores more of the target’s state space than f_2 . Popular fuzzer benchmarks (e.g., FUZZBENCH [144]) exemplify these ideas, relying solely on *aggregate control-flow coverage* accumulated over a fuzzing campaign. However, as we have repeatedly shown throughout this dissertation, this only covers a single dimension of a target’s state space. How can we more-accurately characterize a fuzzer’s state space search?

Chapter outline. We begin by demonstrating the weaknesses of current approaches to measuring and comparing fuzzers based on coverage profiles, motivating the need for alternative techniques (Section 7.2). We then use the dynamic analysis introduced in Chapter 6 to evaluate the performance of 15 coverage metrics—spanning the full range of control- and data-flow-based metrics discussed in Chapter 3 (Section 7.3).

7.2 Motivation

Coverage measures (e.g., number of lines of code) are commonly used to evaluate and compare a fuzzer’s performance, serving as a proxy for a fuzzer’s bug-finding




FIGURE 7.1: Coverage exploration by AFL++ over ten repeated 24h trials with two coverage metrics—coarse-grained probabilistic edge coverage and context-sensitive probabilistic edge coverage—on FUZZBENCH’s *libpng*. The x -axis shows wall-clock time (log scale in seconds) and the y -axis shows the number of covered measurement points (source, edges, and context-sensitive edges in Figs. 7.1a to 7.1c, respectively). Each figure shows mean coverage and 95 % bootstrap confidence interval (CI).

ability (after all, you cannot find bugs in code never executed) [115]. Consequently, we can use coverage to reason about the residual risk of stopping a fuzzing campaign too early (i.e., while the fuzzer is still discovering new program behaviors) [20]. Unfortunately, reasoning about coverage as an approximation for exercised program behaviors is hard. We discuss the difficulties here.

Visualizing fuzzers’ progress via plots showing coverage growth over time are a common sight in fuzzer evaluations [9, 16, 19, 21, 36, 37, 39, 40, 43, 46, 51, 62, 67, 68, 85, 90, 91, 100, 125, 133, 144, 157, 173, 174, 184, 186, 190, 198, 213, 217, 227, 229]. We continue this trend in Fig. 7.1, showing the coverage achieved by AFL++ on FUZZBENCH’s *libpng* target with two coverage metrics: coarse-grained probabilistic edge coverage (AFL edge coverage, denoted \mathcal{P}_{AFL}) and context-sensitive probabilistic edge coverage (AFL edge coverage augmented with calling context, denoted \mathcal{P}_{ctx}).¹

We generate these plots using a “record and replay” (R&R) approach. This approach allows us to view the progress made fuzzing \mathcal{P}_{AFL} through the “point-of-view” (PoV) of \mathcal{P}_{ctx} (and *vice versa*). R&R: (i) records \mathcal{P}_{AFL} ’s current input at regular intervals (here, every 1,000 inputs), storing the results in \mathcal{S}_{AFL} ; and then (ii) replays \mathcal{S}_{AFL} through \mathcal{P}_{ctx} (*post hoc*). This approach, proposed by Aschermann [6], ensures an unbiased comparison between coverage metrics. This is akin to the approach taken by FUZZTASTIC [127], which uses additional instrumentation to measure collision-free block coverage in parallel with the fuzzer (rather than examining the queue *post hoc*).

¹The highlight color reflects the plot colors.

Figure 7.1a echoes the approach used by FUZZBENCH [144], where we replayed \mathcal{S}_{AFL} and \mathcal{S}_{ctx} through \mathcal{P}_{src} instrumented with Clang’s source-based coverage (operating on the Clang AST and preprocessor) [209]. FUZZBENCH uses Clang’s source-based coverage as an independent code coverage metric to enable comparisons *across* fuzzers. While this approach removes bias toward a particular coverage metric, it also means differences between metrics—and their implementations—are lost (e.g., when comparing coverage achieved by fuzzers using more fine-grained metrics).

This “information loss” is evident in Fig. 7.1: viewed through the \mathcal{P}_{AFL} and \mathcal{P}_{src} PoV (Figs. 7.1a and 7.1b, respectively), the coverage achieved across the two campaigns appears identical. However, the \mathcal{P}_{ctx} PoV (Fig. 7.1c) clearly shows that fuzzing \mathcal{P}_{ctx} explores more of \mathcal{P} ’s state space (due to context sensitivity). *How, then, should we make cross-fuzzer coverage comparisons?* The remainder of this chapter answers this question.

7.3 Evaluation

Our evaluation aims to answer the following research questions:

- RQ 1** How do the coverage expansion capabilities of different control- and data-flow-based metrics compare? (Section 7.3.2)
- RQ 2** Can coverage metrics be combined to complement each other, to uncover new states in the target’s state space (and thus share this information with other fuzzers—in an ensemble configuration—to improve their state space search)? (Section 7.3.3)

7.3.1 Methodology

Fuzzer Selection

We evaluate the 15 coverage metrics listed in Table 7.1. These coverage metrics span the full range of coverage metrics discussed throughout this dissertation. All 15 are (re)implemented as LLVM (v14) passes in AFL++ (commit 7101192). The (context-sensitive) edge metrics (E_{coarse} , E_{LTO} , E_{SANCOV} , $E_{1\text{-ctx}}$, $E_{2\text{-ctx}}$, and $E_{3\text{-ctx}}$) are described in Sections 3.2.2 and 3.2.4. The data dependency graph (DDG) (E_{DDG}), memory accesses (E_{MA}), and n -gram paths metrics ($P_{2\text{-gram}}$, $P_{4\text{-gram}}$, and $P_{8\text{-gram}}$) are described in Sections 3.2.3, 3.2.5 and 3.2.7, respectively. Finally, *def-use* chain coverage ($DF_{\text{A+S/A}}$, $DF_{\text{A+S/O}}$, and $DF_{\text{A+S/V}}$) is described in Chapter 5. We ported E_{MA} from its original QEMU-based implementation [214] to operate on the LLVM intermediate representation (IR). E_{DDG} uses DDFuzz’s implementation of the DDG [141], while $DF_{\text{A+S/A}}$,

TABLE 7.1: Evaluated fuzzers and their coverage metrics. B has no C, while E_{LTO} , and E_{SANCOV} uses a variable-sized C.

Fuzzer	Description	Map size (KiB)
B	Blackbox	–
E_{coarse}	Coarse-grain probabilistic edge	64
E_{LTO}	Link-time optimization (LTO) edge	–
E_{SANCOV}	SANCOV LTO fine-grain edge	–
E_{1-ctx}	Edge \oplus calling context ($k = 1$)	64
E_{2-ctx}	Edge \oplus calling context ($k = 2$)	64
E_{3-ctx}	Edge \oplus calling context ($k = 3$)	64
E_{DDG}	Edge \oplus DDG	64
E_{MA}	Edge \oplus memory accesses (reads + writes)	128
P_{2-gram}	Path. 2-gram edge	64
P_{4-gram}	Path. 4-gram edge	64
P_{8-gram}	Path. 8-gram edge	64
$DF_{A+S/A}$	<i>Def-use</i> chain. Array and struct <i>def</i> \oplus access <i>use</i> sites	1,024
$DF_{A+S/O}$	<i>Def-use</i> chain. Array and struct <i>def</i> \oplus access <i>use</i> offsets	1,024
$DF_{A+S/V}$	<i>Def-use</i> chain. Array and struct <i>def</i> \oplus accessed value <i>use</i> sites	1,024

$DF_{A+S/O}$, and $DF_{A+S/V}$ use DATAFLOW’s implementation of *def-use* chain coverage [90]. The other fuzzers listed in Table 7.1 were already implemented in AFL++. Using the same fuzzer (AFL++) and instrumentation framework (LLVM) ensures a fair evaluation [118].

Target Selection

We use the same configuration of FUZZBENCH [144] described in Section 6.3.1.

Experimental Setup

We run one fuzzing *campaign* per {target \times coverage metric} combination. Each campaign consists of ten independent 24 h *trials* (consistent with the recommendations by Klees et al. [115]). This amounts to over 8 CPU-yr of fuzzing. All experiments were carried out using the configuration described in Section 6.3.1. Additionally, we modify FUZZBENCH to support our R&R approach (Section 7.2). R&R ensures: (i) an unbiased comparison of coverage metrics, because inputs have not been retained based on a specific coverage metric; (ii) differences between metrics (and their implementations) are retained; and (iii) the sensitivity of metrics can be compared.

7.3.2 Comparison of Coverage Metrics (RQ 1)

We use our dynamic analysis (described in Section 6.2.2) to compare the coverage expansion (and thus state space search) performance of different control- and data-flow-based coverage metrics. Given the unreliability of our static analysis results

(Section 6.3), we do not quantify coverage with respect to a static interprocedural control-flow graph (ICFG) (as we did in Section 5.5.4). Instead, Table 7.2 summarizes the number of context (in)sensitive edges covered by AFL++ using the 15 coverage metrics listed in Table 7.1. The mean coverage over ten repeated 24 h trials with 95 % CI is shown. We use the Mann-Whitney U -test [139] to statistically compare coverage across fuzzers: two fuzzers cover the same number of context-(in)sensitive edges if the Mann-Whitney U -test's p -value > 0.05 . We first compare performance using context-insensitive edge coverage, before turning our attention to context-sensitive edge coverage.

Context-Insensitive Edge Coverage

Context-insensitive edge coverage is the approach traditionally used to compare fuzzer performance. We use this approach in Table 7.2a to summarize and compare the aggregate coverage achieved by each coverage metric. Based on these results, E_{LTO} and E_{SANCOV} were the best-performing metrics, achieving the most coverage on the most targets. This echoes our result in Chapter 5. Here, both E_{LTO} and E_{SANCOV} consistently outperformed the other edge metric (E_{coarse}), reinforcing the importance of eliminating hash collisions in the coverage map. Notably, E_{SANCOV} 's higher runtime cost (compared to E_{LTO}) did not translate to better performance.

Similar to our results in Section 7.2, adding context sensitivity to coarse-grained probabilistic edge coverage (E_{1-ctx} , E_{2-ctx} , and E_{3-ctx}) failed to meaningfully improve overall coverage. Notably, the performance of the context-sensitive edge coverage metrics improved slightly (mean 4 % increase in edge coverage across all targets) as the number of calling contexts increased (from $k = 1$ to $k = 3$). However, even when $k = 3$, the two LTO edge metrics generally outperformed the context-sensitive edge metrics. We attribute this to hash collisions in the coverage map, due to the large number of context-sensitive edges in the target program (per the results in Table 6.4).

Approximating path coverage with n -gram coverage (P_{2-gram} , P_{4-gram} , and P_{8-gram}) produced similar results to context-sensitive edge coverage, with little statistically-significant improvement over the two LTO edge metrics. Unlike the context-sensitive edge metrics, average performance worsened as n increased (mean 10 % decrease in edge coverage across all targets). Intuitively, we expected targets with fewer basic blocks to perform better, because they would be less prone to hash collisions in the coverage map. However, this intuition was not consistently borne out in our results. While it may explain the coverage increase—as n increased—in *zlib* (the target with the fewest number of basic blocks, per Table 6.1), it does not explain the results for

other targets (e.g., $P_{8\text{-gram}}$ was the best performer on *openthread*, despite it having a higher-than-average basic block count).

The E_{MA} and E_{DDG} metrics both attempt to incorporate data flow information with traditional edge coverage techniques. The approach used by E_{MA} —to xor memory accesses with coarse-grained probabilistic edge coverage—is the less effective of the two metrics. The large number of instrumented memory accesses significantly (a) reduces AFL++’s iteration rate, and (b) increases the size of the fuzzer’s queue, resulting in *queue explosion* and making it harder for the fuzzer’s scheduler to select a suitable seed to fuzz. Ultimately, this makes E_{MA} one of the worst-performing metrics. In contrast, E_{DDG} —which combines a DDG with coarse-grained probabilistic edge coverage—was the (equal) second-best-performing metric, achieving the (equal) highest coverage on ten targets. Unlike E_{MA} , E_{DDG} optimizes the inclusion of data flow information, minimizing run-time overheads and increasing coverage expansion. The E_{DDG} results reinforce one of our key findings from Section 5.5.3: fuzzers guided by control flow and data flow should be combined to maximize fuzzing outcomes.

Unsurprisingly, the worst results occurred when no coverage feedback was used (B). In extreme cases (e.g., *harfbuzz*, *lcms*), this resulted in half the coverage achieved by the control-flow-based metrics. However, to our dismay, *def-use* chain coverage ($DF_{A+S/A}$, $DF_{A+S/O}$, and $DF_{A+S/V}$) performed only marginally better than no coverage at all. Similar to our results in Section 5.5, $DF_{A+S/O}$ was the worst-performing of the three *def-use* chain metrics. Despite $DF_{A+S/V}$ ’s increased sensitivity (per Section 5.3.1), there was no statistically significant improvement over $DF_{A+S/A}$. Again, the run-time overheads associated with tracking *def-use* chains had a significant impact on performance.

Finding

Context-insensitive edge coverage is the traditional approach for measuring and comparing fuzzers’ state space search. Based on this approach, E_{LTO} and E_{SANCOV} were the best-performing coverage metrics. We attribute this result to their (a) low run-time overhead (leading to faster iteration rates), and (b) lack of hash collisions in the coverage map.

Context-Sensitive Edge Coverage

Section 6.2.1 introduced our approach for measuring *context-sensitive* control-flow coverage. Such an approach is important for uniquely identifying execution points. Similar to Table 7.2a, Table 7.2b summarizes aggregate coverage. Importantly, the

previous context-insensitive results are reflected here, with E_{LTO} , E_{SANCOV} , and E_{DDG} once again the best performers, achieving the highest coverage on the most targets. However, many of the other metrics—in particular, E_{1-ctx} , E_{2-ctx} , E_{3-ctx} , and P_{2-gram} —appeared to perform better when accounting for context sensitivity.

Following the results in Section 7.2 (particularly in Fig. 7.1), the improvements made by the context-sensitive edge and n -gram coverage metrics are not surprising. While n -gram coverage does not explicitly incorporate calling context, recording the previous n basic blocks implicitly captures calling context at function call (and return) sites. Once again, B , E_{MA} , $DF_{A+S/A}$, $DF_{A+S/O}$, and $DF_{A+S/V}$ were the worst performers.

Given our results, it is intuitive to assume that “the best performing metric is the one used to compare metrics”. For example, when comparing fuzzers’ ability to expand edge coverage, the best performers used traditional edge coverage as their coverage metric (Section 7.3.2). Similarly, when comparing fuzzers’ ability to expand context-sensitive edge coverage, the best performers incorporated context sensitivity in their coverage metric. However, our previous results in Chapter 5 caution against such thinking. In Section 5.5.4, edge coverage outperformed data-flow coverage when comparing performance using a data-flow-based metric (in this case, *def-use* chains). If anything, our results here and in Chapter 5 lead us to conclude that “the best performing metric is one based on control-flow”.

Finding

Our context-sensitive results largely mirror our context-*insensitive* results: traditional edge coverage (as implemented in E_{LTO} and E_{SANCOV}) is the best-performing coverage metric. While adding context sensitivity is important for uniquely identifying execution points in a target, this ability is not strictly required when comparing fuzzer performance (based on coverage expansion). However, it may still be useful in other circumstances (e.g., to understand fuzz blockers).

7.3.3 Complementary Metrics (RQ 2)

Güler et al. [78] introduced CUPID, a system for automatically identifying fuzzers that complement each other when executed collaboratively in an ensemble configuration. CUPID is based on the intuition that fuzzers covering different parts of a target’s state space should benefit from sharing their progress with other fuzzers. Importantly, CUPID is not simply based on which individual fuzzers achieve the

“best” coverage. Instead, CUPID “measures the degree in which multiple fuzzers are complementary... [computing] the union of the expected mean code coverage”.

We adopt CUPID’s approach to understand how well the 15 coverage metrics in Table 7.1 complement each other, and how much (or little) their state space search overlaps. To do this, we compute a *probability map* of how often a fuzzer was able to cover each control-flow graph (CFG) edge in a 24 h campaign (across ten independent trials). Using this probability map, we can understand (a) which coverage metrics uncover program states that other metrics do not, and (b) which coverage metrics may benefit from an ensemble fuzzing configuration.

Figures 7.2 and 7.3 illustrate the probability maps for context-insensitive and context-sensitive CFG edges (respectively) in *lcms*, while Fig. 7.4 illustrates the probability maps for context-insensitive CFG edges in *sqlite3* (the context-sensitive version contains too many edges to visualize).² Each cell represents a single edge in the context-insensitive CFG, while the color of a cell represents the probability of covering that edge: darker cells correspond to high probabilities. While it is difficult to discern the differences in individual edges (especially as the number of edges increases, becoming even more pronounced when context sensitivity is taken into account), these illustrations help visualize high-level differences across the 15 coverage metrics. In particular, the number of red cells in Figs. 7.2 and 7.3 reinforces the (relative) poor performance of B, E_{MA}, DF_{A+S/A}, DF_{A+S/O}, and DF_{A+S/V}.

We use these probability maps to determine how much a given coverage metric complements another (when combined in an ensemble fuzzing configuration). As described by Güler et al. [78], the combined probability of a set of fuzzers \mathcal{F} covering a given CFG edge is calculated as

$$P_{\mathcal{F}}(e) = 1 - \prod_{f \in \mathcal{F}} (1 - P_f(e)), \quad (7.1)$$

where \mathcal{F} is the set of fuzzers, e is the CFG edge the combined probability is being calculated for, and P_f is the probability of fuzzer f covering edge e .

Using Eq. (7.1), the sum of all edge probabilities can be calculated to determine how well the fuzzers in \mathcal{F} complement each other. This is calculated as

$$\mathbb{E}_{\mathcal{F}} = \sum_{e \in E} P_{\mathcal{F}}(e), \quad (7.2)$$

²We select these targets for illustration because they displayed the highest variance across coverage metrics in Tables 7.2a and 7.2b.

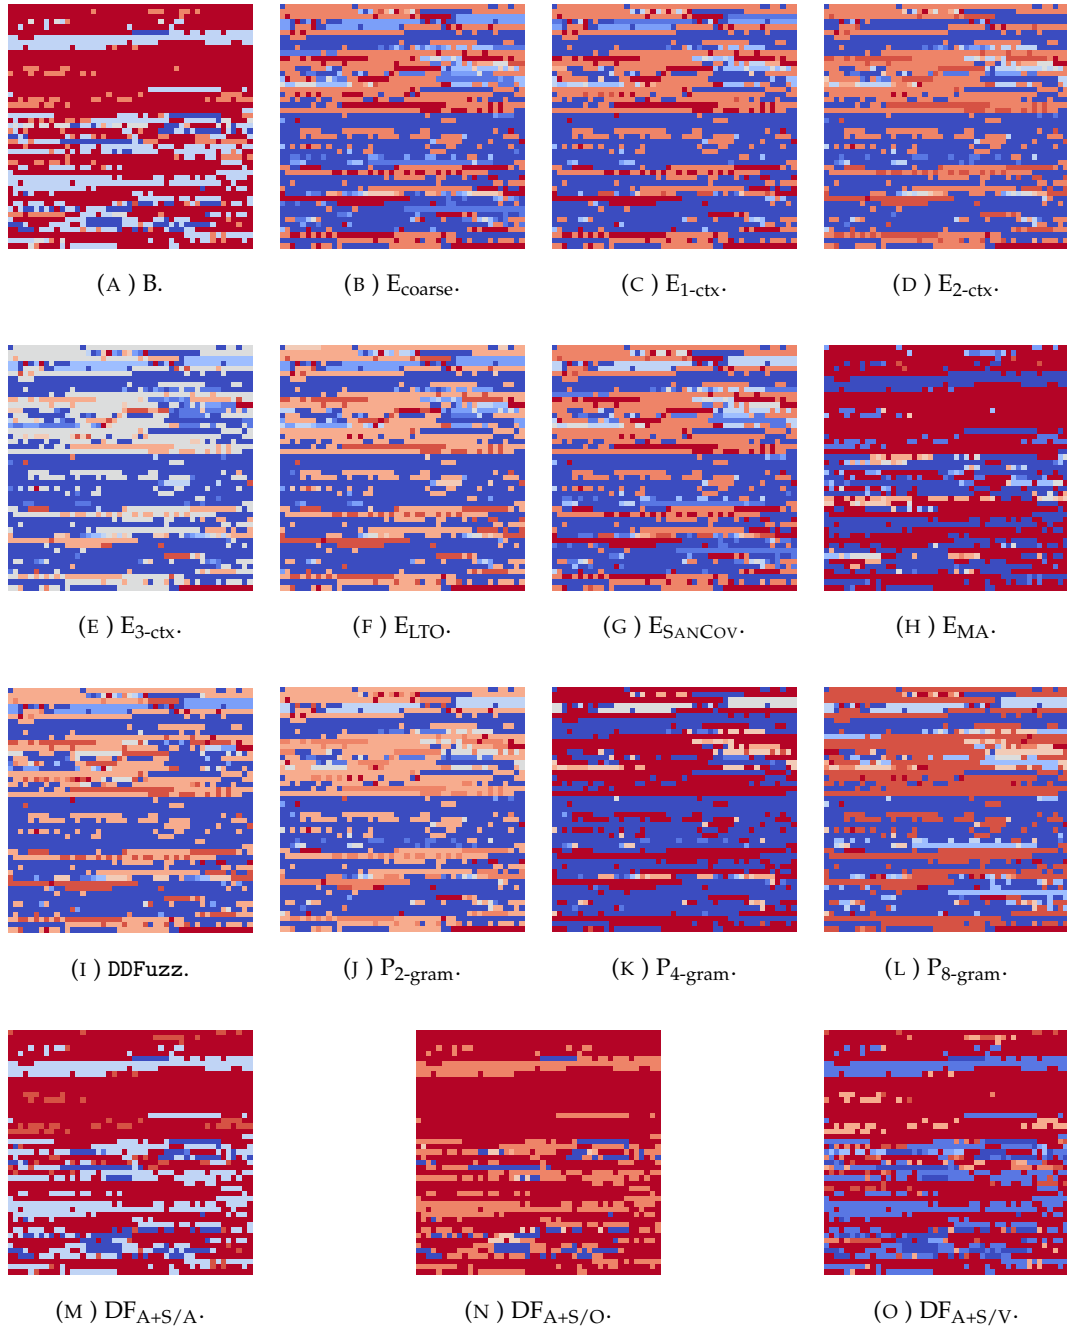


FIGURE 7.2: Probability maps for *lcms*. Each cell represents a single edge in the context-insensitive CFG. The color of a cell represents the probability that this edge will be covered by the given coverage metric (red is probability = 0, blue is probability = 1).

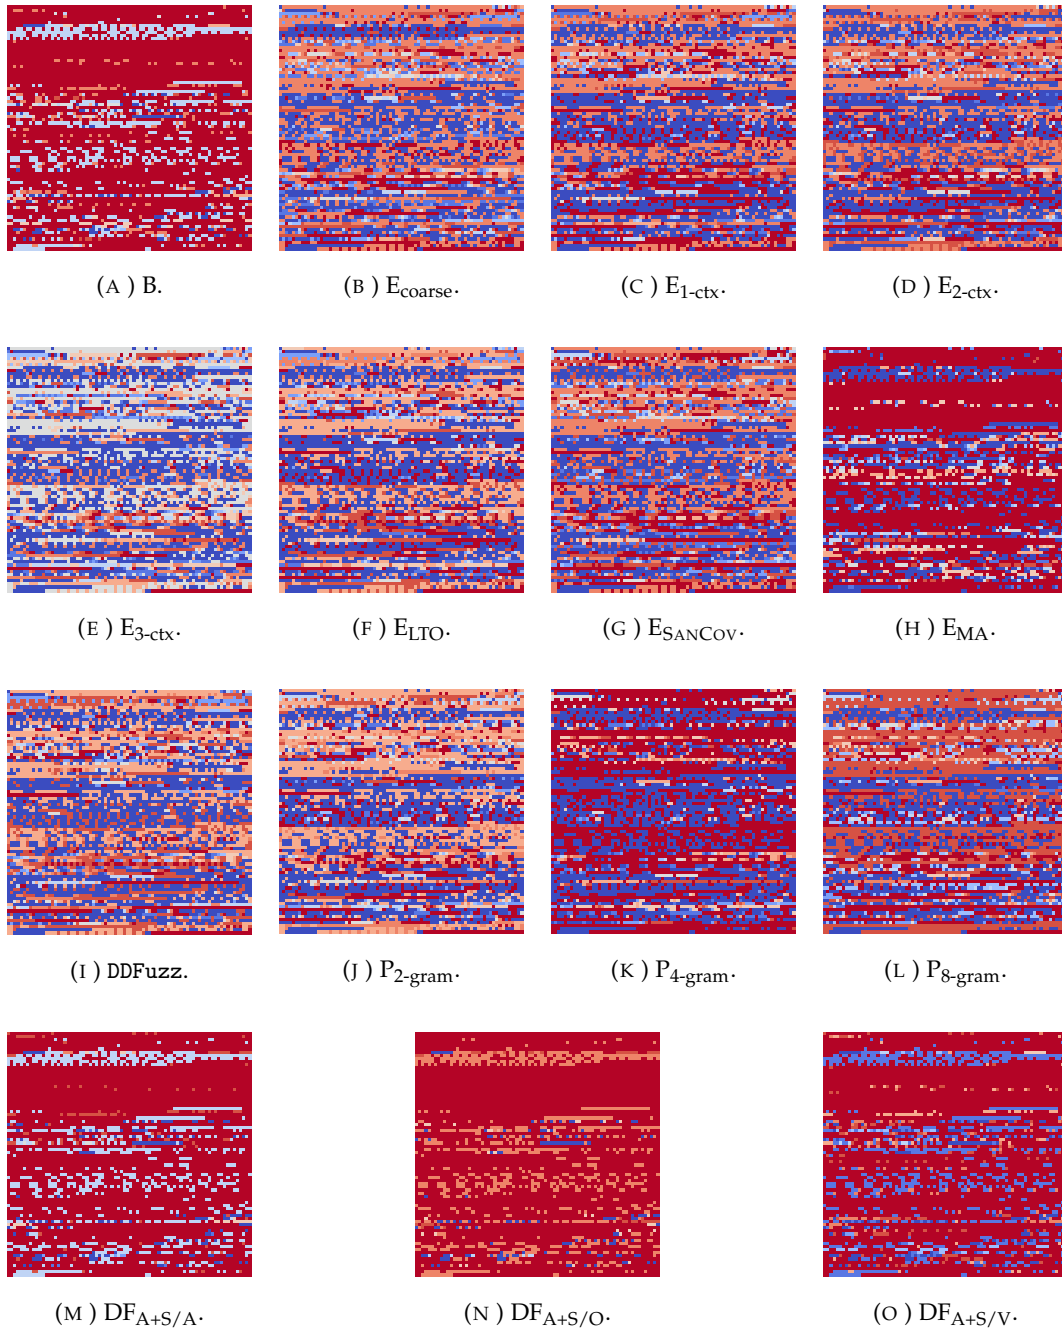


FIGURE 7.3: Probability maps for *lcms*. Each cell represents a single edge in the context-sensitive CFG. The color of a cell represents the probability that this edge will be covered by the given coverage metric (red is probability = 0, blue is probability = 1).

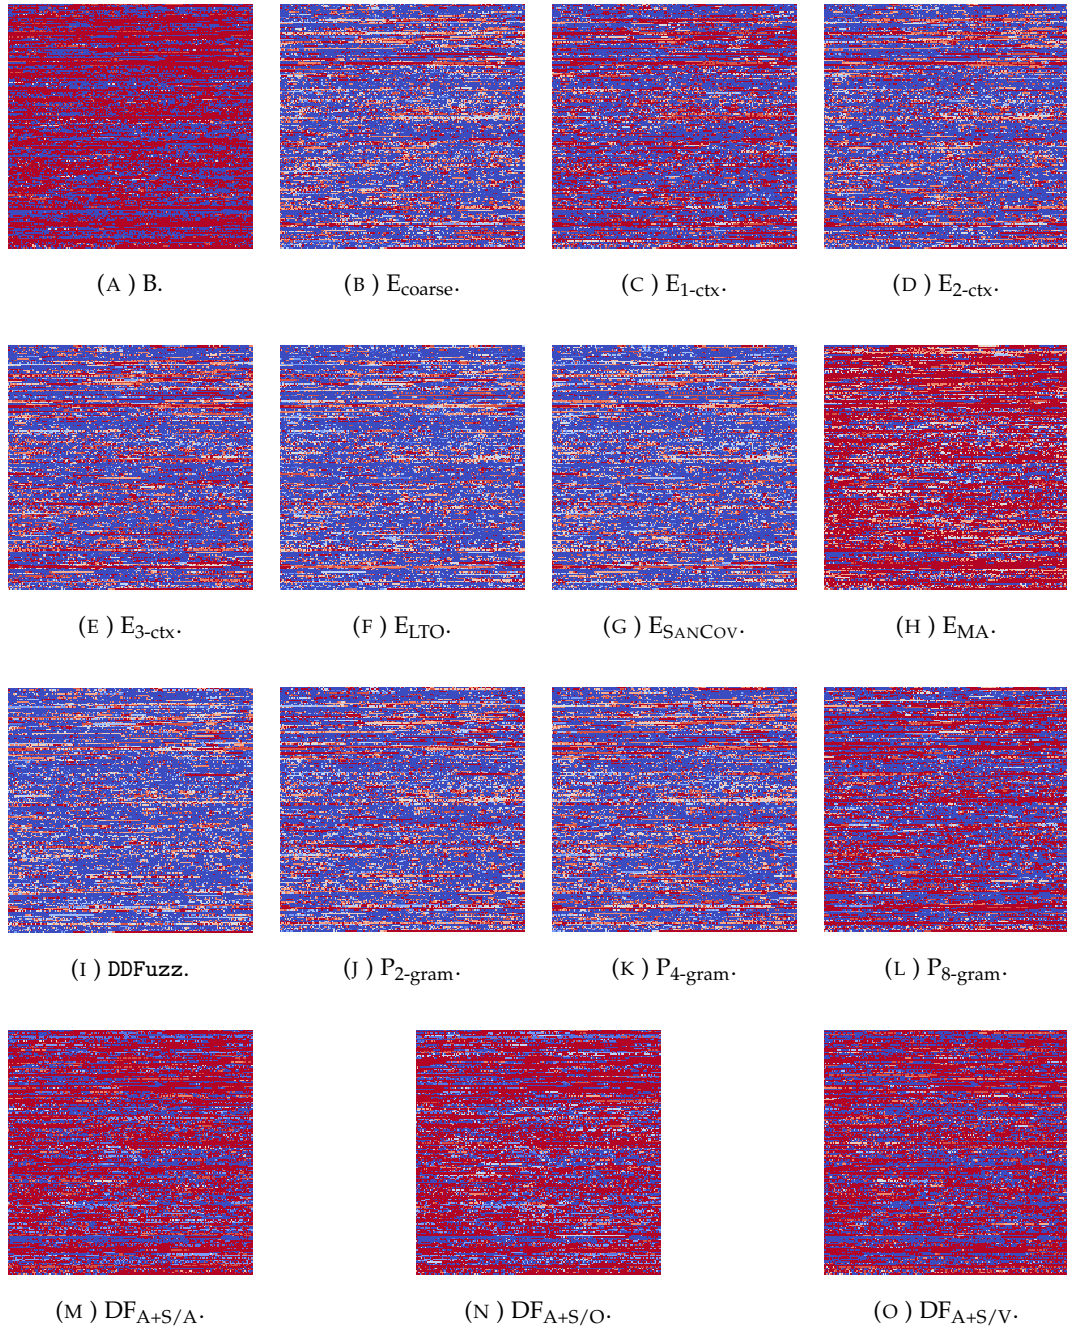


FIGURE 7.4: Probability maps for *sqlite3*. Each cell represents a single edge in the context-insensitive CFG. The color of a cell represents the probability that this edge will be covered by the given coverage metric (red is probability = 0, blue is probability = 1).

where E is the set of all CFG edges in a target. Equation (7.2) represents the expected number of edges that would be covered by the set of fuzzers \mathcal{F} .

Figure 7.5 shows the predicted improvements in context-(in)sensitive CFG edge coverage when pairing coverage metrics. These improvements are based on the coverage results in Tables 7.2a and 7.2b. Each heatmap shows the predicted improvements for a single target, where each cell contains (a) the expected number of covered edges (per Eq. (7.2)), and (b) the percentage increase over only using a single coverage metric. The colors reflect the relative performance of the coverage metric: red is worst performance, blue is best performance.

These heatmaps make it clear—across a majority of targets—that accurate edge coverage metrics (E_{LTO} and E_{SANCOV}) are superior to all other coverage metrics when maximizing context-*insensitive* coverage. In particular, these metrics benefitted little when combined with any of the other 13 metrics (only up to $\sim 3.5\%$). The exceptions to this result were *mbedtls*, *sqlite3*, and *stb*, which were predicted to improve their coverage by 4%, 10%, and 11%, respectively. Interestingly, these improvements occurred when E_{DDG} was incorporated, suggesting that incorporating some form of data-flow-based coverage may yield small improvements to overall coverage.

These results generally hold when considering context-*sensitive* coverage. While E_{LTO} and E_{SANCOV} continue to dominate, the performance gains are much larger (e.g., up to 61% on *libpng*) and occur when pairing with a wider range of coverage metrics (e.g., context-sensitive edge coverage and n -gram coverage).

Finding

The best-performing coverage metrics— E_{LTO} and E_{SANCOV} —were predicted to benefit little when combined with other coverage metrics. When performance was predicted to improve, it was generally when E_{DDG} was incorporated, suggesting that some form of data-flow-based coverage may yield small improvements to overall coverage.

7.4 Comparison to Previous Studies

Salls et al. [183] and Wang et al. [214] have also conducted studies on the performance of fuzzer coverage metrics. However, these studies: (i) did not take into account data-flow-based metrics; (ii) used a smaller set of “real-world” targets; and (iii) did not compare fuzzers based on coverage (instead focusing on counting crashes in synthetic benchmarks such as the Cyber Grand Challenge (CGC) and LAVA-M).

B	5415 (6.81)	5562 (9.72)	5314 (4.83)	5774 (13.91)	7140 (40.85)	7641 (50.74)	7578 (49.48)	7963 (57.09)	7705 (52.00)	5529 (9.06)	7782 (53.50)	7230 (42.61)	7109 (40.23)	7370 (45.38)	6067 (19.67)
DF _{A+S/A}	5562 (7.21)	5582 (7.59)	5390 (3.90)	5772 (11.26)	7158 (37.97)	7642 (47.30)	7571 (45.93)	7962 (53.48)	7702 (48.46)	5643 (8.77)	7780 (49.97)	7224 (39.25)	7108 (37.01)	7366 (41.98)	6095 (17.48)
DF _{A+S/O}	5314 (17.95)	5390 (19.63)	4993 (10.82)	5683 (26.13)	7078 (57.10)	7637 (69.50)	7566 (67.94)	7960 (76.67)	7699 (70.88)	5405 (19.97)	7776 (72.58)	7217 (60.18)	7098 (57.55)	7358 (63.31)	5877 (30.45)
DF _{A+S/V}	5774 (2.55)	5772 (2.51)	5683 (0.92)	5842 (3.76)	7223 (28.27)	7649 (35.84)	7579 (34.60)	7964 (41.43)	7707 (36.87)	5843 (3.76)	7782 (38.22)	7231 (28.43)	7116 (26.38)	7376 (31.00)	6275 (11.45)
E _{1-ctx}	7140 (8.10)	7158 (8.36)	7078 (7.16)	7223 (9.35)	7524 (13.92)	7824 (18.45)	7795 (18.01)	8058 (22.00)	7950 (20.36)	7160 (8.40)	7985 (20.89)	7655 (15.90)	7593 (14.95)	7716 (16.81)	7232 (9.49)
E _{2-ctx}	7641 (0.13)	7642 (0.14)	7637 (0.07)	7649 (0.23)	7824 (2.52)	7917 (3.74)	7938 (4.02)	8128 (6.51)	8047 (5.44)	7647 (0.21)	8080 (5.87)	7819 (2.46)	7782 (1.97)	7880 (3.25)	7664 (0.42)
E _{3-ctx}	7578 (0.21)	7571 (0.11)	7566 (0.06)	7579 (0.22)	7795 (3.08)	7938 (4.97)	7885 (4.27)	8134 (7.56)	8070 (6.72)	7581 (0.25)	8096 (7.07)	7815 (3.34)	7774 (2.80)	7869 (4.06)	7607 (0.60)
E _{DDG}	7963 (0.05)	7962 (0.04)	7960 (0.01)	7964 (0.06)	8058 (1.24)	8128 (2.12)	8134 (2.19)	8229 (3.40)	8220 (3.28)	7965 (0.07)	8239 (3.52)	8054 (1.19)	8026 (0.84)	8091 (1.66)	7970 (0.13)
E _{LTO}	7705 (0.11)	7702 (0.07)	7699 (0.03)	7707 (0.13)	7950 (3.29)	8047 (4.54)	8070 (4.85)	8220 (6.80)	7994 (3.86)	7708 (0.15)	8076 (4.92)	7872 (2.28)	7818 (1.58)	7950 (3.29)	7722 (0.33)
E _{MA}	5529 (7.10)	5643 (9.31)	5405 (4.72)	5843 (13.19)	7160 (38.71)	7647 (48.15)	7581 (46.86)	7965 (54.31)	7708 (49.33)	5588 (8.26)	7785 (50.81)	7235 (40.16)	7118 (37.89)	7377 (42.91)	6089 (17.97)
E _{SANCOV}	7782 (0.11)	7780 (0.09)	7776 (0.03)	7782 (0.12)	7985 (2.75)	8080 (3.94)	8096 (4.16)	8239 (6.00)	8076 (3.89)	7785 (0.14)	8093 (4.11)	7939 (2.13)	7898 (1.60)	8010 (3.04)	7806 (0.42)
E _{coarse}	7230 (0.25)	7224 (0.17)	7217 (0.07)	7231 (0.27)	7655 (6.15)	7819 (8.43)	7815 (8.37)	8054 (11.68)	7872 (9.16)	7235 (0.33)	7939 (10.08)	7538 (4.52)	7496 (3.94)	7670 (6.35)	7266 (0.75)
P _{2-gram}	7109 (0.24)	7108 (0.22)	7098 (0.09)	7116 (0.34)	7593 (7.07)	7782 (9.73)	7774 (9.62)	8026 (13.18)	7818 (10.25)	7118 (0.36)	7898 (11.37)	7496 (5.70)	7380 (4.06)	7597 (7.13)	7146 (0.76)
P _{4-gram}	7370 (0.26)	7366 (0.20)	7358 (0.09)	7376 (0.34)	7716 (4.96)	7880 (7.19)	7869 (7.05)	8091 (10.06)	7950 (8.15)	7377 (0.35)	8010 (8.96)	7670 (4.33)	7597 (3.35)	7702 (4.78)	7384 (0.45)
P _{8-gram}	6067 (8.54)	6095 (9.04)	5877 (5.15)	6275 (12.27)	7232 (29.39)	7664 (37.11)	7607 (36.10)	7970 (42.59)	7722 (38.15)	6089 (8.94)	7806 (39.65)	7266 (30.00)	7146 (27.84)	7384 (32.11)	6081 (8.80)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(A) *bloaty* (context-insensitive).

B	14407 (9.39)	14708 (11.68)	14074 (6.86)	15109 (14.72)	20505 (55.69)	21698 (64.75)	21865 (66.02)	23147 (75.75)	20691 (57.10)	14740 (11.92)	21575 (63.82)	19289 (46.46)	18600 (41.23)	19997 (51.83)	15838 (20.25)
DF _{A+S/A}	14708 (10.66)	14645 (10.18)	14167 (6.60)	15074 (13.42)	20456 (53.91)	21626 (62.71)	21804 (64.05)	23102 (73.82)	20633 (55.24)	14935 (12.37)	21511 (61.85)	19213 (44.56)	18534 (39.45)	19932 (49.97)	15803 (18.90)
DF _{A+S/O}	14074 (22.70)	14167 (23.52)	13064 (13.90)	14707 (28.23)	20206 (76.17)	21579 (88.14)	21611 (88.42)	23033 (100.82)	20521 (78.92)	14338 (25.01)	21431 (86.85)	19156 (67.01)	18464 (60.98)	19818 (72.79)	15159 (32.16)
DF _{A+S/V}	15109 (6.61)	15074 (6.37)	14707 (3.78)	15077 (6.39)	20567 (45.12)	21624 (52.58)	21689 (53.04)	23047 (62.62)	20569 (45.14)	15338 (8.23)	21473 (51.52)	19215 (35.58)	18530 (30.75)	19893 (40.37)	16136 (13.85)
E _{1-ctx}	20505 (9.77)	20456 (9.51)	20206 (8.17)	20567 (10.10)	22355 (19.68)	23193 (24.16)	23278 (24.62)	24217 (29.65)	22808 (22.10)	20590 (10.23)	23267 (24.56)	22031 (17.94)	21698 (16.16)	22235 (19.03)	20570 (10.12)
E _{2-ctx}	21698 (2.30)	21626 (1.96)	21579 (1.74)	21624 (1.95)	23193 (9.35)	23279 (9.75)	23614 (11.34)	24337 (14.74)	23025 (8.56)	21743 (2.51)	23442 (10.52)	22404 (5.63)	22135 (4.36)	22596 (6.54)	21663 (2.14)
E _{3-ctx}	21865 (3.37)	21804 (3.08)	21611 (2.16)	21689 (2.53)	23278 (10.05)	23614 (11.63)	23353 (10.40)	24448 (15.58)	23303 (10.16)	21937 (3.70)	23686 (11.97)	22729 (7.45)	22479 (6.27)	22789 (7.73)	21788 (3.00)
E _{DDG}	23147 (2.02)	23102 (1.82)	23033 (1.51)	23047 (1.57)	24217 (6.73)	24337 (7.26)	24448 (7.75)	24837 (9.46)	24055 (6.02)	23207 (2.28)	24369 (7.40)	23623 (4.11)	23422 (3.22)	23700 (4.45)	23091 (1.77)
E _{LTO}	20691 (2.81)	20633 (2.52)	20521 (1.97)	20569 (2.20)	22808 (13.33)	23025 (14.40)	23303 (15.79)	24055 (19.52)	21882 (8.73)	20748 (3.09)	22634 (12.46)	21561 (7.13)	21153 (5.10)	21927 (8.95)	20638 (2.54)
E _{MA}	14740 (9.70)	14935 (11.15)	14338 (6.70)	15338 (14.15)	20590 (53.23)	21743 (61.82)	21937 (63.26)	23207 (72.71)	20748 (54.41)	14898 (10.88)	21626 (60.94)	19343 (43.95)	18660 (38.87)	20065 (49.33)	15945 (18.67)
E _{SANCOV}	21575 (2.60)	21511 (2.30)	21431 (1.91)	21473 (2.12)	23267 (10.65)	23442 (11.48)	23686 (12.64)	24369 (15.89)	22634 (7.64)	21626 (2.84)	23013 (9.44)	22298 (6.04)	21985 (4.55)	22536 (7.17)	21537 (2.42)
E _{coarse}	19289 (2.73)	19213 (2.33)	19156 (2.02)	19215 (2.34)	22031 (17.34)	22404 (19.32)	22729 (21.06)	23623 (25.82)	21561 (14.83)	19343 (3.02)	22298 (18.76)	20499 (9.18)	20107 (7.09)	21140 (12.59)	19294 (2.76)
P _{2-gram}	18600 (2.88)	18534 (2.51)	18464 (2.13)	18530 (2.49)	21698 (20.01)	22135 (22.43)	22479 (24.33)	23422 (29.55)	21153 (16.99)	18660 (3.21)	21985 (21.60)	20107 (11.21)	19424 (7.44)	20713 (14.56)	18618 (2.97)
P _{4-gram}	19997 (3.17)	19932 (2.83)	19818 (2.24)	19893 (2.63)	22235 (14.71)	22596 (16.58)	22789 (17.57)	23700 (22.27)	21927 (13.12)	20065 (3.52)	22536 (16.27)	21140 (9.06)	20713 (6.86)	21271 (9.74)	19939 (2.87)
P _{8-gram}	15838 (13.66)	15803 (13.41)	15159 (8.78)	16136 (15.79)	20570 (47.62)	21663 (55.46)	21788 (56.36)	23091 (65.71)	20638 (48.11)	15945 (14.43)	21537 (54.56)	19294 (33.61)	18618 (43.09)	19939 (32.11)	15601 (11.95)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(B) *bloaty* (context-sensitive).

FIGURE 7.5: The predicted improvements to ICFG control-flow coverage. These improvements are made by pairing the fuzzer on the y -axis with the fuzzer on the x -axis. Each cell contains the expected number of covered edges and the percentage increase over only using the fuzzer on the y -axis (in brackets). The highest expected coverage is **highlighted**, while the cell color reflects the relative performance of the fuzzer: red is worst performance, blue is best performance.

B	10315 (10.13)	10597 (13.14)	9698 (3.55)	10720 (14.45)	14875 (58.81)	14921 (59.31)	15010 (60.25)	14946 (59.57)	14858 (58.63)	12074 (28.91)	14848 (58.53)	14690 (56.84)	14703 (56.98)	15041 (60.59)	14302 (52.70)
DF _{A+S/A}	10597 (8.81)	10692 (9.78)	9996 (2.64)	10875 (11.66)	14875 (52.73)	14922 (53.21)	15010 (54.12)	14946 (53.46)	14858 (52.56)	12086 (24.09)	14850 (52.47)	14691 (50.84)	14704 (50.97)	15042 (54.44)	14304 (46.87)
DF _{A+S/O}	9698 (37.65)	9996 (41.88)	7944 (12.76)	10250 (45.48)	14874 (111.11)	14921 (111.78)	15010 (113.04)	14946 (112.13)	14857 (110.88)	12072 (71.35)	14848 (110.74)	14690 (108.50)	14703 (108.68)	15041 (113.48)	14296 (102.91)
DF _{A+S/V}	10720 (6.67)	10875 (8.21)	10250 (2.00)	10927 (8.73)	14875 (8.73)	14923 (48.02)	15010 (49.37)	14946 (48.73)	14860 (47.87)	12092 (20.32)	14852 (47.79)	14692 (46.20)	14704 (46.32)	15042 (49.68)	14311 (42.41)
E _{1-ctx}	14875 (0.00)	14875 (0.01)	14874 (0.00)	14875 (0.01)	15086 (1.42)	15149 (1.85)	15188 (2.11)	15164 (1.95)	15210 (2.26)	14885 (0.07)	15195 (2.16)	15064 (1.27)	15074 (1.35)	15200 (2.19)	15004 (0.87)
E _{2-ctx}	14921 (0.00)	14922 (0.01)	14921 (0.00)	14923 (0.02)	15149 (1.53)	15138 (1.45)	15212 (1.95)	15184 (1.77)	15225 (2.04)	14933 (0.08)	15209 (1.93)	15092 (1.15)	15098 (1.19)	15217 (1.99)	15028 (0.72)
E _{3-ctx}	15010 (0.00)	15010 (0.00)	15010 (0.00)	15010 (0.00)	15188 (1.19)	15212 (1.35)	15213 (1.35)	15227 (1.45)	15272 (1.75)	15015 (0.03)	15259 (1.66)	15129 (0.80)	15147 (0.91)	15249 (1.59)	15091 (0.54)
E _{DDG}	14946 (0.00)	14946 (0.00)	14946 (0.00)	14946 (0.01)	15164 (1.46)	15184 (1.60)	15227 (1.88)	15156 (1.41)	15238 (1.96)	14951 (0.04)	15225 (1.87)	15103 (1.05)	15100 (1.04)	15222 (1.85)	15047 (0.68)
E _{LTO}	14858 (0.00)	14858 (0.01)	14857 (0.00)	14860 (0.02)	15210 (2.37)	15225 (2.47)	15272 (2.79)	15238 (2.56)	15140 (1.90)	14871 (0.09)	15161 (2.04)	15122 (1.78)	15133 (1.86)	15275 (2.81)	15004 (0.98)
EMA	12074 (0.05)	12086 (0.15)	12072 (0.04)	12092 (0.20)	14885 (23.34)	14933 (23.74)	15015 (24.42)	14951 (23.90)	14871 (23.23)	12287 (1.82)	14865 (23.18)	14705 (21.86)	14717 (21.95)	15047 (24.69)	14351 (18.92)
E _{SANCOV}	14848 (0.00)	14850 (0.01)	14848 (0.00)	14852 (0.02)	15195 (2.34)	15209 (2.43)	15259 (2.76)	15225 (2.54)	15161 (2.11)	14865 (0.11)	15123 (1.85)	15107 (1.74)	15118 (1.82)	15261 (2.78)	14992 (0.97)
E _{coarse}	14690 (0.00)	14691 (0.01)	14690 (0.00)	14692 (0.01)	15064 (2.55)	15092 (2.74)	15129 (2.99)	15103 (2.81)	15122 (2.94)	14705 (0.10)	15107 (2.84)	14932 (1.65)	14980 (1.97)	15143 (3.08)	14873 (1.25)
P _{2-gram}	14703 (0.00)	14704 (0.01)	14703 (0.00)	14704 (0.01)	15074 (2.53)	15098 (2.69)	15147 (3.02)	15100 (2.70)	15133 (2.93)	14717 (0.10)	15118 (2.83)	14980 (1.89)	14962 (1.76)	15150 (3.04)	14884 (1.23)
P _{4-gram}	15041 (0.00)	15042 (0.00)	15041 (0.01)	15042 (0.01)	15200 (1.05)	15217 (1.17)	15249 (1.38)	15222 (1.56)	15275 (1.56)	15047 (0.04)	15261 (1.46)	15143 (0.68)	15150 (0.72)	15221 (1.19)	15097 (0.37)
P _{8-gram}	14302 (0.05)	14304 (0.06)	14296 (0.01)	14311 (0.11)	15004 (4.96)	15028 (5.13)	15091 (5.56)	15047 (5.26)	15004 (4.95)	14351 (0.39)	14992 (4.87)	14873 (4.04)	14884 (4.12)	15097 (5.61)	14570 (1.92)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	EMA	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(C) *curl* (context-insensitive).

B	23243 (14.96)	24113 (19.26)	22050 (9.05)	24542 (21.38)	36973 (82.86)	37534 (85.64)	36858 (82.29)	38664 (91.22)	35915 (77.63)	27913 (38.05)	36716 (81.59)	37187 (83.92)	37742 (86.67)	37024 (83.11)	35136 (73.77)
DF _{A+S/A}	24113 (12.89)	24477 (14.60)	22942 (7.41)	25045 (17.26)	37129 (73.83)	37700 (76.50)	37039 (73.41)	38808 (81.69)	36067 (68.86)	28071 (31.42)	36872 (72.63)	37316 (74.71)	37902 (77.45)	37181 (74.08)	35292 (65.23)
DF _{A+S/O}	22050 (40.27)	22942 (45.94)	18307 (16.46)	23683 (50.66)	37240 (136.90)	37752 (140.16)	37123 (136.16)	38934 (147.68)	36179 (130.15)	28114 (78.84)	36995 (135.34)	37465 (138.33)	38018 (141.85)	37277 (137.14)	35402 (125.21)
DF _{A+S/V}	24542 (10.74)	25045 (13.01)	23683 (6.86)	25165 (13.55)	37057 (67.21)	37655 (69.91)	36997 (66.94)	38763 (74.91)	36019 (62.53)	28041 (26.53)	36802 (66.06)	37260 (68.13)	37828 (70.69)	37157 (67.66)	35212 (58.88)
E _{1-ctx}	36973 (2.56)	37129 (3.30)	37240 (3.30)	37057 (2.80)	38734 (7.45)	39487 (8.14)	38982 (8.14)	40422 (12.13)	38625 (7.15)	37167 (3.10)	39152 (8.61)	39506 (9.59)	39816 (10.45)	39079 (8.40)	38536 (6.90)
E _{2-ctx}	37534 (2.73)	37700 (3.18)	37752 (3.33)	37655 (3.06)	39487 (8.08)	39597 (8.38)	39414 (7.87)	40682 (11.35)	39081 (6.97)	37807 (3.48)	39549 (8.25)	39846 (9.06)	40039 (9.59)	39480 (8.06)	38990 (6.71)
E _{3-ctx}	36858 (2.74)	37039 (3.24)	37123 (3.48)	36997 (3.12)	38982 (8.66)	39414 (9.86)	38665 (7.77)	40402 (12.61)	38529 (7.40)	37187 (3.65)	39143 (9.11)	39432 (9.91)	39811 (10.97)	38917 (8.47)	38527 (7.39)
E _{DDG}	38664 (2.33)	38808 (2.71)	38934 (3.04)	38763 (2.59)	40422 (6.98)	40682 (7.67)	40402 (6.93)	41086 (8.74)	40012 (5.90)	38869 (2.87)	40445 (7.04)	40643 (7.56)	40827 (8.05)	40488 (7.16)	39890 (5.57)
E _{LTO}	35915 (3.02)	36067 (3.46)	36179 (3.78)	36019 (3.32)	38625 (10.80)	39081 (12.11)	38529 (10.52)	40012 (14.78)	37457 (7.44)	36154 (3.71)	38384 (10.10)	38863 (11.48)	39340 (12.85)	38591 (10.70)	37645 (7.99)
EMA	27913 (3.92)	28071 (4.51)	28114 (4.67)	28041 (4.40)	37167 (38.38)	37807 (40.76)	37187 (38.45)	38869 (44.71)	36154 (34.61)	28536 (6.25)	36944 (37.55)	37378 (39.17)	37911 (41.15)	37281 (38.80)	35425 (31.89)
E _{SANCOV}	36716 (2.62)	36872 (3.06)	36995 (3.40)	36802 (2.86)	39152 (9.43)	39549 (10.54)	39143 (9.41)	40445 (13.04)	38384 (7.28)	36944 (3.26)	38734 (8.26)	39397 (10.12)	39782 (11.19)	39167 (9.47)	38241 (6.89)
E _{coarse}	37187 (2.43)	37316 (2.79)	37465 (3.20)	37260 (2.63)	39506 (8.82)	39846 (9.76)	39432 (8.62)	40643 (11.95)	38863 (7.05)	37378 (2.96)	39397 (8.52)	39389 (8.50)	40001 (10.18)	39475 (8.73)	38600 (6.32)
P _{2-gram}	37742 (2.61)	37902 (3.04)	38018 (3.36)	37828 (2.84)	39816 (8.25)	40039 (8.85)	39811 (8.23)	40827 (10.99)	39340 (6.95)	37911 (3.07)	39782 (8.15)	40001 (8.75)	40002 (8.75)	39852 (8.34)	39118 (6.35)
P _{4-gram}	37024 (2.65)	37181 (3.08)	37277 (3.35)	37157 (3.02)	39079 (8.35)	39480 (9.46)	38917 (7.90)	40488 (12.25)	38591 (6.99)	37281 (3.36)	39167 (8.59)	39475 (9.44)	39852 (10.49)	38707 (7.32)	38628 (7.10)
P _{8-gram}	35136 (2.62)	35292 (3.08)	35402 (3.40)	35212 (2.84)	38536 (12.55)	38990 (13.88)	38527 (12.53)	39890 (16.51)	37645 (9.95)	35425 (3.47)	38241 (11.69)	38600 (12.74)	39118 (14.25)	38628 (12.82)	36874 (7.70)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	EMA	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(D) *curl* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	6471 (2.77)	6587 (4.62)	6391 (1.51)	6778 (7.65)	10051 (59.63)	9982 (58.54)	10063 (59.83)	10633 (68.89)	12090 (92.03)	6758 (7.33)	12163 (93.19)	10097 (60.37)	10010 (58.98)	10203 (62.06)	9496 (50.82)
DF _{A+S/A}	6587 (1.72)	6632 (2.43)	6523 (0.73)	6805 (5.10)	10054 (55.27)	9984 (54.18)	10065 (55.44)	10634 (64.22)	12091 (86.72)	6835 (5.55)	12164 (87.85)	10099 (55.96)	10011 (54.61)	10205 (57.60)	9501 (46.72)
DF _{A+S/O}	6391 (4.91)	6523 (7.07)	6252 (2.62)	6745 (10.72)	10029 (64.61)	9961 (63.50)	10041 (64.82)	10633 (74.54)	12090 (98.45)	6691 (9.83)	12163 (99.65)	10076 (65.38)	9990 (63.97)	10183 (67.14)	9472 (55.48)
DF _{A+S/V}	6778 (0.84)	6805 (1.25)	6745 (0.36)	6881 (2.37)	10104 (50.32)	10035 (49.30)	10118 (50.53)	10633 (58.20)	12089 (79.86)	6945 (3.33)	12163 (80.96)	10149 (51.00)	10061 (49.69)	10254 (52.55)	9562 (42.27)
E _{1-ctx}	10051 (0.59)	10054 (0.63)	10029 (0.37)	10104 (1.12)	10148 (1.57)	10174 (1.83)	10199 (2.08)	10712 (7.22)	12211 (22.22)	10124 (1.33)	12283 (22.94)	10249 (2.58)	10218 (2.27)	10248 (2.57)	10059 (0.67)
E _{2-ctx}	9982 (0.60)	9984 (0.62)	9961 (0.39)	10035 (1.14)	10174 (2.54)	10103 (1.83)	10173 (2.53)	10710 (7.94)	12210 (23.06)	10056 (1.35)	12276 (23.73)	10204 (2.84)	10161 (2.41)	10251 (3.31)	9997 (0.75)
E _{3-ctx}	10063 (0.55)	10065 (0.57)	10041 (0.33)	10118 (1.09)	10199 (1.91)	10173 (1.64)	10177 (1.69)	10727 (7.18)	12232 (22.22)	10138 (1.29)	12283 (22.73)	10240 (2.32)	10212 (2.03)	10259 (2.51)	10072 (0.63)
E _{DDG}	10633 (0.02)	10634 (0.03)	10633 (0.02)	10633 (0.02)	10712 (0.76)	10710 (0.74)	10727 (0.90)	10800 (1.59)	12364 (16.30)	10634 (0.03)	12412 (16.75)	10725 (0.88)	10704 (0.68)	10740 (1.02)	10652 (0.20)
E _{LTO}	12090 (0.02)	12091 (0.03)	12090 (0.02)	12089 (0.01)	12211 (1.03)	12210 (1.01)	12232 (1.20)	12364 (2.29)	12486 (3.30)	12091 (0.03)	12546 (3.79)	12244 (1.29)	12222 (1.11)	12253 (1.37)	12139 (0.43)
E _{MA}	6758 (3.38)	6835 (4.55)	6691 (2.36)	6945 (6.25)	10124 (54.88)	10056 (53.83)	10138 (55.08)	10634 (62.68)	12091 (84.97)	6806 (4.12)	12165 (86.10)	10167 (55.53)	10079 (54.18)	10274 (57.16)	9578 (46.52)
E _{SANCOV}	12163 (0.02)	12164 (0.02)	12163 (0.02)	12163 (0.02)	12283 (1.01)	12276 (0.95)	12283 (1.00)	12412 (2.07)	12546 (3.16)	12165 (0.03)	12527 (3.01)	12288 (1.04)	12273 (0.92)	12307 (1.20)	12209 (0.40)
E _{coarse}	10097 (0.49)	10099 (0.51)	10076 (0.28)	10149 (1.01)	10249 (2.00)	10204 (1.56)	10240 (1.91)	10725 (6.73)	12244 (21.85)	10167 (1.18)	12288 (22.29)	10223 (1.74)	10206 (1.58)	10289 (2.40)	10108 (0.60)
P _{2-gram}	10010 (0.38)	10011 (0.39)	9990 (0.17)	10061 (0.89)	10218 (2.47)	10161 (1.89)	10212 (2.40)	10704 (7.34)	12222 (22.56)	10079 (1.07)	12273 (23.07)	10206 (2.35)	10133 (1.61)	10268 (2.96)	10041 (0.69)
P _{4-gram}	10203 (0.52)	10205 (0.54)	10183 (0.32)	10254 (1.02)	10248 (0.96)	10251 (0.99)	10259 (1.07)	10740 (5.81)	12253 (20.72)	10274 (1.21)	12307 (21.25)	10289 (1.36)	10268 (1.16)	10270 (1.18)	10181 (0.30)
P _{8-gram}	9496 (0.70)	9501 (0.75)	9472 (0.44)	9562 (1.40)	10059 (6.66)	9997 (6.01)	10072 (6.80)	10652 (12.96)	12139 (28.72)	9578 (1.57)	12209 (29.47)	10108 (7.19)	10041 (6.48)	10181 (7.96)	9668 (2.53)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(E) freetype2 (context-insensitive).

B	15378 (8.28)	16148 (13.70)	15420 (8.58)	16977 (19.54)	29419 (107.15)	29046 (104.52)	30120 (112.09)	30906 (117.62)	51366 (261.69)	16206 (14.11)	58674 (313.14)	28941 (103.78)	27834 (95.99)	28978 (104.05)	25853 (82.04)
DF _{A+S/A}	16148 (7.01)	16504 (9.37)	16185 (7.26)	17309 (14.70)	29577 (96.00)	29192 (93.45)	30249 (100.46)	31068 (105.88)	51376 (240.46)	16830 (11.53)	58710 (289.06)	29119 (92.97)	28123 (86.37)	29187 (93.42)	26117 (73.07)
DF _{A+S/O}	15420 (10.40)	16185 (15.87)	15089 (8.03)	17076 (22.25)	29618 (112.04)	29217 (109.17)	30261 (116.65)	31118 (122.78)	51437 (268.25)	16160 (15.69)	58784 (320.85)	29154 (108.72)	28081 (101.04)	29173 (108.86)	26084 (86.74)
DF _{A+S/V}	16977 (5.20)	17309 (7.25)	17076 (5.81)	17648 (9.36)	29713 (84.11)	29297 (81.54)	30377 (88.23)	31153 (93.03)	51397 (218.48)	17487 (8.36)	58709 (263.79)	29372 (82.00)	28493 (76.56)	29403 (82.19)	26490 (64.14)
E _{1-ctx}	29419 (2.48)	29577 (3.03)	29618 (3.17)	29713 (3.50)	30605 (6.61)	30494 (6.22)	31415 (6.22)	32102 (9.43)	51800 (80.44)	29604 (3.12)	59090 (105.83)	30868 (7.53)	30440 (6.03)	30724 (7.02)	29661 (3.32)
E _{2-ctx}	29046 (2.61)	29192 (3.12)	29217 (3.21)	29297 (3.49)	30494 (7.72)	29955 (5.82)	31146 (10.02)	31919 (12.75)	51774 (82.89)	29224 (3.24)	59050 (108.60)	30629 (8.20)	30228 (6.78)	30505 (7.76)	29405 (3.87)
E _{3-ctx}	30120 (2.38)	30249 (2.82)	30261 (2.86)	30377 (3.25)	31415 (6.78)	31146 (5.87)	31759 (7.95)	32745 (11.30)	51951 (76.59)	30323 (3.07)	59183 (101.17)	31467 (6.96)	31065 (5.59)	31350 (6.56)	30400 (3.33)
E _{DDG}	30906 (2.07)	31068 (2.60)	31118 (2.76)	31153 (2.88)	32102 (6.01)	31919 (5.41)	32745 (8.14)	32513 (7.37)	52145 (72.21)	30975 (2.29)	59378 (96.09)	32071 (5.91)	31604 (4.37)	31982 (5.62)	31185 (2.99)
E _{LTO}	51366 (0.68)	51376 (0.70)	51437 (0.82)	51397 (0.74)	51800 (1.53)	51774 (1.48)	51951 (1.83)	52145 (2.21)	55944 (9.65)	51396 (0.74)	62887 (23.26)	51754 (1.44)	51596 (1.13)	51884 (1.70)	51457 (0.86)
E _{MA}	16206 (10.30)	16830 (14.55)	16160 (9.99)	17487 (19.02)	29604 (101.49)	29224 (98.91)	30323 (106.39)	30975 (110.82)	51396 (249.81)	16326 (11.12)	58709 (299.59)	29202 (98.75)	28111 (91.33)	29206 (98.78)	26169 (78.11)
E _{SANCOV}	58674 (0.59)	58710 (0.65)	58784 (0.78)	58709 (0.65)	59090 (1.30)	59050 (1.23)	59183 (1.46)	59378 (1.80)	62887 (7.81)	58709 (0.65)	66438 (13.90)	59002 (1.15)	58820 (0.84)	59110 (1.34)	58772 (0.76)
E _{coarse}	28941 (2.79)	29119 (3.42)	29154 (3.55)	29372 (4.32)	30868 (9.63)	30629 (8.78)	31467 (11.76)	32071 (13.90)	51754 (83.81)	29202 (3.71)	59002 (109.55)	30456 (8.17)	29948 (6.36)	30584 (8.62)	29282 (4.00)
P _{2-gram}	27834 (3.43)	28123 (4.51)	28081 (4.35)	28493 (5.88)	30440 (13.12)	30228 (12.33)	31065 (15.44)	31604 (17.44)	51596 (91.73)	28111 (4.46)	58820 (118.58)	29948 (11.29)	28834 (7.15)	29920 (11.18)	28292 (5.14)
P _{4-gram}	28978 (2.80)	29187 (3.54)	29173 (3.49)	29403 (4.31)	30724 (8.99)	30505 (8.21)	31350 (11.21)	31982 (13.45)	51884 (84.06)	29206 (3.61)	59110 (109.69)	30584 (8.50)	29920 (6.14)	30213 (7.18)	29229 (3.69)
P _{8-gram}	25853 (3.57)	26117 (4.62)	26084 (4.49)	26490 (6.12)	29661 (18.82)	29405 (17.79)	30400 (21.78)	31185 (24.92)	51457 (106.13)	26169 (4.83)	58772 (135.43)	29282 (17.30)	28292 (13.34)	29229 (17.09)	26556 (6.38)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(F) freetype2 (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	21934 (16.10)	23244 (23.04)	21428 (13.42)	21169 (12.05)	37250 (97.17)	39149 (107.23)	39899 (111.20)	42781 (126.45)	43449 (129.99)	28488 (50.80)	43538 (130.46)	41832 (121.42)	41913 (121.85)	41262 (118.41)	35526 (88.05)
DF _{A+S/A}	23244 (11.71)	23875 (14.74)	22710 (9.14)	22407 (7.69)	37304 (79.28)	39282 (88.78)	39911 (91.81)	42783 (105.61)	43450 (108.82)	28812 (38.47)	43539 (109.24)	41833 (101.04)	41918 (101.46)	41270 (98.34)	35620 (71.19)
DF _{A+S/O}	21428 (22.75)	22710 (30.09)	20409 (16.91)	20321 (16.41)	37225 (113.24)	39063 (123.77)	39897 (128.55)	42779 (145.06)	43448 (148.89)	28299 (62.11)	43537 (149.40)	41826 (139.60)	41909 (140.07)	41252 (136.31)	35471 (103.20)
DF _{A+S/V}	21169 (23.07)	22407 (30.27)	20321 (18.14)	19724 (14.67)	37200 (116.27)	39019 (126.85)	39889 (131.90)	42779 (148.70)	43448 (152.59)	28205 (63.98)	43537 (153.11)	41824 (143.16)	41908 (143.64)	41249 (139.81)	35412 (105.88)
E _{1-ctx}	37250 (0.27)	37304 (0.41)	37225 (0.20)	37200 (0.13)	38872 (4.63)	40969 (10.28)	40551 (9.15)	42973 (15.67)	43550 (17.22)	37735 (1.57)	43659 (17.52)	42184 (13.55)	42292 (13.84)	41663 (12.14)	38789 (4.41)
E _{2-ctx}	39149 (1.05)	39282 (1.40)	39063 (0.83)	39019 (0.72)	40969 (5.75)	41702 (7.64)	41708 (7.66)	43281 (11.72)	43762 (12.96)	39827 (2.80)	43842 (13.17)	42724 (10.28)	42797 (10.47)	42371 (9.37)	40729 (5.13)
E _{3-ctx}	39899 (0.07)	39911 (0.10)	39897 (0.06)	39889 (0.04)	40551 (1.70)	41708 (4.60)	41166 (3.24)	43196 (8.33)	43705 (9.61)	40060 (0.47)	43794 (9.83)	42558 (6.74)	42666 (7.01)	42151 (5.71)	40512 (1.60)
E _{DDG}	42781 (0.01)	42783 (0.02)	42779 (0.00)	42779 (0.00)	42973 (0.46)	43281 (1.18)	43196 (0.98)	43697 (2.15)	44105 (3.11)	42815 (0.09)	44150 (3.21)	43523 (1.74)	43536 (1.77)	43425 (1.52)	42924 (0.34)
E _{LTO}	43449 (0.01)	43450 (0.01)	43448 (0.00)	43448 (0.00)	43550 (0.24)	43762 (0.73)	43705 (0.60)	44105 (1.52)	44177 (1.68)	43467 (0.05)	44301 (1.97)	43887 (1.01)	43902 (1.05)	43820 (0.86)	43526 (0.18)
E _{MA}	28488 (2.54)	28812 (3.71)	28299 (1.86)	28205 (1.52)	37735 (35.83)	39827 (43.36)	40060 (44.19)	42815 (54.11)	43467 (56.46)	30289 (9.03)	43555 (56.77)	41899 (50.81)	41975 (51.09)	41348 (48.83)	36323 (30.75)
E _{SANCOV}	43538 (0.00)	43539 (0.01)	43537 (0.00)	43537 (0.00)	43659 (0.28)	43842 (0.70)	43794 (0.59)	44150 (1.41)	44301 (1.76)	43555 (0.04)	44262 (1.67)	43963 (0.98)	43956 (0.96)	43904 (0.85)	43613 (0.18)
E _{coarse}	41832 (0.03)	41833 (0.03)	41826 (0.02)	41824 (0.01)	42184 (0.87)	42724 (2.17)	42558 (1.77)	43523 (4.08)	43887 (4.95)	41899 (0.19)	43963 (5.13)	42943 (2.69)	43106 (3.08)	42915 (2.62)	42100 (0.67)
P _{2-gram}	41913 (0.02)	41918 (0.04)	41909 (0.01)	41908 (0.01)	42292 (0.93)	42797 (2.13)	42666 (1.82)	43536 (3.90)	43902 (4.77)	41976 (0.17)	43956 (4.90)	43106 (2.87)	42999 (2.62)	42980 (2.57)	42185 (0.67)
P _{4-gram}	41262 (0.05)	41270 (0.07)	41252 (0.02)	41249 (0.02)	41663 (1.02)	42371 (2.74)	42151 (2.20)	43425 (6.25)	43820 (6.25)	41348 (0.26)	43904 (6.45)	42915 (4.06)	42980 (4.21)	42459 (2.95)	41564 (0.78)
P _{8-gram}	35526 (0.56)	35620 (0.83)	35471 (0.41)	35412 (0.24)	38789 (9.80)	40729 (15.29)	40512 (14.68)	42924 (21.50)	43526 (23.21)	36323 (2.82)	43613 (23.46)	42100 (19.17)	42185 (19.41)	41564 (17.66)	37643 (6.56)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(G) *harfbuzz* (context-insensitive).

B	66960 (24.24)	71694 (33.02)	64673 (19.99)	64584 (19.83)	125413 (132.69)	129318 (139.93)	133267 (147.26)	143992 (167.16)	145179 (169.36)	89860 (66.72)	146937 (172.62)	139813 (159.40)	140572 (160.81)	137390 (154.91)	114459 (112.36)
DF _{A+S/A}	71694 (17.61)	73626 (20.78)	69508 (14.03)	69063 (13.30)	126288 (107.17)	130490 (114.07)	134212 (120.18)	144761 (137.48)	146042 (139.58)	91296 (49.77)	147752 (142.39)	140681 (130.79)	141397 (131.96)	138217 (126.74)	115396 (89.31)
DF _{A+S/O}	64673 (35.21)	69508 (45.31)	60170 (25.79)	60697 (26.89)	125436 (162.24)	129020 (169.73)	133324 (178.73)	143916 (200.87)	145157 (203.47)	89031 (86.13)	146842 (206.99)	139783 (192.23)	140454 (193.63)	137369 (187.18)	114431 (139.23)
DF _{A+S/V}	64584 (34.16)	69063 (43.46)	60697 (26.09)	59334 (23.26)	126086 (161.92)	129527 (169.07)	134005 (178.37)	144471 (200.11)	145794 (202.86)	89211 (85.32)	147391 (206.17)	140438 (191.73)	140975 (192.85)	137937 (186.54)	115008 (138.91)
E _{1-ctx}	125413 (2.57)	126288 (3.28)	125436 (2.58)	126086 (3.12)	131537 (7.57)	138409 (13.19)	137346 (12.33)	147088 (20.29)	148021 (21.06)	129581 (5.98)	149498 (22.26)	143257 (17.16)	145245 (18.79)	141449 (15.68)	130883 (7.04)
E _{2-ctx}	129318 (3.36)	130490 (4.30)	129020 (3.12)	129527 (3.53)	138409 (10.63)	140178 (12.04)	140784 (12.52)	147711 (18.06)	148464 (18.66)	134434 (7.45)	149966 (19.86)	144798 (15.73)	146470 (17.07)	143380 (14.60)	136343 (8.98)
E _{3-ctx}	133267 (2.31)	134212 (3.04)	133324 (2.36)	134005 (2.88)	137346 (5.44)	140784 (8.08)	138700 (6.48)	147541 (13.27)	148218 (13.79)	136896 (5.10)	149749 (14.97)	144198 (10.71)	146324 (12.34)	143020 (9.80)	136658 (4.92)
E _{DDG}	143992 (1.37)	144761 (1.92)	143916 (1.32)	144471 (1.71)	147088 (3.53)	147711 (3.99)	147541 (3.87)	149210 (5.05)	150565 (6.00)	146345 (3.03)	151922 (6.96)	148507 (4.55)	149833 (5.49)	148227 (4.36)	146082 (2.85)
E _{LTO}	145179 (1.53)	146042 (2.14)	145157 (1.52)	145794 (1.96)	148021 (3.52)	148464 (3.83)	148218 (3.66)	150565 (5.30)	149856 (4.80)	147900 (3.44)	151897 (6.23)	148965 (4.18)	150626 (5.34)	149031 (4.23)	147131 (2.90)
E _{MA}	89860 (5.35)	91296 (7.03)	89031 (4.37)	89211 (4.59)	129581 (51.91)	134434 (57.60)	136896 (60.49)	146345 (71.57)	147900 (73.39)	97548 (14.36)	149447 (75.20)	142774 (67.38)	143192 (67.87)	140345 (64.53)	119626 (40.24)
E _{SANCOV}	146937 (1.39)	147752 (1.95)	146842 (1.33)	147391 (1.71)	149498 (3.16)	149966 (3.48)	149749 (3.33)	151922 (4.83)	151897 (4.81)	149447 (3.12)	152307 (5.10)	150507 (3.86)	151954 (4.85)	150612 (3.95)	149049 (2.85)
E _{coarse}	139813 (1.73)	140681 (2.37)	139783 (1.71)	140438 (2.19)	143257 (4.24)	144798 (5.36)	144198 (4.93)	148507 (8.06)	148965 (8.39)	142774 (3.89)	150507 (9.52)	145399 (5.80)	147889 (7.61)	145916 (6.18)	142252 (3.51)
P _{2-gram}	140572 (1.54)	141397 (2.14)	140454 (1.45)	140975 (1.83)	145245 (4.92)	146470 (5.80)	146324 (5.70)	149833 (8.23)	150626 (8.80)	143192 (3.43)	151954 (9.76)	147889 (6.83)	147777 (6.74)	147131 (6.28)	143360 (3.55)
P _{4-gram}	137390 (1.92)	138217 (2.53)	137369 (1.90)	137937 (2.32)	141449 (4.93)	143380 (6.36)	143020 (6.09)	148227 (9.95)	149031 (10.55)	140345 (4.11)	150612 (11.72)	145916 (8.24)	147131 (9.14)	143845 (6.70)	140026 (3.87)
P _{8-gram}	114459 (3.27)	115396 (4.12)	114431 (3.24)	115008 (3.77)	130883 (18.09)	136343 (23.01)	136658 (23.30)	146082 (31.80)	147131 (32.75)	119626 (7.93)	149049 (34.48)	142252 (28.35)	143360 (29.35)	140026 (26.34)	122241 (10.29)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(H) *harfbuzz* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	670 (3.24)	676 (4.10)	667 (2.69)	702 (8.12)	1036 (59.58)	1036 (59.57)	1037 (59.69)	998 (53.65)	1037 (59.75)	743 (14.50)	1037 (59.80)	1033 (59.10)	1035 (59.49)	1038 (59.81)	975 (50.23)
DF _{A+S/A}	676 (1.51)	678 (1.82)	673 (1.11)	704 (5.70)	1036 (55.60)	1036 (55.59)	1037 (55.71)	999 (50.07)	1037 (55.77)	744 (11.70)	1037 (55.81)	1033 (55.14)	1035 (55.51)	1038 (55.83)	977 (46.73)
DF _{A+S/O}	667 (7.11)	673 (8.16)	656 (5.48)	700 (12.49)	1036 (66.45)	1036 (66.44)	1037 (66.56)	995 (59.84)	1037 (66.63)	742 (19.28)	1037 (66.68)	1033 (65.95)	1035 (66.36)	1038 (66.69)	974 (56.48)
DF _{A+S/V}	702 (0.79)	704 (1.06)	700 (0.53)	717 (2.97)	1036 (48.77)	1036 (48.75)	1037 (48.87)	1002 (43.92)	1037 (48.92)	748 (7.36)	1037 (48.97)	1033 (48.32)	1035 (48.25)	1038 (48.98)	987 (41.76)
E _{1-ctx}	1036 (0.00)	1036 (0.00)	1036 (0.00)	1036 (0.00)	1037 (0.07)	1037 (0.10)	1037 (0.10)	1038 (0.14)	1039 (0.31)	1036 (0.00)	1039 (0.30)	1037 (0.11)	1038 (0.17)	1038 (0.14)	1036 (0.02)
E _{2-ctx}	1036 (0.00)	1036 (0.00)	1036 (0.00)	1036 (0.00)	1037 (0.11)	1037 (0.09)	1037 (0.11)	1037 (0.13)	1039 (0.34)	1036 (0.00)	1039 (0.32)	1037 (0.10)	1038 (0.16)	1038 (0.16)	1037 (0.09)
E _{3-ctx}	1037 (0.00)	1037 (0.00)	1037 (0.00)	1037 (0.00)	1037 (0.04)	1037 (0.04)	1037 (0.04)	1038 (0.08)	1039 (0.27)	1037 (0.00)	1039 (0.25)	1037 (0.05)	1038 (0.11)	1038 (0.08)	1037 (0.03)
E _{DDG}	998 (6.96)	999 (7.14)	995 (6.67)	1002 (7.47)	1038 (11.25)	1037 (11.22)	1038 (11.26)	1027 (10.11)	1040 (11.46)	1007 (7.96)	1039 (11.43)	1037 (11.18)	1038 (11.25)	1038 (11.29)	1024 (9.84)
E _{LTO}	1037 (0.00)	1037 (0.00)	1037 (0.00)	1037 (0.00)	1039 (0.20)	1039 (0.23)	1039 (0.23)	1040 (0.23)	1039 (0.19)	1037 (0.00)	1039 (0.21)	1040 (0.23)	1040 (0.25)	1040 (0.25)	1038 (0.11)
E _{MA}	743 (0.18)	744 (0.23)	742 (0.05)	748 (0.76)	1036 (39.62)	1036 (39.61)	1037 (39.72)	1007 (35.69)	1037 (39.77)	758 (2.22)	1037 (39.81)	1033 (39.20)	1035 (39.54)	1038 (39.82)	998 (34.48)
E _{SANCOV}	1037 (0.00)	1037 (0.00)	1037 (0.00)	1037 (0.00)	1039 (0.16)	1039 (0.17)	1039 (0.18)	1039 (0.18)	1039 (0.18)	1037 (0.00)	1039 (0.12)	1039 (0.17)	1040 (0.22)	1040 (0.22)	1038 (0.06)
E _{coarse}	1033 (0.00)	1033 (0.00)	1033 (0.00)	1033 (0.00)	1037 (0.41)	1037 (0.39)	1037 (0.42)	1037 (0.39)	1040 (0.64)	1033 (0.00)	1039 (0.61)	1036 (0.34)	1038 (0.46)	1038 (0.46)	1036 (0.35)
P _{2-gram}	1035 (0.00)	1035 (0.00)	1035 (0.00)	1035 (0.00)	1038 (0.23)	1038 (0.21)	1038 (0.23)	1038 (0.21)	1040 (0.41)	1035 (0.00)	1040 (0.41)	1038 (0.21)	1038 (0.23)	1038 (0.26)	1037 (0.19)
P _{4-gram}	1038 (0.00)	1038 (0.00)	1038 (0.00)	1038 (0.00)	1038 (0.00)	1038 (0.00)	1038 (0.03)	1038 (0.03)	1040 (0.21)	1038 (0.00)	1040 (0.21)	1038 (0.01)	1038 (0.05)	1038 (0.02)	1038 (0.00)
P _{8-gram}	975 (7.29)	977 (7.47)	974 (7.15)	987 (8.61)	1036 (13.99)	1037 (14.07)	1037 (14.08)	1024 (12.70)	1038 (14.22)	998 (9.77)	1038 (14.19)	1036 (14.03)	1037 (14.12)	1038 (14.14)	1004 (10.41)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(I) *jsoncpp* (context-insensitive).

B	1588 (6.79)	1603 (7.79)	1575 (5.86)	1693 (13.80)	2683 (80.41)	2714 (82.45)	2704 (81.80)	2540 (70.76)	2718 (82.77)	1777 (19.46)	2694 (81.10)	2649 (78.11)	2666 (79.24)	2684 (80.46)	2519 (69.36)
DF _{A+S/A}	1603 (4.43)	1606 (4.57)	1589 (3.48)	1697 (10.53)	2681 (74.66)	2712 (76.61)	2705 (76.21)	2545 (65.78)	2716 (76.92)	1778 (15.82)	2691 (75.28)	2646 (72.37)	2665 (73.58)	2685 (74.86)	2522 (64.29)
DF _{A+S/O}	1575 (11.00)	1589 (12.00)	1540 (8.60)	1683 (18.64)	2677 (88.75)	2708 (90.91)	2702 (90.48)	2530 (78.38)	2712 (91.19)	1771 (24.85)	2687 (89.41)	2642 (86.25)	2661 (87.59)	2681 (89.03)	2510 (76.97)
DF _{A+S/V}	1693 (2.88)	1697 (3.14)	1683 (2.29)	1748 (6.23)	2701 (64.18)	2726 (65.69)	2719 (65.25)	2573 (56.36)	2729 (65.86)	1814 (10.28)	2707 (64.56)	2668 (62.15)	2688 (63.36)	2700 (64.09)	2564 (55.82)
E _{1-ctx}	2683 (1.35)	2681 (1.27)	2677 (1.12)	2701 (2.02)	2712 (2.44)	2743 (3.59)	2734 (3.26)	2702 (2.04)	2749 (3.81)	2678 (1.13)	2729 (3.08)	2700 (1.71)	2702 (2.06)	2723 (2.84)	2718 (2.67)
E _{2-ctx}	2714 (1.12)	2712 (1.03)	2708 (0.90)	2726 (1.58)	2743 (2.20)	2756 (2.69)	2756 (2.68)	2732 (1.81)	2766 (3.07)	2709 (0.95)	2750 (2.46)	2730 (1.71)	2732 (1.79)	2748 (2.41)	2740 (2.10)
E _{3-ctx}	2704 (1.31)	2705 (1.36)	2702 (1.23)	2719 (1.87)	2734 (2.44)	2756 (3.25)	2741 (2.71)	2720 (1.90)	2762 (3.50)	2698 (1.08)	2745 (2.86)	2720 (1.90)	2721 (1.96)	2736 (2.49)	2736 (2.51)
E _{DDG}	2540 (7.93)	2545 (8.16)	2530 (7.53)	2573 (9.32)	2702 (14.81)	2732 (16.11)	2720 (15.58)	2646 (12.43)	2740 (16.43)	2556 (8.61)	2718 (15.48)	2680 (13.91)	2686 (14.15)	2705 (14.95)	2674 (13.63)
E _{LTO}	2718 (0.80)	2716 (0.72)	2712 (0.56)	2729 (1.19)	2749 (1.92)	2766 (2.57)	2762 (2.43)	2740 (1.59)	2765 (2.52)	2716 (0.71)	2754 (2.13)	2735 (1.41)	2741 (1.65)	2754 (2.11)	2746 (1.83)
E _{MA}	1777 (3.35)	1778 (3.43)	1771 (3.01)	1814 (5.53)	2678 (55.75)	2709 (57.58)	2698 (56.91)	2556 (48.66)	2716 (57.98)	1812 (5.38)	2690 (56.46)	2643 (53.74)	2657 (54.56)	2678 (55.75)	2572 (49.59)
E _{SANCOV}	2694 (1.15)	2691 (1.06)	2687 (0.90)	2707 (1.67)	2729 (2.50)	2750 (3.26)	2745 (3.10)	2718 (2.05)	2754 (3.43)	2690 (1.02)	2732 (2.59)	2714 (1.92)	2718 (2.07)	2733 (2.65)	2727 (2.40)
E _{coarse}	2649 (1.34)	2646 (1.23)	2642 (1.06)	2668 (2.05)	2700 (3.29)	2730 (4.42)	2720 (4.04)	2680 (2.53)	2735 (4.61)	2643 (1.11)	2714 (3.82)	2677 (2.40)	2687 (2.79)	2706 (3.52)	2703 (3.41)
P _{2-gram}	2666 (1.84)	2665 (1.80)	2661 (1.65)	2688 (2.67)	2702 (3.23)	2732 (4.35)	2721 (3.96)	2686 (2.62)	2741 (4.72)	2657 (1.51)	2718 (3.83)	2687 (2.65)	2683 (2.49)	2709 (3.49)	2708 (3.43)
P _{4-gram}	2684 (1.28)	2685 (1.30)	2681 (1.17)	2700 (1.87)	2723 (2.75)	2748 (3.70)	2736 (3.22)	2705 (2.07)	2754 (3.91)	2678 (1.04)	2733 (3.13)	2706 (2.11)	2709 (2.22)	2722 (2.69)	2725 (2.81)
P _{8-gram}	2519 (7.60)	2522 (7.74)	2510 (7.23)	2564 (9.51)	2718 (16.12)	2740 (17.04)	2736 (16.87)	2674 (14.22)	2746 (17.31)	2572 (9.86)	2727 (16.47)	2703 (15.47)	2708 (15.65)	2725 (16.38)	2635 (12.57)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(J) *jsoncpp* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	625 (32.84)	615 (30.74)	516 (9.79)	698 (48.43)	1342 (185.29)	1362 (189.52)	1629 (246.32)	1468 (212.12)	1451 (208.46)	827 (75.73)	1302 (176.76)	1322 (180.95)	1413 (200.32)	1159 (146.36)	1195 (154.04)
DF _{A+S/A}	615 (33.00)	598 (29.22)	505 (9.13)	684 (47.82)	1342 (190.22)	1362 (194.53)	1629 (252.31)	1468 (217.52)	1451 (213.80)	820 (77.40)	1302 (181.55)	1322 (185.81)	1413 (205.51)	1159 (150.63)	1195 (158.43)
DF _{A+S/O}	516 (124.34)	505 (119.21)	319 (38.69)	654 (184.19)	1342 (482.97)	1362 (491.62)	1629 (607.69)	1468 (537.79)	1451 (530.32)	819 (255.86)	1302 (465.55)	1322 (474.11)	1413 (513.68)	1159 (403.43)	1195 (419.11)
DF _{A+S/V}	698 (8.59)	684 (6.30)	654 (1.74)	719 (11.79)	1342 (108.71)	1362 (111.80)	1629 (153.36)	1468 (128.34)	1451 (125.66)	826 (28.43)	1302 (102.47)	1322 (105.54)	1413 (119.70)	1159 (80.23)	1195 (85.85)
E _{1-ctx}	1342 (0.00)	1342 (0.00)	1342 (0.00)	1342 (0.00)	1487 (10.79)	1511 (12.57)	1727 (28.66)	1582 (17.90)	1576 (17.45)	1342 (0.00)	1485 (10.63)	1489 (10.93)	1572 (17.14)	1367 (1.84)	1438 (7.14)
E _{2-ctx}	1362 (0.00)	1362 (0.00)	1362 (0.00)	1362 (0.00)	1511 (10.93)	1526 (12.03)	1744 (28.09)	1605 (17.82)	1596 (17.22)	1362 (0.00)	1507 (10.65)	1510 (10.85)	1590 (16.72)	1388 (1.95)	1455 (6.86)
E _{3-ctx}	1629 (0.00)	1629 (0.00)	1629 (0.00)	1629 (0.00)	1727 (5.98)	1744 (7.08)	1876 (15.15)	1784 (9.53)	1782 (9.39)	1629 (0.00)	1721 (5.66)	1726 (5.95)	1780 (9.27)	1648 (1.16)	1695 (4.05)
E _{DDG}	1468 (0.00)	1468 (0.00)	1468 (0.00)	1468 (0.00)	1582 (9.29)	1605 (9.29)	1784 (21.54)	1655 (12.70)	1656 (12.77)	1468 (0.00)	1580 (7.64)	1582 (7.73)	1659 (13.01)	1481 (0.89)	1545 (5.25)
E _{LTO}	1451 (0.00)	1451 (0.00)	1451 (0.00)	1451 (0.00)	1576 (8.63)	1596 (10.02)	1782 (22.81)	1656 (14.11)	1650 (13.70)	1451 (0.00)	1571 (8.30)	1576 (8.61)	1650 (13.72)	1470 (1.32)	1533 (5.68)
E _{MA}	827 (0.91)	820 (0.13)	819 (0.00)	826 (0.81)	1342 (63.82)	1362 (66.25)	1629 (98.87)	1468 (79.23)	1451 (77.13)	860 (5.01)	1302 (58.94)	1322 (61.34)	1413 (72.46)	1159 (41.48)	1195 (45.89)
E _{SANCOV}	1302 (0.00)	1302 (0.00)	1302 (0.00)	1302 (0.00)	1485 (14.04)	1507 (15.75)	1721 (32.22)	1580 (21.39)	1571 (20.70)	1302 (0.01)	1472 (13.09)	1482 (13.81)	1562 (20.01)	1355 (4.09)	1421 (9.15)
E _{coarse}	1322 (0.00)	1322 (0.00)	1322 (0.00)	1322 (0.00)	1489 (12.64)	1510 (14.23)	1726 (30.60)	1582 (19.68)	1576 (19.24)	1322 (0.01)	1482 (12.11)	1485 (12.38)	1571 (18.85)	1365 (3.27)	1431 (8.27)
P _{2-gram}	1413 (0.00)	1413 (0.00)	1413 (0.00)	1413 (0.00)	1572 (11.28)	1590 (12.52)	1780 (26.01)	1659 (17.44)	1650 (16.80)	1413 (0.01)	1562 (10.60)	1571 (11.18)	1633 (15.58)	1454 (2.91)	1515 (7.22)
P _{4-gram}	1159 (0.00)	1159 (0.00)	1159 (0.00)	1159 (0.00)	1367 (17.93)	1388 (19.81)	1648 (42.21)	1481 (27.82)	1470 (26.86)	1159 (0.01)	1355 (16.94)	1365 (17.76)	1454 (25.45)	1202 (3.74)	1289 (11.25)
P _{8-gram}	1195 (0.00)	1195 (0.00)	1195 (0.00)	1195 (0.00)	1438 (20.32)	1455 (21.78)	1695 (41.84)	1545 (29.31)	1533 (28.32)	1195 (0.01)	1421 (18.91)	1431 (19.74)	1515 (26.75)	1289 (7.89)	1349 (12.87)

(K) *lcms* (context-insensitive).

B	1096 (36.59)	1073 (33.76)	891 (11.07)	1226 (52.71)	2939 (266.19)	2927 (264.71)	3751 (367.46)	3250 (305.00)	3240 (303.76)	1577 (96.55)	2872 (257.91)	2990 (272.57)	3148 (292.32)	2518 (213.73)	2552 (218.02)
DF _{A+S/A}	1073 (37.17)	1034 (32.16)	864 (10.39)	1193 (52.46)	2941 (275.89)	2931 (274.58)	3753 (379.57)	3255 (315.92)	3243 (314.42)	1562 (99.64)	2876 (267.53)	2995 (282.71)	3150 (302.60)	2520 (222.08)	2556 (226.69)
DF _{A+S/O}	891 (149.38)	864 (141.68)	525 (46.79)	1135 (217.70)	2933 (720.64)	2920 (717.09)	3746 (948.21)	3243 (807.40)	3235 (805.17)	1552 (334.12)	2865 (701.62)	2984 (734.91)	3144 (779.64)	2512 (602.91)	2546 (612.44)
DF _{A+S/V}	1226 (10.34)	1193 (7.41)	1135 (2.23)	1259 (13.33)	2945 (165.18)	2935 (164.25)	3756 (238.18)	3259 (193.41)	3246 (192.26)	1579 (42.16)	2881 (159.36)	2997 (169.87)	3154 (183.92)	2524 (127.25)	2560 (130.48)
E _{1-ctx}	2939 (0.38)	2941 (0.47)	2933 (0.18)	2945 (0.60)	3338 (14.01)	3360 (14.01)	4028 (37.58)	3596 (22.83)	3580 (22.27)	2950 (0.76)	3355 (14.59)	3456 (18.06)	3562 (21.66)	3046 (4.04)	3191 (8.99)
E _{2-ctx}	2927 (0.44)	2931 (0.59)	2920 (0.22)	2935 (0.72)	3360 (15.32)	3347 (14.85)	4045 (38.83)	3609 (23.85)	3596 (23.39)	2939 (0.84)	3363 (15.40)	3464 (18.89)	3569 (22.47)	3055 (4.82)	3184 (9.27)
E _{3-ctx}	3751 (0.25)	3753 (0.28)	3746 (0.11)	3756 (0.37)	4028 (7.64)	4045 (8.10)	4476 (19.61)	4197 (12.16)	4185 (11.84)	3761 (0.51)	4036 (7.87)	4120 (10.09)	4173 (11.52)	3835 (2.48)	3931 (5.06)
E _{DDG}	3250 (0.43)	3255 (0.56)	3243 (0.21)	3259 (0.70)	3596 (11.12)	3609 (11.51)	4197 (29.69)	3787 (17.02)	3803 (17.50)	3265 (0.90)	3603 (11.33)	3694 (14.15)	3799 (17.39)	3344 (3.33)	3471 (7.25)
E _{LTO}	3240 (0.31)	3243 (0.39)	3235 (0.15)	3246 (0.50)	3580 (10.82)	3596 (11.31)	4185 (29.57)	3803 (17.73)	3770 (16.70)	3249 (0.57)	3591 (11.18)	3688 (14.18)	3777 (16.93)	3319 (2.75)	3457 (7.02)
E _{MA}	1577 (2.04)	1562 (1.06)	1552 (0.37)	1579 (2.14)	2950 (90.83)	2939 (90.10)	3761 (143.32)	3265 (111.24)	3249 (110.16)	1714 (10.87)	2884 (86.57)	3006 (94.43)	3159 (104.39)	2528 (63.57)	2566 (65.98)
E _{SANCOV}	2872 (0.48)	2876 (0.61)	2865 (0.23)	2881 (0.78)	3355 (17.37)	3363 (17.64)	4036 (41.21)	3603 (26.05)	3591 (25.63)	2884 (0.89)	3334 (16.65)	3447 (20.59)	3569 (24.84)	3045 (6.52)	3170 (10.89)
E _{coarse}	2990 (0.40)	2995 (0.57)	2984 (0.21)	2997 (0.66)	3456 (16.07)	3464 (16.34)	4120 (38.34)	3694 (24.06)	3688 (23.86)	3006 (0.93)	3447 (15.76)	3530 (18.53)	3669 (23.21)	3152 (5.87)	3281 (10.20)
P _{2-gram}	3148 (0.28)	3150 (0.35)	3144 (0.14)	3154 (0.45)	3562 (13.46)	3569 (13.68)	4173 (32.92)	3799 (21.01)	3777 (20.31)	3159 (0.64)	3569 (13.67)	3669 (16.86)	3722 (18.56)	3277 (4.39)	3404 (8.42)
P _{4-gram}	2518 (0.43)	2520 (0.53)	2512 (0.21)	2524 (0.68)	3046 (21.50)	3055 (21.84)	3835 (52.96)	3344 (33.39)	3319 (32.39)	2528 (0.86)	3045 (21.45)	3152 (25.75)	3277 (30.73)	2652 (5.77)	2837 (13.16)
P _{8-gram}	2552 (0.47)	2556 (0.64)	2546 (0.24)	2560 (0.78)	3191 (25.62)	3184 (25.35)	3931 (54.77)	3471 (36.64)	3457 (36.08)	2566 (1.00)	3170 (24.79)	3281 (29.18)	3404 (34.00)	2837 (11.68)	2935 (15.54)

(L) *lcms* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	2145 (1.62)	2175 (3.04)	2139 (1.33)	2207 (4.55)	2585 (22.46)	2585 (22.44)	2585 (22.43)	2585 (22.44)	2584 (22.41)	2282 (8.10)	2584 (22.37)	2583 (22.34)	2582 (22.31)	2587 (22.51)	2583 (22.34)
DF _{A+S/A}	2175 (0.89)	2190 (1.57)	2170 (0.63)	2219 (2.93)	2585 (19.91)	2585 (19.89)	2585 (19.88)	2585 (19.90)	2584 (19.86)	2287 (6.06)	2584 (19.83)	2583 (19.80)	2582 (19.76)	2587 (19.96)	2583 (19.80)
DF _{A+S/O}	2139 (4.18)	2170 (5.65)	2117 (3.07)	2205 (7.38)	2585 (25.89)	2585 (25.87)	2585 (25.86)	2585 (25.87)	2584 (25.84)	2282 (11.11)	2583 (25.80)	2583 (25.76)	2582 (25.73)	2586 (25.95)	2583 (25.76)
DF _{A+S/V}	2207 (0.45)	2219 (1.00)	2205 (0.35)	2236 (1.75)	2586 (17.66)	2585 (17.64)	2585 (17.63)	2585 (17.65)	2584 (17.61)	2296 (4.51)	2584 (17.59)	2583 (17.55)	2582 (17.51)	2587 (17.71)	2583 (17.57)
E _{1-ctx}	2585 (0.02)	2585 (0.03)	2585 (0.02)	2586 (0.03)	2587 (0.09)	2587 (0.08)	2587 (0.08)	2587 (0.08)	2587 (0.09)	2586 (0.04)	2587 (0.10)	2587 (0.08)	2587 (0.09)	2587 (0.10)	2587 (0.09)
E _{2-ctx}	2585 (0.05)	2585 (0.05)	2585 (0.05)	2585 (0.05)	2587 (0.12)	2586 (0.08)	2586 (0.09)	2586 (0.10)	2586 (0.10)	2585 (0.05)	2587 (0.12)	2586 (0.09)	2586 (0.10)	2587 (0.13)	2586 (0.09)
E _{3-ctx}	2585 (0.00)	2585 (0.00)	2585 (0.00)	2585 (0.01)	2587 (0.09)	2586 (0.05)	2586 (0.04)	2586 (0.06)	2587 (0.09)	2585 (0.01)	2587 (0.10)	2586 (0.06)	2586 (0.06)	2587 (0.09)	2587 (0.08)
E _{DDG}	2585 (0.02)	2585 (0.03)	2585 (0.02)	2585 (0.03)	2587 (0.10)	2586 (0.07)	2586 (0.07)	2586 (0.07)	2587 (0.08)	2585 (0.04)	2587 (0.10)	2586 (0.08)	2586 (0.08)	2587 (0.10)	2586 (0.08)
E _{LTO}	2584 (0.06)	2584 (0.06)	2584 (0.06)	2584 (0.07)	2587 (0.17)	2586 (0.13)	2587 (0.16)	2587 (0.15)	2586 (0.13)	2585 (0.09)	2586 (0.14)	2586 (0.14)	2587 (0.15)	2587 (0.18)	2586 (0.13)
E _{MA}	2282 (0.11)	2287 (0.31)	2282 (0.09)	2296 (0.74)	2586 (13.42)	2585 (13.39)	2585 (13.38)	2585 (13.41)	2585 (13.39)	2321 (1.79)	2585 (13.37)	2583 (13.31)	2583 (13.29)	2587 (13.46)	2584 (13.34)
E _{SANCOV}	2584 (0.06)	2584 (0.06)	2583 (0.05)	2584 (0.07)	2587 (0.20)	2587 (0.18)	2587 (0.20)	2587 (0.19)	2586 (0.17)	2585 (0.10)	2586 (0.16)	2586 (0.17)	2587 (0.19)	2588 (0.22)	2586 (0.17)
E _{coarse}	2583 (0.09)	2583 (0.09)	2583 (0.08)	2583 (0.10)	2587 (0.24)	2586 (0.20)	2586 (0.22)	2586 (0.22)	2586 (0.22)	2583 (0.10)	2586 (0.23)	2585 (0.18)	2586 (0.22)	2587 (0.26)	2586 (0.20)
P _{2-gram}	2582 (0.02)	2582 (0.02)	2582 (0.02)	2582 (0.03)	2587 (0.21)	2586 (0.17)	2586 (0.18)	2587 (0.19)	2587 (0.20)	2583 (0.04)	2587 (0.21)	2586 (0.18)	2585 (0.14)	2587 (0.21)	2587 (0.20)
P _{4-gram}	2587 (0.00)	2587 (0.00)	2586 (0.00)	2587 (0.03)	2587 (0.03)	2587 (0.02)	2587 (0.02)	2587 (0.03)	2587 (0.03)	2587 (0.00)	2588 (0.05)	2587 (0.03)	2587 (0.03)	2587 (0.03)	2587 (0.03)
P _{8-gram}	2583 (0.11)	2583 (0.12)	2583 (0.10)	2583 (0.14)	2587 (0.27)	2586 (0.23)	2587 (0.26)	2586 (0.25)	2586 (0.24)	2584 (0.15)	2586 (0.25)	2586 (0.23)	2587 (0.27)	2587 (0.29)	2585 (0.20)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(M) *libpng* (context-insensitive).

B	3517 (4.62)	3587 (6.72)	3501 (4.14)	3709 (10.34)	5240 (55.89)	5248 (56.12)	5176 (53.99)	5235 (55.74)	5215 (55.15)	3922 (16.69)	5190 (54.40)	5238 (55.82)	5236 (55.76)	5209 (54.96)	5198 (54.63)
DF _{A+S/A}	3587 (5.51)	3588 (5.54)	3553 (4.50)	3717 (9.34)	5192 (52.70)	5202 (53.01)	5125 (50.74)	5190 (52.65)	5166 (51.95)	3886 (14.31)	5138 (51.12)	5188 (52.59)	5188 (52.60)	5160 (51.77)	5152 (51.53)
DF _{A+S/O}	3501 (7.38)	3553 (8.99)	3436 (5.41)	3687 (13.10)	5240 (60.74)	5255 (61.19)	5174 (58.72)	5237 (60.64)	5216 (59.98)	3909 (19.91)	5186 (59.08)	5240 (60.74)	5241 (60.75)	5206 (59.69)	5195 (59.35)
DF _{A+S/V}	3709 (4.24)	3717 (4.48)	3687 (3.63)	3791 (6.56)	5212 (46.50)	5226 (46.88)	5142 (44.52)	5204 (46.28)	5189 (45.83)	3960 (11.29)	5158 (44.97)	5208 (46.38)	5209 (46.42)	5179 (45.56)	5163 (45.11)
E _{1-ctx}	5240 (3.09)	5192 (2.13)	5240 (3.09)	5212 (2.54)	5210 (2.49)	5222 (2.74)	5166 (2.75)	5223 (2.75)	5195 (2.19)	5178 (1.86)	5185 (2.01)	5209 (2.48)	5207 (2.43)	5186 (2.02)	5193 (2.16)
E _{2-ctx}	5248 (3.34)	5202 (2.45)	5255 (3.49)	5226 (2.91)	5222 (2.85)	5206 (2.53)	5177 (1.95)	5231 (3.00)	5200 (2.41)	5177 (1.94)	5197 (2.35)	5206 (2.52)	5206 (2.53)	5200 (2.41)	5211 (2.62)
E _{3-ctx}	5176 (3.61)	5125 (2.59)	5174 (3.57)	5142 (2.92)	5166 (3.40)	5177 (3.63)	5096 (2.00)	5169 (3.47)	5146 (3.02)	5106 (2.21)	5128 (2.65)	5159 (3.27)	5159 (3.26)	5134 (2.76)	5136 (2.80)
E _{DDG}	5235 (3.04)	5190 (2.15)	5237 (3.30)	5204 (2.77)	5223 (2.81)	5231 (2.95)	5169 (1.75)	5220 (2.75)	5206 (2.47)	5171 (1.79)	5191 (2.17)	5220 (2.74)	5218 (2.71)	5196 (2.26)	5195 (2.26)
E _{LTO}	5215 (3.30)	5166 (2.33)	5216 (3.30)	5189 (2.77)	5195 (2.89)	5200 (3.00)	5146 (1.93)	5206 (3.12)	5167 (2.35)	5154 (2.08)	5164 (2.28)	5189 (2.77)	5185 (2.70)	5164 (2.29)	5174 (2.48)
E _{MA}	3922 (5.80)	3886 (4.83)	3909 (5.45)	3960 (6.80)	5178 (39.66)	5177 (39.63)	5106 (37.74)	5171 (39.49)	5154 (39.01)	3979 (7.33)	5128 (38.32)	5157 (39.10)	5165 (39.31)	5156 (39.08)	5148 (38.85)
E _{SANCOV}	5190 (3.28)	5138 (2.24)	5186 (3.20)	5158 (2.64)	5185 (3.18)	5197 (3.42)	5128 (2.04)	5191 (3.29)	5164 (2.76)	5128 (2.04)	5143 (2.34)	5182 (3.11)	5179 (3.07)	5151 (2.51)	5157 (2.62)
E _{coarse}	5238 (3.63)	5188 (2.64)	5240 (3.68)	5208 (3.05)	5209 (3.07)	5206 (3.00)	5159 (2.08)	5220 (3.28)	5189 (2.66)	5157 (2.03)	5182 (2.52)	5186 (2.61)	5198 (2.84)	5192 (2.72)	5196 (2.80)
P _{2-gram}	5236 (3.32)	5188 (2.38)	5241 (3.42)	5209 (2.79)	5207 (2.74)	5206 (2.74)	5159 (1.79)	5218 (2.97)	5185 (2.32)	5165 (1.91)	5179 (2.21)	5198 (2.56)	5185 (2.31)	5178 (2.19)	5191 (2.43)
P _{4-gram}	5209 (3.08)	5160 (2.12)	5206 (3.03)	5179 (2.49)	5186 (2.63)	5200 (2.91)	5134 (1.60)	5196 (2.82)	5164 (2.20)	5156 (2.04)	5151 (1.95)	5192 (2.74)	5178 (2.48)	5143 (1.78)	5154 (2.00)
P _{8-gram}	5198 (3.18)	5152 (2.27)	5195 (3.12)	5163 (2.49)	5193 (3.09)	5211 (3.44)	5136 (1.95)	5195 (3.13)	5174 (2.71)	5148 (2.19)	5157 (2.38)	5196 (3.14)	5191 (3.05)	5154 (2.32)	5143 (2.09)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(N) *libpng* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	6692 (5.24)	7259 (14.16)	6662 (4.76)	7855 (23.53)	9117 (43.36)	13119 (106.30)	7076 (11.27)	13298 (109.11)	9318 (46.54)	13233 (108.10)	12924 (103.23)	13063 (105.43)	13364 (110.15)	12042 (89.36)
DF _{A+S/A}	7259 (2.91)	7591 (7.61)	7222 (2.37)	8100 (14.82)	9322 (32.14)	13163 (86.59)	7705 (9.22)	13342 (89.14)	9410 (33.40)	13281 (88.26)	12961 (83.73)	13100 (85.71)	13405 (90.02)	12106 (71.61)
DF _{A+S/O}	6662 (9.33)	7222 (18.52)	6500 (6.67)	7833 (28.56)	9093 (49.23)	13114 (115.22)	6835 (12.18)	13289 (118.10)	9300 (52.63)	13227 (117.08)	12916 (111.97)	13057 (114.28)	13357 (119.21)	12027 (97.39)
DF _{A+S/V}	7855 (1.46)	8100 (4.62)	7833 (1.18)	8342 (7.75)	9564 (23.53)	13214 (70.68)	8328 (7.57)	13413 (73.25)	9532 (23.12)	13358 (72.54)	13018 (68.14)	13165 (70.05)	13462 (73.88)	12197 (57.54)
E _{1-ctx}	9117 (1.24)	9322 (3.52)	9093 (0.98)	9564 (6.21)	9609 (6.71)	13174 (46.30)	9464 (5.10)	13454 (49.41)	10384 (15.32)	13398 (48.79)	13072 (45.17)	13213 (46.74)	13454 (49.41)	12231 (35.83)
E _{2-ctx}	13119 (0.15)	13163 (0.49)	13114 (0.11)	13214 (0.88)	13174 (0.58)	13978 (0.89)	13215 (0.89)	14334 (9.43)	13396 (2.27)	14296 (9.14)	14099 (7.63)	14152 (8.04)	14198 (8.39)	13724 (4.77)
E _{DDG}	7076 (424.12)	7705 (470.73)	6835 (406.31)	8328 (516.87)	9464 (601.02)	13215 (878.89)	2565 (90.00)	13370 (890.37)	9726 (620.47)	13312 (886.06)	13020 (864.48)	13148 (873.93)	13431 (894.87)	12196 (803.38)
E _{LTO}	13298 (0.19)	13342 (0.53)	13289 (0.12)	13413 (1.06)	13454 (1.37)	14334 (8.00)	13370 (0.74)	13975 (5.29)	13636 (2.74)	14014 (5.59)	14035 (5.74)	14068 (5.99)	14219 (7.13)	13835 (4.24)
E _{MA}	9318 (0.50)	9410 (1.49)	9300 (0.30)	9532 (2.80)	10384 (11.99)	13396 (44.48)	9726 (4.90)	13636 (47.06)	9935 (7.15)	13595 (46.62)	13234 (42.73)	13379 (44.29)	13670 (47.43)	12531 (35.14)
E _{SANCOV}	13233 (0.18)	13281 (0.54)	13227 (0.14)	13358 (1.13)	13398 (1.43)	14296 (8.23)	13312 (0.78)	14014 (6.09)	13595 (2.92)	13897 (5.21)	14024 (6.17)	14059 (6.43)	14173 (7.30)	13781 (4.33)
E _{coarse}	12924 (0.17)	12961 (0.45)	12916 (0.11)	13018 (0.89)	13072 (1.31)	14099 (0.92)	13020 (0.92)	14035 (8.78)	13234 (2.57)	14024 (8.69)	13681 (6.03)	13817 (7.09)	14026 (8.71)	13562 (5.11)
P _{2-gram}	13063 (0.17)	13100 (0.46)	13057 (0.12)	13165 (0.95)	13213 (1.32)	14152 (8.52)	13148 (0.82)	14068 (7.87)	13379 (2.59)	14059 (7.80)	13817 (5.95)	13772 (5.61)	14059 (7.81)	13647 (4.65)
P _{4-gram}	13364 (0.15)	13405 (0.46)	13357 (0.10)	13462 (0.89)	13454 (0.84)	14198 (6.41)	13431 (0.66)	14219 (6.56)	13670 (2.45)	14173 (6.22)	14026 (5.12)	14059 (5.37)	13996 (4.90)	13734 (2.93)
P _{8-gram}	12042 (0.30)	12106 (0.84)	12027 (0.18)	12197 (1.60)	12231 (1.88)	13724 (14.32)	12196 (1.58)	13835 (15.24)	12531 (4.38)	13781 (14.79)	13562 (12.97)	13647 (13.67)	13734 (14.40)	12943 (7.81)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(O) *libxslt* (context-insensitive).

B	15894 (12.90)	17988 (27.77)	15577 (10.65)	20647 (46.66)	35550 (152.52)	51395 (265.06)	18311 (30.06)	52662 (274.07)	27049 (92.13)	53055 (276.86)	50012 (255.24)	51153 (263.35)	52255 (271.18)	46505 (230.33)
DF _{A+S/A}	17988 (8.95)	19288 (16.83)	17666 (7.01)	21662 (31.21)	36054 (118.38)	51530 (212.12)	20803 (24.19)	52789 (219.75)	27505 (66.60)	53185 (222.14)	50120 (203.58)	51268 (210.54)	52384 (217.29)	46691 (182.81)
DF _{A+S/O}	15577 (18.67)	17666 (34.59)	14830 (12.98)	20378 (55.24)	35500 (170.44)	51392 (291.52)	17458 (33.00)	52648 (301.08)	26855 (104.59)	53046 (304.11)	49982 (280.77)	51139 (289.59)	52235 (297.93)	46474 (254.05)
DF _{A+S/V}	20647 (5.91)	21662 (11.11)	20378 (4.53)	22859 (17.25)	36626 (87.87)	51684 (165.11)	23198 (18.99)	52985 (171.78)	27843 (42.82)	53389 (173.85)	50319 (158.10)	51469 (164.01)	52557 (169.58)	46918 (140.66)
E _{1-ctx}	35550 (1.50)	36054 (2.94)	35500 (1.36)	36626 (4.58)	38447 (9.78)	52927 (51.12)	37252 (6.36)	54607 (55.91)	38601 (10.21)	54923 (56.82)	52259 (49.21)	53163 (51.79)	53979 (54.12)	48634 (38.86)
E _{2-ctx}	51395 (0.65)	51530 (0.92)	51392 (0.65)	51684 (1.22)	52927 (3.65)	55705 (9.09)	51920 (1.68)	57713 (13.03)	52149 (2.13)	58042 (13.67)	56009 (9.69)	56540 (10.73)	56414 (10.48)	54295 (6.33)
E _{DDG}	18311 (226.16)	20503 (265.20)	17458 (210.97)	23198 (313.21)	37252 (563.55)	51920 (824.81)	10667 (90.00)	53053 (844.99)	29371 (423.16)	53406 (851.29)	50564 (800.66)	51612 (819.32)	52663 (838.05)	47330 (743.05)
E _{LTO}	52662 (0.69)	52789 (0.93)	52648 (0.66)	52985 (1.31)	54607 (4.41)	57713 (10.34)	53053 (1.43)	56351 (7.74)	53593 (2.47)	57010 (9.00)	56265 (7.58)	56560 (8.14)	56978 (8.94)	55248 (5.63)
E _{MA}	27049 (2.88)	27505 (4.61)	26855 (2.14)	27843 (5.90)	38601 (46.81)	52149 (98.34)	29371 (11.71)	53593 (103.83)	29346 (11.61)	54002 (105.39)	50898 (93.58)	52073 (98.05)	53096 (101.94)	47756 (81.63)
E _{SANCOV}	53055 (0.65)	53185 (0.90)	53046 (0.63)	53389 (1.28)	54923 (4.19)	58042 (10.11)	53406 (1.32)	57010 (8.15)	54002 (2.45)	56791 (7.74)	56728 (7.62)	56998 (8.13)	57265 (8.64)	55603 (5.48)
E _{coarse}	50012 (0.74)	50120 (0.96)	49982 (0.68)	50319 (1.36)	52259 (5.27)	56009 (12.82)	50564 (1.86)	56265 (13.34)	50898 (2.53)	56728 (14.27)	54036 (8.85)	54975 (10.74)	55561 (11.92)	53437 (7.64)
P _{2-gram}	51153 (0.71)	51268 (0.94)	51139 (0.68)	51469 (1.33)	53163 (4.67)	56540 (11.32)	51612 (1.61)	56560 (11.35)	52073 (2.52)	56998 (12.22)	54975 (8.23)	54969 (8.22)	55968 (10.19)	54052 (6.42)
P _{4-gram}	52255 (0.66)	52384 (0.91)	52235 (0.62)	52557 (1.24)	53979 (3.98)	56414 (8.67)	52663 (1.45)	56978 (9.76)	53096 (2.28)	57265 (10.31)	55561 (7.03)	55968 (7.81)	55292 (6.51)	54116 (4.24)
P _{8-gram}	46505 (0.84)	46691 (1.24)	46474 (0.77)	46918 (1.73)	48634 (5.46)	54295 (17.73)	47330 (2.63)	55248 (19.80)	47756 (3.55)	55603 (20.57)	53437 (15.87)	54052 (17.20)	54116 (17.34)	50531 (9.57)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(P) *libxslt* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	3534 (1.46)	3591 (3.10)	3522 (1.10)	3644 (4.62)	4078 (17.08)	4076 (17.02)	4123 (18.38)	4277 (22.80)	4159 (19.39)	3563 (2.30)	4143 (18.96)	4090 (17.42)	4073 (16.93)	4120 (18.28)	3910 (12.26)
DF _{A+S/A}	3591 (0.83)	3616 (1.53)	3583 (0.61)	3661 (2.80)	4080 (14.55)	4078 (14.49)	4124 (15.79)	4278 (20.12)	4159 (16.77)	3608 (1.30)	4144 (16.34)	4091 (14.86)	4075 (14.40)	4122 (15.73)	3922 (10.13)
DF _{A+S/O}	3522 (2.49)	3583 (4.29)	3493 (1.65)	3640 (5.94)	4078 (18.68)	4075 (18.61)	4123 (20.00)	4277 (24.48)	4159 (21.04)	3551 (3.36)	4144 (20.60)	4089 (19.02)	4072 (18.52)	4120 (19.90)	3908 (13.74)
DF _{A+S/V}	3644 (0.50)	3661 (0.97)	3640 (0.38)	3680 (1.48)	4082 (12.57)	4079 (12.50)	4126 (13.78)	4280 (18.04)	4159 (14.70)	3653 (0.74)	4144 (14.29)	4093 (12.87)	4077 (12.43)	4123 (13.70)	3931 (8.40)
E _{1-ctx}	4078 (0.03)	4080 (0.08)	4078 (0.03)	4082 (0.13)	4142 (1.61)	4150 (1.79)	4169 (2.26)	4318 (5.93)	4196 (2.94)	4079 (0.05)	4192 (2.83)	4152 (1.84)	4150 (1.81)	4184 (2.64)	4113 (0.88)
E _{2-ctx}	4076 (0.03)	4078 (0.08)	4075 (0.03)	4079 (0.12)	4150 (1.84)	4146 (1.76)	4169 (2.33)	4320 (6.04)	4195 (2.97)	4076 (0.05)	4190 (2.83)	4152 (1.89)	4152 (1.91)	4181 (2.63)	4109 (0.86)
E _{3-ctx}	4123 (0.03)	4124 (0.05)	4123 (0.02)	4126 (0.08)	4169 (1.13)	4169 (1.15)	4179 (1.38)	4331 (5.07)	4210 (2.14)	4123 (0.03)	4206 (2.03)	4168 (1.11)	4173 (1.23)	4198 (1.85)	4144 (0.52)
E _{DDG}	4277 (0.02)	4278 (0.04)	4277 (0.01)	4280 (0.09)	4318 (0.98)	4320 (1.03)	4331 (1.28)	4454 (4.15)	4358 (1.92)	4278 (0.03)	4353 (1.79)	4317 (0.95)	4318 (0.98)	4347 (1.64)	4299 (0.53)
E _{LTO}	4159 (0.01)	4159 (0.01)	4159 (0.01)	4159 (0.02)	4196 (0.89)	4195 (0.89)	4210 (1.25)	4358 (4.81)	4219 (1.47)	4159 (0.01)	4222 (1.54)	4198 (0.95)	4200 (1.00)	4224 (1.57)	4175 (0.41)
E _{MA}	3563 (1.17)	3608 (2.44)	3551 (0.83)	3653 (3.72)	4079 (15.80)	4076 (15.74)	4123 (17.07)	4278 (21.45)	4159 (18.07)	3567 (1.27)	4143 (17.64)	4090 (16.14)	4073 (15.65)	4120 (16.98)	3909 (10.99)
E _{SANCOV}	4143 (0.01)	4144 (0.01)	4144 (0.01)	4144 (0.02)	4192 (1.18)	4190 (1.13)	4206 (1.51)	4353 (5.06)	4222 (1.91)	4143 (0.01)	4211 (1.64)	4190 (1.14)	4192 (1.18)	4212 (1.67)	4161 (0.43)
E _{coarse}	4090 (0.04)	4091 (0.05)	4089 (0.02)	4093 (0.10)	4152 (1.55)	4152 (1.54)	4168 (1.94)	4317 (5.59)	4198 (2.67)	4090 (0.05)	4190 (2.49)	4144 (1.35)	4151 (1.53)	4180 (2.24)	4116 (0.67)
P _{2-gram}	4073 (0.04)	4075 (0.09)	4072 (0.03)	4077 (0.14)	4150 (1.95)	4152 (1.99)	4173 (2.50)	4318 (6.07)	4200 (3.16)	4073 (0.05)	4192 (2.97)	4151 (1.96)	4144 (1.78)	4184 (2.77)	4109 (0.92)
P _{4-gram}	4120 (0.03)	4122 (0.08)	4120 (0.03)	4123 (0.11)	4184 (1.59)	4181 (1.53)	4198 (1.94)	4347 (5.54)	4224 (2.55)	4120 (0.04)	4212 (2.28)	4180 (1.50)	4184 (1.59)	4198 (1.62)	4144 (0.61)
P _{8-gram}	3910 (0.19)	3922 (0.49)	3908 (0.13)	3931 (0.71)	4113 (5.37)	4109 (5.29)	4144 (6.17)	4299 (10.15)	4175 (6.98)	3909 (0.16)	4161 (6.62)	4116 (5.45)	4109 (5.27)	4144 (6.17)	3998 (2.43)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(Q) *mbedtls* (context-insensitive).

B	6757 (3.76)	6869 (5.48)	6722 (3.22)	6985 (7.27)	7732 (18.74)	7616 (16.95)	7761 (19.18)	8385 (28.76)	7928 (21.75)	6773 (4.00)	7863 (20.75)	7728 (18.67)	7632 (17.19)	7879 (20.99)	7454 (14.45)
DF _{A+S/A}	6869 (3.45)	6900 (3.91)	6828 (2.83)	7024 (5.77)	7732 (16.44)	7618 (14.73)	7759 (16.85)	8388 (26.33)	7933 (19.47)	6860 (3.31)	7869 (18.51)	7726 (16.36)	7636 (15.00)	7882 (18.71)	7470 (12.50)
DF _{A+S/O}	6722 (5.68)	6828 (7.36)	6632 (4.27)	6956 (9.36)	7708 (21.19)	7590 (19.33)	7734 (21.61)	8361 (31.46)	7904 (24.28)	6731 (5.83)	7843 (23.32)	7703 (21.12)	7604 (19.56)	7855 (23.51)	7426 (16.76)
DF _{A+S/V}	6985 (2.35)	7024 (2.91)	6956 (1.92)	7059 (3.43)	7758 (13.67)	7678 (12.50)	7801 (14.31)	8423 (23.41)	7959 (16.62)	6968 (2.10)	7903 (15.80)	7769 (13.83)	7683 (12.58)	7915 (15.97)	7517 (10.14)
E _{1-ctx}	7732 (1.95)	7732 (1.95)	7708 (1.63)	7758 (2.29)	7850 (3.50)	7804 (3.92)	7881 (5.60)	8490 (11.95)	8039 (6.00)	7698 (1.50)	7987 (5.31)	7869 (3.76)	7813 (3.02)	8015 (5.69)	7817 (3.07)
E _{2-ctx}	7616 (2.87)	7618 (2.90)	7590 (2.51)	7678 (3.70)	7804 (5.41)	7658 (3.44)	7779 (5.07)	8407 (13.56)	7942 (7.28)	7571 (2.26)	7886 (6.52)	7771 (4.96)	7697 (3.97)	7924 (7.03)	7717 (4.24)
E _{3-ctx}	7761 (2.27)	7759 (2.24)	7734 (1.91)	7801 (2.79)	7881 (3.85)	7779 (2.50)	7839 (3.29)	8463 (11.52)	8015 (5.62)	7725 (1.79)	7963 (4.92)	7845 (3.36)	7796 (2.73)	7993 (5.33)	7828 (3.14)
E _{DDG}	8385 (1.88)	8388 (1.92)	8361 (1.59)	8423 (2.34)	8490 (3.16)	8407 (2.15)	8463 (2.83)	8990 (9.23)	8615 (4.68)	8351 (1.46)	8564 (4.06)	8463 (2.83)	8413 (2.22)	8597 (4.46)	8449 (2.66)
E _{LTO}	7928 (2.06)	7933 (2.12)	7904 (1.75)	7959 (2.45)	8039 (3.49)	7942 (2.24)	8015 (3.18)	8615 (10.90)	8095 (4.20)	7889 (1.55)	8088 (4.11)	8017 (3.20)	7957 (2.43)	8139 (4.77)	7985 (2.79)
E _{MA}	6773 (3.92)	6860 (5.25)	6731 (3.28)	6968 (6.92)	7698 (18.12)	7571 (16.17)	7725 (18.53)	8351 (28.13)	7889 (21.05)	6728 (3.23)	7820 (19.99)	7694 (18.06)	7591 (16.48)	7842 (20.33)	7410 (13.70)
E _{SANCOV}	7863 (2.10)	7869 (2.18)	7843 (1.84)	7903 (2.61)	7987 (3.71)	7886 (2.40)	7963 (3.39)	8564 (11.20)	8088 (5.01)	7820 (1.54)	8014 (4.06)	7959 (3.34)	7901 (2.59)	8072 (4.81)	7915 (2.77)
E _{coarse}	7728 (2.23)	7726 (2.20)	7703 (1.89)	7769 (2.76)	7869 (4.09)	7771 (2.79)	7845 (3.76)	8463 (11.95)	8017 (6.05)	7694 (1.77)	7959 (5.27)	7818 (3.41)	7782 (2.94)	7984 (5.61)	7803 (3.22)
P _{2-gram}	7632 (2.83)	7636 (2.89)	7604 (2.46)	7683 (3.52)	7813 (5.27)	7697 (3.71)	7796 (5.05)	8413 (13.36)	7957 (7.21)	7591 (2.28)	7901 (6.46)	7782 (4.86)	7694 (3.66)	7935 (6.92)	7730 (4.15)
P _{4-gram}	7879 (2.23)	7882 (2.27)	7855 (1.91)	7915 (2.69)	8015 (3.99)	7924 (2.81)	7993 (3.71)	8597 (11.54)	8139 (5.60)	7842 (1.75)	8072 (4.73)	7984 (3.59)	7935 (2.95)	8081 (4.84)	7933 (2.93)
P _{8-gram}	7454 (2.40)	7470 (2.63)	7426 (2.03)	7517 (3.27)	7817 (7.40)	7717 (6.03)	7828 (7.54)	8449 (16.08)	7985 (9.70)	7410 (1.81)	7915 (8.74)	7803 (7.20)	7730 (6.20)	7933 (9.00)	7605 (4.48)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(R) *mbedtls* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	12379 (10.89)	12360 (10.72)	12322 (10.38)	12383 (10.92)	12663 (13.43)	12745 (14.16)	12708 (13.84)	12847 (15.08)	12915 (15.69)	12362 (10.74)	12862 (15.21)	12568 (12.58)	12800 (14.66)	12643 (13.26)	12298 (10.16)
DF _{A+S/A}	12360 (4.98)	11963 (1.60)	11892 (1.00)	12062 (2.44)	12557 (6.64)	12674 (7.64)	12629 (7.26)	12810 (8.79)	12887 (9.45)	12001 (1.93)	12823 (8.91)	12547 (6.56)	12737 (8.17)	12542 (6.52)	11981 (1.75)
DF _{A+S/O}	12322 (7.16)	11892 (3.43)	11734 (2.05)	11993 (4.30)	12530 (8.98)	12658 (10.09)	12610 (9.67)	12802 (11.34)	12883 (12.04)	11917 (3.64)	12816 (11.46)	12501 (8.72)	12721 (10.63)	12517 (8.86)	11823 (2.82)
DF _{A+S/V}	12383 (3.89)	12062 (1.20)	11993 (0.62)	12091 (1.44)	12577 (5.52)	12689 (6.46)	12645 (6.09)	12818 (7.54)	12895 (8.19)	12076 (1.32)	12833 (7.67)	12578 (5.53)	12753 (7.00)	12566 (5.43)	12076 (1.32)
E _{1-ctx}	12663 (1.56)	12557 (0.70)	12530 (0.49)	12577 (0.87)	12694 (1.80)	12774 (2.45)	12741 (2.18)	12859 (3.13)	12921 (3.62)	12557 (0.71)	12876 (3.26)	12734 (2.13)	12814 (2.77)	12699 (1.84)	12521 (0.42)
E _{2-ctx}	12745 (1.06)	12674 (0.50)	12658 (0.37)	12689 (0.62)	12774 (1.29)	12802 (1.51)	12792 (1.43)	12884 (2.16)	12934 (2.55)	12677 (0.52)	12901 (2.29)	12794 (1.45)	12846 (1.86)	12760 (1.18)	12647 (0.28)
E _{3-ctx}	12708 (1.15)	12629 (0.52)	12610 (0.36)	12645 (0.65)	12741 (1.41)	12792 (1.81)	12740 (1.40)	12870 (2.43)	12928 (2.89)	12632 (0.54)	12886 (1.62)	12767 (1.62)	12829 (2.10)	12724 (1.27)	12600 (0.28)
E _{DDG}	12847 (0.46)	12810 (0.17)	12802 (0.11)	12818 (0.24)	12859 (0.56)	12884 (0.75)	12870 (0.64)	12899 (0.87)	12949 (1.26)	12811 (0.18)	12930 (1.11)	12866 (0.61)	12887 (0.78)	12850 (0.49)	12805 (0.13)
E _{LTO}	12915 (0.35)	12887 (0.13)	12883 (0.09)	12895 (0.19)	12921 (0.39)	12934 (0.49)	12928 (0.44)	12949 (0.61)	12959 (0.68)	12889 (0.14)	12957 (0.68)	12921 (0.39)	12936 (0.51)	12911 (0.32)	12883 (0.10)
E _{MA}	12362 (5.05)	12001 (1.98)	11917 (1.26)	12076 (2.62)	12557 (6.70)	12677 (7.72)	12632 (7.34)	12811 (8.86)	12889 (9.52)	11968 (1.70)	12825 (8.98)	12550 (6.64)	12739 (8.25)	12545 (6.60)	11976 (1.77)
E _{SANCOV}	12862 (0.50)	12823 (0.20)	12816 (0.14)	12833 (0.28)	12876 (0.61)	12901 (0.81)	12886 (0.69)	12930 (1.04)	12957 (1.25)	12825 (0.21)	12928 (1.02)	12882 (0.66)	12911 (0.89)	12863 (0.51)	12816 (0.14)
E _{coarse}	12568 (11.07)	12547 (10.89)	12501 (10.48)	12578 (11.16)	12734 (12.54)	12794 (13.07)	12767 (12.83)	12866 (13.70)	12921 (14.19)	12550 (10.91)	12882 (13.85)	12613 (11.47)	12820 (13.30)	12721 (12.43)	12455 (10.07)
P _{2-gram}	12800 (0.86)	12737 (0.37)	12721 (0.24)	12753 (0.50)	12814 (0.98)	12846 (1.23)	12829 (1.09)	12887 (1.56)	12936 (1.94)	12739 (0.39)	12911 (1.74)	12820 (1.02)	12842 (1.20)	12801 (0.87)	12723 (0.26)
P _{4-gram}	12643 (1.60)	12542 (0.78)	12517 (0.58)	12566 (0.98)	12699 (2.04)	12760 (2.24)	12724 (2.24)	12850 (3.26)	12911 (3.75)	12545 (0.81)	12863 (3.36)	12721 (2.22)	12801 (2.86)	12652 (1.67)	12502 (0.46)
P _{8-gram}	12298 (12.12)	11981 (9.23)	11823 (7.79)	12076 (10.10)	12521 (14.16)	12647 (15.31)	12600 (14.88)	12805 (16.74)	12883 (17.46)	11976 (9.19)	12816 (16.85)	12455 (13.56)	12723 (16.00)	12502 (13.98)	11550 (5.31)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(S) *openh264* (context-insensitive).

B	17666 (11.49)	17711 (11.78)	17666 (11.50)	17790 (12.28)	18083 (14.13)	18220 (14.99)	18082 (14.12)	18320 (15.62)	18399 (16.12)	17724 (11.86)	18329 (15.68)	17947 (13.27)	18295 (15.46)	18067 (14.02)	17609 (11.13)
DF _{A+S/A}	17711 (5.25)	17233 (2.40)	17166 (2.01)	17414 (3.48)	17976 (6.82)	18147 (7.84)	18005 (6.99)	18277 (8.61)	18372 (9.18)	17306 (2.84)	18292 (8.70)	17946 (6.64)	18218 (8.26)	17951 (6.68)	17259 (2.56)
DF _{A+S/O}	17666 (7.00)	17166 (3.97)	16957 (2.70)	17333 (4.98)	17949 (8.71)	18129 (9.81)	17996 (9.00)	18287 (10.76)	18376 (11.30)	17207 (4.22)	18294 (10.80)	17895 (8.39)	18204 (10.26)	17918 (8.52)	17093 (3.53)
DF _{A+S/V}	17790 (3.98)	17414 (1.78)	17333 (1.31)	17474 (2.14)	18048 (5.49)	18215 (6.46)	18096 (5.77)	18346 (7.23)	18433 (7.74)	17448 (1.98)	18359 (7.31)	18038 (5.43)	18284 (6.87)	18034 (5.41)	17427 (1.86)
E _{1-ctx}	18083 (1.72)	17976 (1.65)	17949 (1.22)	18048 (2.06)	18114 (1.96)	18262 (2.19)	18131 (1.81)	18338 (2.74)	18414 (3.03)	17979 (1.67)	18349 (1.31)	18173 (2.76)	18308 (3.53)	18143 (2.59)	17915 (1.30)
E _{2-ctx}	18220 (1.72)	18147 (1.32)	18129 (1.22)	18215 (1.70)	18262 (1.96)	18302 (2.19)	18235 (1.81)	18401 (2.74)	18453 (3.03)	18145 (1.31)	18405 (2.76)	18278 (2.05)	18378 (2.61)	18249 (1.89)	18091 (1.01)
E _{3-ctx}	18082 (2.17)	18005 (1.74)	17996 (1.69)	18096 (2.25)	18131 (2.46)	18235 (3.04)	18052 (2.01)	18298 (3.40)	18376 (3.84)	18017 (1.81)	18315 (3.49)	18168 (2.66)	18284 (3.32)	18123 (2.41)	17945 (1.40)
E _{DDG}	18320 (1.45)	18277 (1.21)	18287 (1.26)	18346 (1.59)	18338 (1.55)	18401 (1.73)	18298 (1.32)	18367 (1.71)	18444 (2.13)	18295 (1.31)	18416 (1.98)	18340 (1.56)	18405 (1.92)	18341 (1.56)	18243 (1.02)
E _{LTO}	18399 (1.29)	18372 (1.14)	18376 (1.16)	18433 (1.47)	18414 (1.37)	18453 (1.58)	18376 (1.16)	18444 (1.54)	18449 (1.57)	18383 (1.20)	18453 (1.58)	18404 (1.32)	18462 (1.64)	18410 (1.35)	18337 (0.95)
E _{MA}	17724 (5.07)	17306 (2.59)	17207 (2.00)	17448 (3.43)	17979 (6.58)	18145 (7.56)	18017 (6.80)	18295 (8.45)	18383 (8.97)	17256 (2.29)	18292 (8.44)	17956 (6.44)	18226 (8.04)	17952 (6.42)	17278 (2.42)
E _{SANCOV}	18329 (1.31)	18292 (1.10)	18294 (1.12)	18359 (1.48)	18349 (1.42)	18405 (1.73)	18315 (1.23)	18416 (1.79)	18453 (1.99)	18292 (1.11)	18392 (1.66)	18353 (1.44)	18430 (1.87)	18345 (1.40)	18261 (0.94)
E _{coarse}	17947 (11.91)	17946 (11.91)	17895 (11.59)	18038 (12.49)	18173 (13.32)	18278 (13.98)	18168 (13.29)	18340 (14.37)	18404 (14.77)	17956 (11.97)	18353 (14.44)	17985 (12.15)	18309 (14.17)	18165 (13.27)	17804 (11.02)
P _{2-gram}	18295 (1.69)	18218 (1.27)	18204 (1.19)	18284 (1.63)	18308 (1.77)	18378 (2.16)	18284 (1.64)	18405 (2.31)	18462 (2.62)	18226 (1.31)	18430 (2.44)	18309 (1.77)	18360 (2.05)	18301 (1.73)	18182 (1.07)
P _{4-gram}	18067 (2.42)	17951 (1.77)	17918 (1.58)	18034 (2.24)	18143 (2.85)	18249 (3.46)	18123 (2.74)	18341 (3.97)	18410 (4.37)	17952 (1.77)	18345 (4.00)	18165 (2.98)	18301 (3.75)	18074 (2.46)	17882 (1.38)
P _{8-gram}	17609 (11.74)	17259 (9.52)	17093 (8.47)	17427 (10.59)	17915 (13.68)	18091 (14.80)	17945 (13.87)	18243 (15.76)	18337 (16.36)	17278 (9.64)	18261 (15.88)	17804 (12.98)	18182 (15.38)	17882 (13.48)	16694 (5.93)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(T) *openh264* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	8920 (5.34)	9124 (7.75)	8929 (5.44)	9629 (13.71)	10300 (21.64)	10324 (21.92)	10418 (23.03)	10405 (22.88)	10362 (22.36)	9107 (7.54)	10358 (22.33)	10222 (20.72)	10279 (21.39)	10431 (23.18)	10319 (21.87)
DF _{A+S/A}	9124 (13.77)	9163 (14.24)	9107 (13.55)	9686 (20.77)	10309 (28.54)	10330 (28.80)	10422 (29.95)	10409 (29.78)	10368 (29.27)	9276 (15.66)	10366 (29.25)	10233 (27.59)	10289 (28.28)	10436 (30.12)	10327 (28.76)
DF _{A+S/O}	8929 (6.96)	9107 (9.09)	8851 (6.02)	9623 (15.27)	10295 (23.32)	10318 (23.60)	10414 (24.75)	10408 (24.68)	10360 (24.10)	9105 (9.07)	10356 (24.05)	10217 (22.39)	10276 (23.10)	10429 (24.93)	10315 (23.57)
DF _{A+S/V}	9629 (0.98)	9686 (1.57)	9623 (0.91)	9840 (3.19)	10341 (8.44)	10359 (8.63)	10440 (9.48)	10435 (9.42)	10390 (8.96)	9714 (1.86)	10388 (8.93)	10283 (7.83)	10328 (8.30)	10448 (9.56)	10357 (8.61)
E _{1-ctx}	10300 (0.20)	10309 (0.28)	10295 (0.14)	10341 (0.59)	10412 (1.28)	10440 (1.55)	10486 (2.00)	10486 (2.01)	10458 (1.73)	10326 (0.45)	10455 (1.70)	10416 (1.33)	10427 (1.43)	10491 (2.05)	10447 (1.63)
E _{2-ctx}	10324 (0.18)	10330 (0.25)	10318 (0.13)	10359 (0.53)	10440 (1.31)	10424 (1.16)	10489 (1.79)	10483 (1.73)	10462 (1.52)	10341 (0.36)	10454 (1.46)	10416 (1.09)	10426 (1.18)	10487 (1.77)	10453 (1.44)
E _{3-ctx}	10418 (0.12)	10422 (0.16)	10414 (0.08)	10440 (0.33)	10486 (0.77)	10489 (0.80)	10510 (1.00)	10523 (1.13)	10490 (1.00)	10438 (0.31)	10504 (0.95)	10476 (0.68)	10484 (0.75)	10527 (1.17)	10501 (0.92)
E _{DDG}	10405 (0.13)	10409 (0.16)	10408 (0.15)	10435 (0.41)	10486 (0.91)	10483 (0.88)	10523 (1.26)	10490 (0.94)	10493 (0.97)	10414 (0.21)	10484 (0.88)	10462 (0.68)	10462 (0.68)	10510 (1.14)	10491 (0.96)
E _{LTO}	10362 (0.14)	10368 (0.20)	10360 (0.12)	10390 (0.42)	10458 (1.08)	10462 (1.11)	10510 (1.58)	10493 (1.41)	10454 (1.04)	10378 (0.30)	10454 (1.04)	10440 (0.90)	10442 (0.92)	10494 (1.42)	10467 (1.16)
E _{MA}	9107 (5.10)	9276 (7.05)	9105 (5.08)	9714 (12.10)	10326 (19.17)	10341 (19.34)	10438 (20.46)	10414 (20.18)	10378 (19.77)	9037 (4.29)	10375 (19.74)	10245 (18.24)	10299 (18.85)	10443 (20.52)	10335 (19.28)
E _{SANCOV}	10358 (0.15)	10366 (0.22)	10356 (0.12)	10388 (0.43)	10455 (1.08)	10454 (1.08)	10504 (1.55)	10484 (1.36)	10454 (1.08)	10375 (0.31)	10432 (0.86)	10427 (0.81)	10431 (0.85)	10483 (1.35)	10459 (1.12)
E _{coarse}	10222 (0.31)	10233 (0.42)	10217 (0.26)	10283 (0.90)	10416 (2.22)	10416 (2.22)	10476 (2.80)	10462 (2.67)	10440 (2.45)	10245 (0.54)	10427 (2.33)	10358 (1.65)	10392 (1.98)	10474 (2.78)	10429 (2.34)
P _{2-gram}	10279 (0.27)	10289 (0.36)	10276 (0.24)	10328 (0.75)	10427 (1.71)	10426 (1.71)	10484 (2.27)	10462 (2.05)	10442 (1.86)	10299 (0.46)	10431 (1.75)	10392 (1.37)	10386 (1.32)	10473 (2.16)	10438 (1.82)
P _{4-gram}	10431 (0.07)	10436 (0.12)	10429 (0.06)	10448 (0.24)	10491 (0.61)	10487 (0.65)	10527 (1.00)	10510 (0.83)	10494 (0.67)	10443 (0.19)	10483 (0.57)	10474 (0.48)	10473 (0.48)	10501 (0.74)	10493 (0.66)
P _{8-gram}	10319 (0.16)	10327 (0.23)	10315 (0.12)	10357 (0.53)	10447 (1.41)	10453 (1.46)	10501 (1.93)	10491 (1.83)	10467 (1.59)	10335 (0.32)	10459 (1.52)	10429 (1.23)	10438 (1.31)	10493 (1.84)	10438 (1.32)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(U) *openssl* (context-insensitive).

B	58935 (9.71)	59778 (11.28)	59715 (11.16)	61518 (14.52)	70654 (31.52)	71816 (33.69)	72870 (35.65)	71399 (32.91)	70170 (30.62)	60822 (13.22)	70196 (30.67)	68397 (27.32)	70538 (31.31)	70914 (32.01)	71997 (34.02)
DF _{A+S/A}	59778 (20.85)	58565 (18.40)	60042 (21.38)	60629 (22.57)	69203 (39.90)	70698 (42.92)	71585 (44.72)	69994 (39.15)	68830 (39.15)	60905 (23.13)	68671 (38.83)	67110 (35.67)	69324 (40.15)	69577 (40.66)	70761 (43.05)
DF _{A+S/O}	59715 (11.30)	60042 (11.91)	59590 (11.07)	61820 (15.22)	70890 (32.13)	72130 (34.44)	73169 (36.38)	71682 (33.60)	70446 (31.30)	61288 (14.23)	70433 (31.28)	68628 (27.91)	70761 (31.89)	71152 (32.62)	72338 (34.83)
DF _{A+S/V}	61518 (9.08)	60629 (7.50)	61820 (9.62)	60519 (7.31)	69667 (23.53)	71033 (25.95)	71965 (27.60)	70285 (24.63)	69142 (22.60)	62559 (10.93)	69099 (22.52)	67394 (19.50)	69792 (23.75)	69867 (23.88)	70999 (25.89)
E _{1-ctx}	70654 (14.93)	69203 (12.57)	70890 (15.31)	69667 (13.32)	66735 (8.55)	68278 (11.06)	68940 (12.14)	67436 (9.70)	66651 (8.42)	71384 (16.12)	66407 (8.02)	66336 (7.91)	67339 (9.54)	67068 (9.54)	68528 (11.47)
E _{2-ctx}	71816 (13.94)	70698 (12.17)	72130 (14.44)	71033 (12.70)	68278 (8.33)	68563 (8.78)	69753 (10.67)	68370 (8.48)	67509 (7.11)	72570 (15.14)	67558 (7.19)	67409 (6.95)	68320 (8.40)	67957 (7.82)	69035 (9.53)
E _{3-ctx}	72870 (14.29)	71585 (12.27)	73169 (14.76)	71965 (12.87)	68940 (8.12)	69753 (9.40)	70015 (9.81)	69115 (8.40)	68429 (7.32)	73600 (15.43)	68366 (7.22)	68344 (7.19)	69155 (8.46)	68793 (7.89)	69898 (9.63)
E _{DDG}	71399 (14.54)	69994 (12.29)	71682 (14.99)	70285 (12.75)	67436 (8.18)	68370 (9.68)	69115 (10.87)	67230 (7.85)	66714 (7.02)	72163 (15.76)	66600 (6.84)	66602 (6.84)	67627 (8.49)	67040 (7.55)	68363 (9.67)
E _{LTO}	70170 (15.11)	68830 (12.91)	70446 (15.56)	69142 (13.43)	66651 (9.34)	67509 (10.75)	68429 (12.26)	66714 (9.44)	65505 (7.46)	70935 (16.37)	65875 (8.07)	65547 (7.53)	66686 (9.40)	66202 (8.60)	67627 (10.94)
E _{MA}	60822 (12.20)	60905 (12.35)	61288 (13.06)	62559 (15.40)	71384 (31.68)	72570 (33.87)	73600 (35.77)	72163 (33.11)	70935 (30.85)	60415 (11.44)	70867 (30.72)	69321 (27.87)	71309 (31.54)	71690 (32.24)	72858 (34.40)
E _{SANCOV}	70196 (15.24)	68671 (12.74)	70433 (15.63)	69099 (13.44)	66407 (9.02)	67558 (10.91)	68366 (12.24)	66600 (9.34)	65875 (8.15)	70867 (16.35)	65174 (7.00)	65599 (7.70)	66648 (9.42)	66198 (8.68)	67728 (11.19)
E _{coarse}	68397 (14.16)	67110 (12.02)	68628 (14.55)	67394 (12.49)	66336 (10.72)	67409 (12.52)	68344 (14.08)	66602 (11.17)	65547 (9.41)	69321 (15.71)	65599 (9.50)	64654 (7.92)	66321 (10.70)	66103 (10.34)	67482 (12.64)
P _{2-gram}	70538 (14.37)	69324 (12.40)	70761 (14.73)	69792 (13.16)	67339 (9.18)	68320 (10.78)	69155 (12.13)	67627 (9.65)	66686 (8.13)	71309 (15.62)	66648 (8.06)	66321 (7.53)	66767 (8.26)	67148 (8.88)	68663 (11.33)
P _{4-gram}	70914 (14.39)	69577 (12.24)	71152 (14.78)	69867 (12.70)	67068 (8.19)	67957 (9.62)	68793 (10.97)	67040 (8.14)	66202 (6.79)	71690 (15.64)	66198 (6.79)	66103 (6.63)	67148 (8.32)	66272 (6.91)	68022 (9.73)
P _{8-gram}	71997 (14.15)	70761 (12.19)	72338 (14.69)	70999 (12.57)	68528 (8.65)	69035 (9.45)	69898 (10.82)	68363 (8.39)	67627 (7.22)	72858 (15.51)	67728 (7.38)	67482 (6.99)	68663 (8.86)	68022 (7.85)	68519 (8.63)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(V) *openssl* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	6954 (0.00)	6954 (0.00)	6954 (0.00)	6954 (0.00)	7989 (14.88)	7971 (14.62)	7974 (14.68)	7996 (14.99)	7980 (14.75)	6962 (0.12)	7976 (14.69)	7961 (14.48)	7960 (14.46)	8045 (15.69)	8228 (18.32)
DF _{A+S/A}	6954 (0.00)	6954 (0.00)	6954 (0.00)	6954 (0.00)	7989 (14.88)	7971 (14.62)	7974 (14.68)	7996 (14.99)	7980 (14.75)	6962 (0.12)	7976 (14.69)	7961 (14.48)	7960 (14.46)	8045 (15.69)	8228 (18.32)
DF _{A+S/O}	6954 (0.00)	6954 (0.00)	6954 (0.00)	6954 (0.00)	7989 (14.88)	7971 (14.62)	7974 (14.68)	7996 (14.99)	7980 (14.75)	6962 (0.12)	7976 (14.69)	7961 (14.48)	7960 (14.46)	8045 (15.69)	8228 (18.32)
DF _{A+S/V}	6954 (0.00)	6954 (0.00)	6954 (0.00)	6954 (0.00)	7989 (14.88)	7971 (14.62)	7974 (14.68)	7996 (14.99)	7980 (14.75)	6962 (0.12)	7976 (14.69)	7961 (14.48)	7960 (14.46)	8045 (15.69)	8228 (18.32)
E _{1-ctx}	7989 (0.00)	7989 (0.00)	7989 (0.00)	7989 (0.00)	8021 (0.41)	8004 (0.20)	8006 (0.21)	8026 (0.47)	8011 (0.28)	7989 (0.00)	8007 (0.23)	8001 (0.16)	8000 (0.14)	8071 (1.03)	8237 (3.10)
E _{2-ctx}	7971 (0.00)	7971 (0.00)	7971 (0.00)	7971 (0.00)	8004 (0.42)	7978 (0.09)	7980 (0.12)	8004 (0.42)	7988 (0.22)	7971 (0.00)	7983 (0.16)	7976 (0.06)	7975 (0.05)	8054 (1.04)	8235 (3.31)
E _{3-ctx}	7974 (0.00)	7974 (0.00)	7974 (0.00)	7974 (0.00)	8006 (0.39)	7980 (0.07)	7980 (0.07)	8006 (0.39)	7988 (0.17)	7974 (0.00)	7984 (0.12)	7979 (0.05)	7978 (0.04)	8055 (1.00)	8235 (3.27)
E _{DDG}	7996 (0.00)	7996 (0.00)	7996 (0.00)	7996 (0.00)	8026 (0.37)	8004 (0.10)	8006 (0.12)	8027 (0.38)	8012 (0.20)	7996 (0.00)	8008 (0.15)	8002 (0.08)	8002 (0.07)	8072 (0.95)	8238 (3.03)
E _{LTO}	7980 (0.00)	7980 (0.00)	7980 (0.00)	7980 (0.00)	8011 (0.39)	7988 (0.11)	7988 (0.11)	8012 (0.41)	7989 (0.12)	7980 (0.00)	7988 (0.11)	7986 (0.07)	7985 (0.07)	8060 (1.01)	8238 (3.23)
EMA	6962 (0.00)	6962 (0.00)	6962 (0.00)	6962 (0.00)	7989 (14.75)	7971 (14.49)	7974 (14.54)	7996 (14.85)	7980 (14.62)	6969 (0.10)	7976 (14.56)	7961 (14.35)	7960 (14.33)	8045 (15.56)	8228 (18.18)
ES _{ANCOV}	7976 (0.00)	7976 (0.00)	7976 (0.00)	7976 (0.00)	8007 (0.39)	7983 (0.10)	7984 (0.11)	8008 (0.40)	7988 (0.16)	7976 (0.00)	7984 (0.11)	7981 (0.07)	7980 (0.06)	8057 (1.02)	8236 (3.26)
E _{coarse}	7961 (0.00)	7961 (0.00)	7961 (0.00)	7961 (0.00)	8001 (0.50)	7976 (0.19)	7979 (0.22)	8002 (0.52)	7986 (0.31)	7961 (0.00)	7981 (0.25)	7971 (0.12)	7971 (0.12)	8051 (1.15)	8234 (3.42)
P _{2-gram}	7960 (0.00)	7960 (0.00)	7960 (0.00)	7960 (0.00)	8000 (0.50)	7975 (0.20)	7978 (0.23)	8002 (0.53)	7985 (0.32)	7960 (0.00)	7980 (0.26)	7971 (0.14)	7969 (0.12)	8051 (1.14)	8233 (3.43)
P _{4-gram}	8045 (0.00)	8045 (0.00)	8045 (0.00)	8045 (0.00)	8071 (0.32)	8054 (0.11)	8055 (0.12)	8072 (0.33)	8060 (0.19)	8045 (0.00)	8057 (0.15)	8051 (0.07)	8051 (0.07)	8105 (0.75)	8243 (2.46)
P _{8-gram}	8228 (0.00)	8228 (0.00)	8228 (0.00)	8228 (0.00)	8237 (0.10)	8235 (0.09)	8235 (0.09)	8238 (0.12)	8238 (0.12)	8228 (0.00)	8236 (0.10)	8234 (0.07)	8233 (0.06)	8243 (0.18)	8288 (0.72)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	EMA	ES _{ANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(W) *openthread* (context-insensitive).

B	16318 (2.15)	16385 (2.57)	16349 (2.34)	16461 (3.05)	20398 (27.69)	20314 (27.17)	20181 (26.33)	20335 (27.30)	20358 (27.44)	16452 (2.99)	20357 (27.44)	20223 (26.60)	20273 (26.91)	20486 (28.24)	21112 (32.16)
DF _{A+S/A}	16385 (2.00)	16399 (2.09)	16406 (2.13)	16483 (2.61)	20419 (27.11)	20348 (26.67)	20246 (26.03)	20390 (26.93)	20401 (27.00)	16495 (2.68)	20398 (26.98)	20262 (26.14)	20300 (26.37)	20519 (27.74)	21148 (31.66)
DF _{A+S/O}	16349 (2.24)	16406 (2.60)	16324 (2.09)	16486 (3.11)	20400 (27.58)	20329 (27.14)	20152 (26.03)	20321 (27.08)	20357 (27.31)	16460 (2.94)	20359 (27.32)	20213 (26.41)	20266 (26.74)	20475 (28.05)	21098 (31.95)
DF _{A+S/V}	16461 (1.72)	16483 (1.86)	16486 (1.88)	16482 (1.85)	20471 (26.50)	20398 (26.05)	20349 (25.74)	20474 (26.52)	20480 (26.55)	16568 (2.38)	20428 (26.23)	20338 (25.67)	20356 (25.79)	20591 (27.24)	21200 (31.00)
E _{1-ctx}	20398 (1.39)	20419 (1.49)	20400 (1.40)	20471 (1.75)	20628 (2.53)	20571 (2.25)	20501 (1.90)	20614 (2.54)	20629 (2.42)	20403 (1.42)	20606 (2.42)	20526 (2.03)	20559 (2.19)	20756 (3.17)	21300 (5.87)
E _{2-ctx}	20314 (1.63)	20348 (1.80)	20329 (1.70)	20398 (2.05)	20571 (2.91)	20418 (2.15)	20374 (1.93)	20509 (2.60)	20516 (2.63)	20310 (1.60)	20487 (2.49)	20419 (2.15)	20458 (2.34)	20665 (3.38)	21265 (6.38)
E _{3-ctx}	20181 (2.09)	20246 (2.42)	20152 (1.95)	20349 (2.94)	20501 (3.71)	20374 (3.07)	20142 (1.90)	20343 (2.91)	20386 (3.13)	20178 (2.08)	20421 (3.31)	20295 (2.67)	20358 (2.99)	20539 (3.90)	21149 (6.99)
E _{DDG}	20335 (1.73)	20390 (2.01)	20321 (1.66)	20474 (2.43)	20614 (3.13)	20509 (2.60)	20343 (1.77)	20453 (2.32)	20529 (2.70)	20334 (1.73)	20533 (2.72)	20440 (2.26)	20495 (2.53)	20670 (3.41)	21220 (6.16)
E _{LTO}	20358 (1.72)	20401 (1.93)	20357 (1.71)	20480 (2.33)	20629 (3.07)	20516 (2.51)	20386 (1.86)	20529 (2.57)	20497 (2.41)	20350 (1.68)	20547 (2.66)	20456 (2.21)	20511 (2.48)	20713 (3.49)	21298 (6.41)
EMA	16452 (2.33)	16495 (2.59)	16460 (2.38)	16568 (3.04)	20403 (26.90)	20310 (26.32)	20178 (25.50)	20334 (26.47)	20350 (26.57)	16503 (2.64)	20362 (26.65)	20220 (25.76)	20267 (26.05)	20488 (27.43)	21113 (31.31)
ES _{ANCOV}	20357 (1.53)	20398 (1.74)	20359 (1.54)	20428 (1.89)	20606 (2.77)	20487 (2.18)	20421 (1.85)	20533 (2.41)	20547 (2.48)	20362 (1.56)	20471 (2.10)	20462 (2.05)	20502 (2.26)	20720 (3.34)	21290 (6.19)
E _{coarse}	20223 (1.80)	20262 (2.00)	20213 (1.53)	20338 (2.38)	20526 (3.33)	20419 (2.79)	20295 (2.17)	20440 (2.90)	20456 (2.97)	20220 (1.79)	20462 (3.00)	20330 (2.34)	20402 (2.70)	20621 (3.81)	21221 (6.83)
P _{2-gram}	20273 (1.57)	20300 (1.70)	20266 (1.53)	20356 (1.98)	20458 (3.00)	20358 (2.49)	20495 (1.99)	20495 (2.68)	20511 (2.76)	20267 (1.53)	20502 (2.71)	20402 (2.21)	20416 (2.28)	20657 (3.49)	21252 (6.47)
P _{4-gram}	20486 (1.56)	20519 (1.73)	20475 (1.51)	20591 (2.08)	20756 (2.90)	20665 (2.45)	20539 (1.82)	20670 (2.48)	20713 (2.69)	20488 (1.58)	20720 (2.72)	20621 (2.23)	20657 (2.41)	20757 (2.91)	21245 (5.33)
P _{8-gram}	21112 (1.53)	21148 (1.71)	21098 (1.46)	21200 (1.95)	21300 (2.44)	21265 (2.27)	21149 (1.71)	21220 (2.05)	21298 (2.42)	21113 (1.53)	21290 (2.39)	21221 (2.06)	21252 (2.20)	21245 (2.17)	21377 (2.81)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	EMA	ES _{ANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(X) *openthread* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	3410 (5.86)	3596 (11.63)	3366 (4.47)	3682 (14.29)	5256 (63.13)	5258 (63.19)	5270 (63.57)	5083 (57.77)	5280 (63.89)	4026 (24.97)	5279 (63.87)	5256 (63.14)	5260 (63.27)	5285 (64.03)	5248 (62.91)
DF _{A+S/A}	3596 (1.81)	3670 (3.88)	3562 (0.82)	3733 (5.68)	5256 (48.78)	5258 (48.84)	5270 (49.18)	5114 (44.77)	5280 (49.47)	4036 (14.25)	5279 (49.45)	5256 (48.79)	5260 (48.90)	5285 (49.60)	5248 (48.57)
DF _{A+S/O}	3366 (10.86)	3562 (17.32)	3246 (6.92)	3660 (20.55)	5256 (73.11)	5258 (73.18)	5270 (73.58)	5064 (66.81)	5280 (73.92)	4019 (32.39)	5279 (73.90)	5256 (73.12)	5260 (73.26)	5285 (74.07)	5248 (72.87)
DF _{A+S/V}	3682 (0.82)	3733 (2.22)	3660 (0.21)	3752 (2.74)	5256 (43.90)	5258 (43.96)	5270 (44.29)	5126 (40.35)	5280 (44.57)	4046 (10.78)	5279 (44.56)	5256 (43.92)	5260 (44.03)	5285 (44.70)	5248 (43.71)
E _{1-ctx}	5256 (0.00)	5256 (0.00)	5256 (0.00)	5256 (0.00)	5286 (0.59)	5291 (0.68)	5294 (0.74)	5299 (0.82)	5300 (0.84)	5256 (0.00)	5299 (0.83)	5289 (0.64)	5290 (0.66)	5301 (0.86)	5289 (0.64)
E _{2-ctx}	5258 (0.00)	5258 (0.00)	5258 (0.00)	5258 (0.00)	5291 (0.64)	5292 (0.65)	5297 (0.74)	5300 (0.80)	5301 (0.83)	5258 (0.00)	5301 (0.83)	5291 (0.64)	5293 (0.67)	5303 (0.86)	5291 (0.64)
E _{3-ctx}	5270 (0.00)	5270 (0.00)	5270 (0.00)	5270 (0.00)	5294 (0.47)	5297 (0.52)	5296 (0.51)	5303 (0.63)	5304 (0.65)	5270 (0.00)	5304 (0.65)	5294 (0.46)	5294 (0.46)	5304 (0.65)	5294 (0.46)
E _{DDG}	5083 (6.77)	5114 (7.42)	5064 (6.38)	5126 (7.67)	5299 (11.30)	5300 (11.33)	5303 (11.39)	5253 (10.35)	5307 (11.48)	5162 (8.44)	5307 (11.48)	5299 (11.30)	5300 (11.33)	5307 (11.49)	5299 (11.31)
E _{LTO}	5280 (0.00)	5280 (0.00)	5280 (0.00)	5280 (0.00)	5300 (0.37)	5301 (0.41)	5304 (0.46)	5307 (0.52)	5305 (0.47)	5280 (0.00)	5305 (0.48)	5301 (0.40)	5302 (0.41)	5309 (0.55)	5302 (0.43)
E _{MA}	4026 (0.20)	4036 (0.44)	4019 (0.02)	4046 (0.68)	5256 (30.79)	5258 (30.84)	5270 (31.14)	5162 (28.47)	5280 (31.39)	4145 (3.15)	5279 (31.38)	5256 (30.80)	5260 (30.90)	5285 (31.51)	5248 (30.61)
E _{SANCOV}	5279 (0.00)	5279 (0.00)	5279 (0.00)	5279 (0.00)	5299 (0.42)	5301 (0.42)	5304 (0.46)	5307 (0.52)	5305 (0.49)	5279 (0.00)	5304 (0.46)	5300 (0.39)	5301 (0.41)	5308 (0.55)	5301 (0.42)
E _{coarse}	5256 (0.00)	5256 (0.00)	5256 (0.00)	5256 (0.00)	5289 (0.63)	5291 (0.68)	5294 (0.72)	5299 (0.81)	5301 (0.86)	5256 (0.00)	5300 (0.84)	5287 (0.60)	5289 (0.63)	5300 (0.84)	5287 (0.60)
P _{2-gram}	5260 (0.00)	5260 (0.00)	5260 (0.00)	5260 (0.00)	5290 (0.57)	5293 (0.62)	5294 (0.65)	5300 (0.76)	5302 (0.79)	5260 (0.00)	5301 (0.78)	5289 (0.55)	5289 (0.54)	5301 (0.78)	5288 (0.54)
P _{4-gram}	5285 (0.00)	5285 (0.00)	5285 (0.00)	5285 (0.00)	5301 (0.31)	5303 (0.34)	5304 (0.37)	5307 (0.43)	5309 (0.46)	5285 (0.00)	5308 (0.45)	5300 (0.31)	5301 (0.31)	5306 (0.41)	5300 (0.29)
P _{8-gram}	5248 (0.00)	5248 (0.00)	5248 (0.00)	5248 (0.00)	5289 (0.77)	5291 (0.82)	5294 (0.87)	5299 (0.96)	5302 (1.03)	5248 (0.00)	5301 (1.01)	5287 (0.74)	5288 (0.77)	5300 (0.98)	5282 (0.64)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(Y) *re2* (context-insensitive).

B	6239 (11.36)	6902 (23.19)	6147 (9.71)	7202 (28.55)	11149 (98.99)	10863 (93.88)	11029 (96.84)	10482 (87.08)	11236 (100.54)	8070 (44.03)	11137 (98.76)	11210 (100.08)	11250 (100.79)	11483 (104.94)	10996 (96.25)
DF _{A+S/A}	6902 (5.79)	7246 (11.05)	6808 (4.34)	7492 (14.83)	11254 (72.48)	10992 (68.46)	11149 (70.87)	10674 (63.59)	11345 (73.87)	8241 (26.30)	11242 (72.29)	11305 (73.27)	11340 (73.79)	11565 (77.24)	11097 (70.08)
DF _{A+S/O}	6147 (17.77)	6808 (30.44)	5861 (12.30)	7146 (36.91)	11127 (113.20)	10844 (107.76)	11015 (111.04)	10412 (99.50)	11222 (115.01)	8044 (54.12)	11108 (112.83)	11194 (114.47)	11234 (115.24)	11467 (119.69)	10973 (110.23)
DF _{A+S/V}	7202 (3.41)	7492 (7.57)	7146 (2.60)	7547 (8.35)	11329 (62.66)	11085 (59.16)	11225 (61.17)	10820 (55.35)	11410 (63.82)	8362 (20.05)	11327 (62.64)	11374 (63.30)	11408 (63.79)	11619 (66.82)	11189 (60.64)
E _{1-ctx}	11149 (1.50)	11254 (2.45)	11127 (1.30)	11329 (3.14)	11661 (6.16)	11549 (5.14)	11658 (6.13)	11559 (5.23)	11753 (7.00)	11327 (3.12)	11687 (6.40)	11754 (7.01)	11783 (7.27)	11892 (8.27)	11616 (5.75)
E _{2-ctx}	10863 (1.82)	10992 (3.02)	10844 (1.64)	11085 (3.90)	11549 (8.24)	11269 (5.62)	11408 (6.93)	11323 (6.13)	11563 (8.37)	11052 (3.59)	11470 (7.51)	11567 (8.41)	11590 (8.63)	11709 (9.74)	11415 (6.99)
E _{3-ctx}	11029 (1.76)	11149 (2.87)	11015 (1.63)	11225 (3.57)	11658 (7.57)	11408 (5.26)	11483 (5.95)	11446 (5.61)	11651 (7.50)	11201 (3.35)	11590 (6.94)	11663 (7.61)	11676 (7.73)	11797 (8.85)	11549 (6.56)
E _{DDG}	10482 (7.57)	10674 (9.54)	10412 (6.86)	10820 (11.04)	11559 (18.62)	11323 (16.21)	11446 (17.47)	11190 (14.84)	11597 (19.02)	10879 (11.64)	11486 (17.88)	11596 (19.01)	11611 (19.16)	11747 (20.56)	11420 (17.20)
E _{LTO}	11236 (1.56)	11345 (2.54)	11222 (1.43)	11410 (3.13)	11753 (6.23)	11563 (4.51)	11651 (5.30)	11597 (4.82)	11713 (5.87)	11412 (3.14)	11708 (5.83)	11798 (6.64)	11824 (6.87)	11909 (7.64)	11691 (5.67)
E _{MA}	8070 (2.56)	8241 (4.73)	8044 (2.23)	8362 (6.27)	11327 (43.96)	11052 (40.46)	11201 (42.36)	10879 (38.26)	11412 (45.03)	8517 (8.24)	11302 (43.64)	11351 (44.26)	11391 (44.77)	11598 (47.40)	11159 (41.82)
E _{SANCOV}	11137 (1.69)	11242 (2.65)	11108 (1.43)	11327 (3.43)	11687 (6.72)	11470 (4.74)	11590 (5.83)	11486 (4.88)	11708 (6.91)	11302 (3.21)	11589 (5.82)	11713 (6.96)	11729 (7.10)	11836 (8.08)	11580 (5.74)
E _{coarse}	11210 (1.59)	11305 (2.45)	11194 (1.44)	11374 (3.07)	11754 (6.51)	11567 (4.81)	11663 (5.69)	11596 (5.08)	11798 (6.91)	11351 (2.86)	11713 (6.14)	11745 (6.43)	11789 (6.83)	11894 (7.78)	11642 (5.50)
P _{2-gram}	11250 (1.47)	11340 (2.28)	11234 (1.32)	11408 (2.89)	11783 (6.28)	11590 (4.53)	11676 (5.31)	11611 (4.72)	11824 (6.65)	11391 (2.74)	11729 (5.79)	11789 (6.33)	11771 (6.17)	11913 (7.45)	11663 (5.19)
P _{4-gram}	11483 (1.19)	11565 (1.92)	11467 (1.05)	11619 (2.39)	11892 (4.80)	11709 (3.19)	11797 (3.97)	11747 (3.52)	11909 (4.95)	11598 (2.21)	11836 (4.31)	11894 (4.82)	11913 (4.99)	11958 (5.38)	11794 (3.94)
P _{8-gram}	10996 (1.75)	11097 (2.69)	10973 (1.54)	11189 (3.54)	11616 (7.49)	11415 (5.63)	11549 (6.87)	11420 (5.68)	11691 (8.19)	11159 (3.26)	11580 (7.16)	11642 (7.74)	11663 (7.92)	11794 (9.14)	11444 (5.90)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(Z) *re2* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	13483 (1.64)	13664 (3.00)	13541 (2.07)	14238 (7.33)	20250 (52.65)	23208 (74.95)	22779 (71.72)	25755 (94.15)	25449 (91.84)	14094 (6.24)	24899 (87.70)	23665 (78.40)	23703 (78.68)	23512 (77.24)	18585 (40.10)
DF _{A+S/A}	13664 (5.02)	13723 (5.47)	13633 (4.78)	14364 (10.40)	20325 (56.22)	23248 (78.68)	22821 (75.41)	25780 (98.14)	25482 (95.86)	13964 (7.33)	24934 (91.64)	23706 (82.20)	23743 (82.49)	23557 (81.06)	18686 (43.62)
DF _{A+S/O}	13541 (9.48)	13633 (10.23)	13305 (7.58)	14305 (15.67)	20268 (63.88)	23214 (87.70)	22792 (84.28)	25760 (108.28)	25454 (105.81)	13614 (10.08)	24904 (101.36)	23675 (91.42)	23711 (91.71)	23523 (90.19)	18629 (50.63)
DF _{A+S/V}	14238 (0.84)	14364 (1.73)	14305 (1.32)	14667 (3.88)	20376 (44.31)	23256 (64.70)	22846 (61.80)	25777 (82.56)	25473 (80.41)	14696 (4.08)	24936 (76.60)	23707 (67.90)	23745 (68.17)	23577 (66.98)	18826 (33.33)
E _{1-ctx}	20250 (0.05)	20325 (0.42)	20268 (0.14)	20376 (0.68)	21271 (5.10)	23655 (16.88)	23317 (15.21)	26003 (28.48)	25732 (27.14)	20380 (0.70)	25250 (24.76)	24210 (19.62)	24219 (19.67)	24048 (18.82)	20805 (2.80)
E _{2-ctx}	23208 (0.03)	23248 (0.20)	23214 (0.06)	23256 (0.24)	23655 (1.96)	24784 (6.82)	24772 (6.77)	26797 (15.50)	26587 (14.60)	23257 (0.24)	26207 (12.96)	25415 (9.55)	25442 (9.66)	25306 (9.07)	23373 (0.75)
E _{3-ctx}	22779 (0.03)	22821 (0.22)	22792 (0.09)	22846 (0.33)	23317 (2.39)	24772 (8.79)	24372 (7.03)	26681 (17.17)	26469 (16.24)	22857 (0.38)	26082 (14.54)	25243 (10.86)	25265 (10.95)	25130 (10.36)	22993 (0.97)
E _{DDG}	25755 (0.01)	25780 (0.10)	25760 (0.03)	25777 (0.09)	26003 (0.97)	26797 (4.05)	26681 (3.60)	27750 (7.75)	27876 (8.24)	25778 (0.10)	27683 (7.50)	27092 (5.20)	27036 (4.98)	27048 (5.03)	25851 (0.38)
E _{LTO}	25449 (0.01)	25482 (0.14)	25454 (0.03)	25473 (0.10)	25732 (1.12)	26587 (4.48)	26469 (4.02)	27876 (9.54)	27408 (7.70)	25471 (0.09)	27436 (7.81)	26919 (5.78)	26906 (5.73)	26838 (5.47)	25541 (0.37)
E _{MA}	14094 (23.59)	13964 (22.46)	13614 (19.39)	14696 (28.87)	20380 (78.71)	23257 (103.94)	22857 (100.44)	25778 (126.05)	25471 (123.36)	12810 (12.33)	24925 (118.66)	23705 (107.87)	23739 (108.17)	23570 (106.69)	18829 (65.11)
E _{SANCOV}	24899 (0.01)	24934 (0.15)	24904 (0.03)	24936 (0.16)	25250 (1.42)	26207 (5.26)	26082 (4.76)	27683 (11.19)	27436 (10.20)	24935 (0.15)	26950 (8.25)	26575 (6.74)	26571 (6.73)	26476 (6.34)	25008 (0.45)
E _{coarse}	23665 (0.02)	23706 (0.19)	23675 (0.06)	23707 (0.19)	24210 (2.32)	25415 (7.41)	25243 (6.69)	27092 (14.50)	26919 (13.77)	23705 (0.18)	26575 (12.31)	25563 (8.04)	25795 (9.02)	25681 (8.53)	23852 (0.81)
P _{2-gram}	23703 (0.02)	23743 (0.19)	23711 (0.05)	23745 (0.20)	24219 (2.20)	25442 (7.36)	25265 (6.61)	27036 (14.08)	26906 (13.54)	23739 (0.17)	26571 (12.12)	25795 (8.85)	25565 (7.88)	25711 (8.49)	23892 (0.81)
P _{4-gram}	23512 (0.02)	23557 (0.21)	23523 (0.06)	23577 (0.30)	24048 (2.30)	25306 (7.65)	25130 (6.90)	27048 (15.06)	26838 (14.17)	23570 (0.27)	26476 (12.63)	25681 (9.25)	25711 (9.37)	25419 (8.13)	23682 (0.74)
P _{8-gram}	18585 (0.12)	18686 (0.67)	18629 (0.36)	18826 (1.42)	20805 (12.08)	23373 (25.92)	22993 (23.87)	25851 (39.27)	25541 (37.60)	18829 (1.44)	25008 (34.73)	23852 (28.50)	23892 (28.71)	23682 (27.59)	19480 (4.95)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(AA) *sqlite3* (context-insensitive).

B	88585 (17.42)	96210 (27.53)	87643 (16.17)	110430 (46.38)	248756 (229.73)	337790 (347.74)	323842 (329.25)	321820 (326.58)	342156 (353.53)	105901 (40.37)	364941 (383.73)	338557 (348.76)	311222 (312.53)	311436 (312.81)	218292 (189.35)
DF _{A+S/A}	96210 (19.12)	101436 (25.58)	94113 (16.52)	115579 (43.10)	249208 (208.54)	337521 (317.88)	323547 (300.58)	322046 (298.72)	342601 (324.17)	110602 (36.93)	364723 (351.55)	338579 (319.19)	311326 (285.44)	311813 (286.05)	219251 (171.45)
DF _{A+S/O}	87643 (31.65)	94113 (41.37)	82697 (24.22)	108832 (63.48)	247682 (272.06)	336829 (405.97)	323090 (385.33)	320806 (381.90)	341110 (412.40)	100814 (51.44)	363841 (446.55)	337414 (406.85)	310339 (366.18)	310274 (366.08)	217387 (226.55)
DF _{A+S/V}	110430 (13.16)	115579 (18.44)	108832 (11.52)	125324 (28.42)	253370 (159.63)	339998 (248.40)	325998 (234.05)	325877 (233.93)	346494 (255.06)	124184 (27.25)	347327 (276.41)	341609 (250.05)	314415 (222.19)	315739 (223.54)	223772 (129.30)
E _{1-ctx}	248756 (3.18)	249208 (3.36)	247682 (2.73)	253370 (3.36)	286818 (18.96)	363897 (50.93)	349695 (50.93)	358135 (48.54)	378938 (57.17)	253730 (5.24)	389915 (61.72)	370059 (53.49)	343793 (42.59)	348946 (44.73)	278954 (15.70)
E _{2-ctx}	337790 (2.14)	337521 (2.06)	336829 (1.85)	339998 (2.80)	363897 (10.03)	396002 (19.74)	389888 (17.89)	403547 (28.14)	419653 (22.02)	341104 (3.14)	424168 (28.25)	408333 (23.47)	389134 (17.66)	395959 (19.72)	358347 (8.35)
E _{3-ctx}	323842 (1.97)	323547 (1.88)	323090 (1.73)	325998 (2.65)	349695 (10.11)	389888 (22.77)	372552 (17.31)	391849 (23.39)	410074 (29.12)	327570 (3.15)	415543 (30.85)	399040 (25.65)	376891 (18.68)	386561 (21.72)	343339 (8.11)
E _{DDG}	321820 (2.19)	322046 (2.26)	320806 (1.87)	325877 (3.48)	358135 (13.72)	403547 (28.14)	391849 (24.42)	384104 (21.96)	411607 (30.70)	325884 (3.48)	423215 (34.38)	408795 (29.80)	382935 (21.59)	391851 (24.42)	350727 (11.37)
E _{LTO}	342156 (2.14)	342601 (2.27)	341110 (1.83)	346494 (3.43)	378938 (13.12)	419653 (25.27)	410074 (22.41)	411607 (22.87)	423872 (26.53)	346428 (3.41)	439211 (31.11)	425415 (26.99)	403994 (20.60)	409769 (22.32)	371748 (10.97)
E _{MA}	105901 (32.59)	110602 (38.48)	100814 (26.22)	124184 (55.48)	253730 (217.68)	341104 (327.08)	327570 (310.13)	325884 (308.02)	346428 (333.75)	109271 (36.81)	368512 (22.75)	436127 (328.33)	430716 (294.79)	411545 (295.37)	223866 (180.29)
E _{SANCOV}	364941 (1.99)	364723 (1.93)	363841 (1.68)	367327 (2.66)	389915 (8.97)	424168 (18.54)	415543 (16.13)	423215 (18.28)	439211 (22.75)	368512 (2.99)	436127 (21.88)	430716 (20.37)	411545 (15.01)	417787 (16.76)	385368 (7.70)
E _{coarse}	338557 (2.38)	338579 (2.39)	337414 (2.03)	341609 (3.30)	370059 (11.91)	408333 (23.48)	399040 (20.67)	408795 (23.62)	425415 (28.65)	342103 (3.45)	430716 (30.25)	408176 (23.43)	395849 (19.71)	400903 (21.23)	364177 (10.13)
P _{2-gram}	311222 (2.18)	311326 (2.22)	310339 (1.89)	314415 (3.23)	343793 (12.88)	389134 (27.76)	376891 (23.74)	382935 (25.73)	403994 (32.64)	315318 (3.53)	411545 (35.12)	395849 (29.97)	365981 (20.16)	380874 (25.05)	336162 (10.37)
P _{4-gram}	311436 (2.72)	311813 (2.84)	310274 (2.34)	315739 (4.14)	348946 (15.09)	395959 (30.60)	386561 (27.50)	391851 (29.24)	409769 (35.15)	315774 (4.15)	417787 (37.80)	400903 (32.23)	380874 (25.62)	380230 (25.41)	340839 (12.42)
P _{8-gram}	218292 (3.20)	219251 (3.65)	217387 (2.77)	223772 (5.79)	278954 (31.88)	358347 (69.41)	343339 (62.32)	350727 (65.81)	371748 (75.75)	223866 (5.83)	385368 (82.19)	364177 (72.17)	336162 (58.92)	340839 (61.13)	253038 (19.63)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(AB) *sqlite3* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	2396 (2.93)	2430 (4.38)	2372 (1.87)	2485 (6.74)	2754 (18.30)	2758 (18.46)	2772 (19.05)	2753 (18.25)	2794 (20.00)	2512 (7.90)	2794 (20.02)	2756 (18.36)	2764 (18.70)	2745 (17.90)	2680 (15.13)
DF _{A+S/A}	2430 (2.47)	2442 (2.95)	2403 (1.31)	2495 (5.21)	2754 (16.13)	2758 (16.29)	2772 (16.86)	2757 (16.26)	2794 (17.80)	2521 (6.30)	2794 (17.81)	2756 (16.19)	2764 (16.53)	2745 (15.74)	2681 (13.04)
DF _{A+S/O}	2372 (7.79)	2403 (9.20)	2301 (4.56)	2470 (12.26)	2754 (25.15)	2758 (25.33)	2771 (25.96)	2740 (24.53)	2794 (26.97)	2480 (12.70)	2794 (26.98)	2755 (25.22)	2763 (25.60)	2745 (24.73)	2676 (21.61)
DF _{A+S/V}	2485 (1.12)	2495 (1.53)	2470 (0.50)	2519 (2.48)	2755 (12.11)	2759 (12.27)	2772 (12.80)	2766 (12.56)	2794 (13.70)	2553 (3.88)	2795 (13.72)	2757 (12.19)	2764 (12.47)	2746 (11.73)	2687 (9.35)
E _{1-ctx}	2754 (0.03)	2754 (0.03)	2754 (0.01)	2755 (0.07)	2778 (0.89)	2780 (0.96)	2788 (1.26)	2801 (1.72)	2806 (1.92)	2756 (0.10)	2807 (1.94)	2779 (0.91)	2782 (1.03)	2772 (0.68)	2759 (0.21)
E _{2-ctx}	2758 (0.01)	2758 (0.01)	2758 (0.00)	2759 (0.06)	2780 (0.80)	2779 (0.76)	2788 (1.10)	2802 (1.62)	2805 (1.73)	2759 (0.07)	2806 (1.76)	2782 (0.82)	2782 (0.89)	2772 (0.55)	2760 (0.10)
E _{3-ctx}	2772 (0.02)	2772 (0.02)	2771 (0.01)	2772 (0.03)	2788 (0.61)	2788 (0.61)	2793 (0.78)	2806 (1.26)	2809 (1.37)	2772 (0.02)	2810 (1.40)	2789 (0.65)	2788 (0.60)	2781 (0.37)	2773 (0.09)
E _{DDG}	2753 (9.24)	2757 (9.42)	2740 (8.73)	2766 (9.77)	2801 (11.14)	2802 (11.20)	2806 (11.35)	2783 (10.42)	2814 (11.68)	2764 (9.69)	2815 (11.70)	2800 (11.21)	2802 (11.21)	2799 (11.06)	2788 (10.64)
E _{LTO}	2794 (0.00)	2794 (0.00)	2794 (0.00)	2794 (0.03)	2806 (0.45)	2805 (0.42)	2809 (0.55)	2814 (0.73)	2812 (0.64)	2795 (0.03)	2812 (0.67)	2808 (0.50)	2806 (0.43)	2804 (0.36)	2796 (0.07)
E _{MA}	2512 (3.02)	2521 (3.39)	2480 (1.70)	2553 (4.70)	2756 (13.03)	2759 (13.17)	2772 (13.67)	2764 (13.36)	2795 (14.60)	2506 (2.76)	2795 (14.64)	2758 (13.12)	2765 (13.38)	2748 (12.69)	2696 (10.55)
E _{SANCOV}	2794 (0.01)	2794 (0.01)	2794 (0.00)	2795 (0.03)	2807 (0.46)	2806 (0.43)	2810 (0.57)	2815 (0.75)	2812 (0.66)	2795 (0.06)	2812 (0.65)	2808 (0.52)	2806 (0.44)	2804 (0.37)	2796 (0.09)
E _{coarse}	2756 (0.03)	2756 (0.03)	2755 (0.01)	2757 (0.09)	2779 (0.86)	2780 (0.92)	2789 (1.25)	2800 (1.65)	2808 (1.92)	2758 (0.12)	2808 (1.95)	2773 (0.67)	2783 (1.03)	2773 (0.65)	2759 (0.16)
P _{2-gram}	2764 (0.00)	2764 (0.01)	2763 (0.00)	2764 (0.02)	2782 (0.66)	2782 (0.67)	2788 (0.88)	2802 (1.41)	2806 (1.54)	2765 (0.05)	2806 (1.55)	2773 (0.72)	2781 (0.63)	2774 (0.37)	2765 (0.07)
P _{4-gram}	2745 (0.01)	2745 (0.01)	2745 (0.01)	2746 (0.05)	2772 (1.02)	2773 (1.04)	2781 (1.35)	2799 (1.98)	2804 (2.16)	2748 (0.12)	2804 (2.18)	2773 (1.03)	2774 (1.07)	2762 (0.64)	2749 (0.15)
P _{8-gram}	2680 (0.40)	2681 (0.41)	2676 (0.23)	2687 (0.66)	2759 (3.35)	2760 (3.39)	2773 (3.89)	2788 (4.43)	2796 (4.72)	2696 (1.97)	2796 (4.74)	2759 (3.36)	2765 (3.58)	2749 (2.95)	2703 (1.24)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(AC) *stb* (context-insensitive).

B	4437 (4.84)	4476 (5.75)	4394 (3.81)	4594 (8.54)	4997 (18.06)	5036 (18.98)	5030 (18.84)	4997 (18.07)	5080 (20.03)	4619 (9.13)	5106 (20.64)	4987 (17.83)	5029 (18.81)	5012 (18.42)	4888 (15.48)
DF _{A+S/A}	4476 (4.42)	4462 (4.09)	4416 (3.03)	4579 (6.82)	4961 (15.73)	5006 (16.78)	5006 (16.78)	4975 (16.06)	5051 (17.84)	4603 (7.39)	5082 (18.55)	4961 (15.75)	4999 (16.62)	4988 (16.37)	4856 (13.29)
DF _{A+S/O}	4394 (8.84)	4416 (9.40)	4264 (5.63)	4553 (12.78)	4968 (23.06)	5014 (24.20)	5013 (24.19)	4955 (22.75)	5055 (25.21)	4553 (12.79)	5090 (26.10)	4968 (23.06)	5006 (24.00)	4991 (23.63)	4857 (20.32)
DF _{A+S/V}	4594 (2.72)	4579 (2.38)	4553 (1.80)	4639 (3.72)	4991 (11.59)	5029 (12.45)	5034 (12.57)	5015 (12.12)	5074 (13.45)	4679 (4.62)	5102 (14.07)	4997 (11.72)	5024 (12.34)	5013 (12.09)	4892 (9.38)
E _{1-ctx}	4997 (2.07)	4961 (1.33)	4968 (1.47)	4991 (1.94)	4991 (1.95)	5038 (2.89)	5037 (2.89)	5044 (3.03)	5063 (3.42)	4965 (1.41)	5100 (4.18)	4999 (2.12)	5023 (2.60)	5028 (2.70)	4982 (1.77)
E _{2-ctx}	5036 (1.28)	5006 (0.68)	5014 (0.84)	5029 (1.15)	5038 (1.33)	5060 (1.77)	5072 (2.00)	5082 (2.21)	5094 (2.45)	5013 (0.82)	5119 (2.95)	5043 (1.43)	5056 (1.68)	5059 (1.75)	5017 (0.90)
E _{3-ctx}	5030 (2.23)	5006 (1.74)	5013 (1.89)	5034 (2.32)	5037 (2.38)	5072 (3.08)	5053 (2.69)	5067 (2.99)	5093 (3.50)	5008 (1.78)	5121 (4.08)	5029 (2.20)	5054 (2.71)	5058 (2.79)	5028 (2.18)
E _{DDG}	4997 (11.19)	4975 (10.69)	4955 (10.26)	5015 (11.58)	5044 (12.23)	5082 (13.09)	5067 (12.75)	5016 (11.61)	5088 (13.21)	4990 (11.03)	5119 (13.91)	5042 (12.18)	5065 (12.70)	5076 (12.95)	5038 (12.10)
E _{LTO}	5080 (1.55)	5051 (0.96)	5055 (1.04)	5074 (1.42)	5063 (1.20)	5094 (1.82)	5093 (1.79)	5088 (1.69)	5088 (1.70)	5080 (0.94)	5122 (2.39)	5071 (1.35)	5079 (1.51)	5091 (1.76)	5059 (1.13)
E _{MA}	4619 (4.89)	4603 (4.54)	4553 (3.40)	4679 (6.26)	4965 (12.75)	5013 (13.84)	5008 (13.72)	4990 (13.32)	5050 (14.68)	4583 (4.08)	5087 (15.52)	4964 (12.73)	5002 (13.60)	4990 (13.31)	4883 (10.90)
E _{SANCOV}	5106 (1.19)	5082 (0.70)	5090 (0.88)	5102 (1.10)	5100 (1.07)	5119 (1.45)	5121 (1.49)	5119 (1.45)	5122 (1.51)	5087 (0.81)	5140 (1.86)	5106 (1.20)	5109 (1.24)	5123 (1.53)	5091 (0.88)
E _{coarse}	4987 (2.50)	4961 (1.97)	4968 (2.10)	4997 (2.69)	4999 (2.75)	5043 (3.65)	5029 (3.35)	5042 (3.62)	5071 (4.21)	4964 (2.02)	5106 (4.95)	4979 (2.33)	5028 (3.34)	5024 (3.26)	4987 (2.49)
P _{2-gram}	5029 (1.62)	4999 (1.02)	5006 (1.15)	5024 (1.53)	5023 (1.50)	5056 (2.16)	5054 (2.12)	5065 (2.35)	5079 (2.63)	5002 (1.09)	5109 (3.23)	5028 (1.61)	5036 (1.77)	5045 (1.95)	5011 (1.26)
P _{4-gram}	5012 (1.70)	4988 (1.21)	4991 (1.27)	5013 (1.73)	5028 (2.02)	5059 (2.66)	5058 (2.63)	5076 (3.01)	5091 (3.30)	4990 (1.25)	5123 (3.96)	5024 (1.95)	5045 (2.38)	5038 (2.22)	5003 (1.51)
P _{8-gram}	4888 (2.10)	4856 (1.43)	4857 (1.46)	4892 (2.19)	4982 (4.07)	5017 (4.79)	5028 (5.02)	5038 (5.23)	5059 (5.68)	4883 (2.01)	5091 (6.34)	4987 (4.16)	5011 (4.67)	5003 (4.50)	4899 (2.32)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(AD) *stb* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	1558 (2.48)	1596 (4.95)	1557 (2.38)	1569 (3.17)	1897 (24.78)	1902 (25.10)	1904 (25.25)	1903 (25.15)	1908 (25.47)	1566 (3.03)	1908 (25.51)	1902 (25.08)	1901 (25.02)	1905 (25.32)	1906 (25.39)
DF _{A+S/A}	1596 (1.33)	1617 (2.69)	1595 (1.26)	1597 (1.41)	1897 (20.46)	1902 (20.77)	1904 (20.92)	1903 (20.82)	1908 (21.13)	1598 (1.46)	1908 (21.17)	1902 (20.75)	1901 (20.69)	1905 (20.98)	1906 (21.06)
DF _{A+S/O}	1557 (3.55)	1595 (6.08)	1549 (3.06)	1568 (4.27)	1897 (26.19)	1902 (26.52)	1904 (26.67)	1903 (26.57)	1908 (26.89)	1563 (3.96)	1908 (26.94)	1902 (26.49)	1901 (26.44)	1905 (26.73)	1906 (26.82)
DF _{A+S/V}	1569 (1.75)	1597 (3.59)	1568 (1.68)	1561 (1.27)	1897 (23.07)	1902 (23.38)	1904 (23.53)	1903 (23.43)	1908 (23.74)	1563 (1.40)	1908 (23.79)	1902 (23.37)	1901 (23.32)	1905 (23.60)	1906 (23.67)
E _{1-ctx}	1897 (0.01)	1897 (0.00)	1897 (0.00)	1897 (0.01)	1904 (0.35)	1904 (0.39)	1906 (0.47)	1905 (0.41)	1909 (0.61)	1897 (0.00)	1909 (0.65)	1905 (0.43)	1905 (0.40)	1908 (0.56)	1909 (0.64)
E _{2-ctx}	1902 (0.01)	1902 (0.00)	1902 (0.00)	1902 (0.01)	1904 (0.13)	1904 (0.10)	1906 (0.20)	1905 (0.14)	1909 (0.35)	1902 (0.00)	1909 (0.39)	1905 (0.16)	1905 (0.16)	1907 (0.29)	1909 (0.37)
E _{3-ctx}	1904 (0.00)	1904 (0.00)	1904 (0.00)	1904 (0.01)	1906 (0.09)	1906 (0.08)	1907 (0.12)	1906 (0.09)	1910 (0.29)	1904 (0.00)	1910 (0.32)	1906 (0.12)	1906 (0.10)	1909 (0.25)	1910 (0.30)
E _{DDG}	1903 (0.01)	1903 (0.00)	1903 (0.00)	1903 (0.01)	1905 (0.11)	1905 (0.11)	1906 (0.17)	1904 (0.09)	1909 (0.33)	1903 (0.00)	1910 (0.37)	1905 (0.13)	1905 (0.12)	1908 (0.26)	1909 (0.32)
E _{LTO}	1908 (0.00)	1908 (0.00)	1908 (0.00)	1908 (0.01)	1909 (0.06)	1909 (0.06)	1910 (0.11)	1909 (0.07)	1909 (0.07)	1908 (0.00)	1910 (0.11)	1910 (0.11)	1909 (0.09)	1910 (0.12)	1910 (0.16)
E _{MA}	1566 (2.39)	1598 (4.43)	1563 (2.15)	1563 (2.18)	1897 (24.00)	1902 (24.32)	1904 (24.47)	1903 (24.36)	1908 (24.68)	1556 (1.69)	1908 (24.73)	1902 (24.30)	1901 (24.25)	1905 (24.54)	1906 (24.62)
E _{SANCOV}	1908 (0.00)	1908 (0.00)	1908 (0.00)	1908 (0.01)	1909 (0.06)	1909 (0.06)	1910 (0.11)	1910 (0.07)	1910 (0.07)	1908 (0.00)	1910 (0.09)	1910 (0.10)	1910 (0.08)	1910 (0.11)	1911 (0.15)
E _{coarse}	1902 (0.01)	1902 (0.00)	1902 (0.00)	1902 (0.02)	1905 (0.19)	1905 (0.18)	1906 (0.26)	1905 (0.19)	1910 (0.42)	1902 (0.01)	1910 (0.45)	1905 (0.20)	1905 (0.19)	1908 (0.34)	1909 (0.41)
P _{2-gram}	1901 (0.01)	1901 (0.00)	1901 (0.00)	1901 (0.02)	1905 (0.20)	1905 (0.22)	1906 (0.29)	1905 (0.22)	1909 (0.45)	1901 (0.01)	1910 (0.48)	1905 (0.23)	1904 (0.20)	1908 (0.37)	1909 (0.45)
P _{4-gram}	1905 (0.01)	1905 (0.00)	1905 (0.01)	1905 (0.02)	1908 (0.13)	1907 (0.12)	1909 (0.20)	1908 (0.13)	1910 (0.24)	1905 (0.01)	1910 (0.28)	1908 (0.15)	1908 (0.14)	1908 (0.18)	1910 (0.24)
P _{8-gram}	1906 (0.00)	1906 (0.00)	1906 (0.00)	1906 (0.00)	1909 (0.14)	1909 (0.13)	1910 (0.18)	1909 (0.12)	1910 (0.21)	1906 (0.00)	1911 (0.24)	1909 (0.15)	1909 (0.15)	1910 (0.17)	1910 (0.16)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(AE) *vorbis* (context-insensitive).

B	2326 (4.37)	2397 (7.57)	2313 (3.81)	2345 (5.22)	3612 (62.09)	3629 (62.84)	3630 (62.88)	3643 (63.49)	3676 (64.95)	2330 (4.55)	3634 (63.08)	3606 (61.82)	3617 (62.30)	3663 (64.37)	3643 (63.48)
DF _{A+S/A}	2397 (3.34)	2438 (5.09)	2383 (2.72)	2400 (3.46)	3618 (55.98)	3639 (56.90)	3637 (56.81)	3655 (57.59)	3680 (58.66)	2391 (3.09)	3644 (57.08)	3618 (55.99)	3621 (56.11)	3671 (58.24)	3649 (57.32)
DF _{A+S/O}	2313 (6.08)	2383 (9.26)	2287 (4.88)	2332 (6.92)	3614 (65.71)	3632 (66.55)	3632 (66.54)	3647 (67.22)	3676 (68.58)	2312 (6.02)	3637 (66.77)	3610 (65.53)	3618 (65.91)	3665 (68.07)	3645 (67.14)
DF _{A+S/V}	2345 (3.93)	2400 (6.37)	2332 (3.34)	2333 (3.40)	3612 (60.09)	3629 (60.84)	3630 (60.88)	3643 (61.47)	3675 (62.90)	2328 (3.19)	3634 (61.08)	3606 (59.84)	3617 (60.31)	3663 (62.36)	3643 (61.47)
E _{1-ctx}	3612 (0.07)	3618 (0.23)	3614 (0.11)	3612 (0.07)	3655 (1.27)	3669 (1.64)	3666 (1.56)	3683 (2.02)	3701 (2.53)	3612 (0.07)	3673 (1.75)	3656 (1.28)	3657 (1.30)	3696 (2.40)	3678 (1.89)
E _{2-ctx}	3629 (0.09)	3639 (0.38)	3632 (0.18)	3629 (0.09)	3669 (1.19)	3670 (1.22)	3673 (1.32)	3680 (1.49)	3706 (2.23)	3629 (0.10)	3679 (1.46)	3660 (0.94)	3671 (1.26)	3698 (1.99)	3687 (1.69)
E _{3-ctx}	3630 (0.09)	3637 (0.31)	3632 (0.15)	3630 (0.09)	3666 (1.09)	3673 (1.30)	3671 (1.24)	3684 (1.60)	3706 (2.18)	3630 (0.11)	3678 (1.43)	3662 (0.99)	3667 (1.12)	3700 (2.04)	3684 (1.60)
E _{DDG}	3643 (0.13)	3655 (0.47)	3647 (0.23)	3643 (0.13)	3683 (1.21)	3680 (1.13)	3684 (1.26)	3682 (1.20)	3712 (2.02)	3644 (0.15)	3689 (1.40)	3669 (0.84)	3686 (1.29)	3703 (1.77)	3698 (1.63)
E _{LTO}	3676 (0.10)	3680 (0.22)	3676 (0.12)	3675 (0.09)	3701 (0.78)	3706 (0.93)	3706 (0.91)	3712 (1.09)	3718 (1.25)	3676 (0.10)	3705 (0.88)	3699 (0.73)	3702 (0.80)	3719 (1.27)	3710 (1.02)
E _{MA}	2330 (4.72)	2391 (7.48)	2312 (3.92)	2328 (4.65)	3612 (62.36)	3629 (63.13)	3630 (63.17)	3644 (63.78)	3676 (65.23)	2304 (3.56)	3635 (63.38)	3607 (62.11)	3617 (62.57)	3663 (64.67)	3644 (63.77)
E _{SANCOV}	3634 (0.09)	3644 (0.35)	3637 (0.17)	3634 (0.09)	3673 (1.15)	3679 (1.31)	3678 (1.31)	3689 (1.61)	3705 (2.03)	3635 (0.11)	3676 (1.23)	3667 (0.99)	3675 (1.21)	3701 (1.92)	3684 (1.47)
E _{coarse}	3606 (0.11)	3618 (0.44)	3610 (0.21)	3606 (0.11)	3656 (1.49)	3660 (1.60)	3662 (1.66)	3669 (1.86)	3699 (2.69)	3607 (0.12)	3667 (1.80)	3645 (1.18)	3658 (1.54)	3689 (2.39)	3675 (2.02)
P _{2-gram}	3617 (0.05)	3621 (0.17)	3618 (0.08)	3617 (0.05)	3657 (1.15)	3671 (1.55)	3667 (1.44)	3686 (1.95)	3702 (2.40)	3617 (0.05)	3675 (1.65)	3658 (1.18)	3655 (1.10)	3697 (2.25)	3678 (1.74)
P _{4-gram}	3663 (0.08)	3671 (0.28)	3665 (0.13)	3663 (0.08)	3696 (0.98)	3698 (1.02)	3700 (1.09)	3703 (1.16)	3719 (1.60)	3663 (0.09)	3701 (1.10)	3689 (0.77)	3697 (0.99)	3712 (1.42)	3706 (1.25)
P _{8-gram}	3643 (0.07)	3649 (0.23)	3645 (0.11)	3643 (0.06)	3678 (1.01)	3687 (1.27)	3684 (1.20)	3698 (1.57)	3710 (1.90)	3644 (0.07)	3684 (1.19)	3675 (0.94)	3678 (1.02)	3706 (1.79)	3687 (1.26)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(AF) *vorbis* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	1769 (2.35)	1802 (4.24)	1768 (2.26)	1851 (7.06)	2402 (38.98)	2363 (36.71)	2411 (39.47)	2399 (38.76)	2411 (39.50)	2022 (16.99)	2407 (39.27)	2374 (37.36)	2349 (35.91)	2428 (40.47)	2368 (37.00)
DF _{A+S/A}	1802 (1.60)	1820 (2.60)	1800 (1.49)	1863 (5.02)	2403 (35.50)	2366 (33.40)	2411 (35.94)	2400 (35.31)	2412 (36.02)	2024 (14.14)	2410 (35.90)	2375 (33.90)	2353 (32.67)	2428 (36.92)	2369 (33.54)
DF _{A+S/O}	1768 (4.27)	1800 (6.18)	1751 (3.29)	1851 (9.19)	2402 (41.71)	2364 (39.44)	2411 (42.22)	2399 (41.50)	2411 (42.24)	2026 (19.49)	2407 (42.00)	2374 (40.06)	2350 (38.61)	2428 (43.24)	2368 (39.70)
DF _{A+S/V}	1851 (0.63)	1863 (1.29)	1851 (0.65)	1880 (2.25)	2403 (30.65)	2365 (28.59)	2411 (31.13)	2401 (30.54)	2412 (31.18)	2035 (10.64)	2409 (30.99)	2376 (29.19)	2350 (27.80)	2428 (32.05)	2369 (28.82)
E _{1-ctx}	2402 (0.01)	2403 (0.04)	2402 (0.00)	2403 (0.02)	2448 (1.89)	2446 (1.84)	2457 (2.29)	2454 (2.17)	2461 (2.46)	2413 (0.46)	2481 (3.26)	2445 (1.76)	2440 (1.59)	2461 (2.44)	2441 (1.62)
E _{2-ctx}	2363 (0.03)	2366 (0.14)	2364 (0.06)	2365 (0.10)	2446 (3.55)	2430 (2.86)	2451 (3.74)	2446 (3.53)	2455 (3.92)	2395 (1.38)	2468 (4.48)	2438 (3.19)	2426 (2.68)	2455 (3.93)	2432 (2.96)
E _{3-ctx}	2411 (0.02)	2411 (0.03)	2411 (0.02)	2411 (0.04)	2457 (1.94)	2451 (1.67)	2456 (1.90)	2459 (2.01)	2468 (2.40)	2417 (0.29)	2485 (3.09)	2450 (1.64)	2447 (1.53)	2463 (2.18)	2446 (1.46)
E _{DDG}	2399 (0.06)	2400 (0.11)	2399 (0.10)	2401 (0.14)	2454 (2.38)	2446 (2.04)	2459 (2.57)	2450 (2.19)	2463 (2.76)	2417 (0.83)	2482 (3.54)	2448 (2.14)	2442 (1.86)	2461 (2.66)	2443 (1.91)
E _{LTO}	2411 (0.03)	2412 (0.07)	2411 (0.03)	2412 (0.07)	2461 (2.10)	2455 (1.84)	2468 (2.39)	2463 (2.19)	2463 (2.18)	2424 (0.55)	2482 (2.95)	2456 (1.90)	2450 (1.65)	2471 (2.49)	2451 (1.69)
E _{MA}	2022 (0.07)	2024 (0.17)	2026 (0.24)	2035 (0.68)	2413 (19.42)	2395 (18.51)	2417 (19.62)	2417 (19.61)	2424 (19.94)	2048 (1.35)	2431 (20.27)	2387 (18.11)	2386 (18.07)	2433 (20.39)	2374 (17.47)
E _{SANCOV}	2407 (0.06)	2410 (0.18)	2407 (0.05)	2409 (0.12)	2481 (3.10)	2468 (3.29)	2485 (3.17)	2482 (3.16)	2482 (3.16)	2431 (1.02)	2476 (2.93)	2472 (2.74)	2461 (2.30)	2491 (3.53)	2465 (2.47)
E _{coarse}	2374 (0.05)	2375 (0.06)	2374 (0.04)	2376 (0.10)	2445 (3.00)	2438 (2.72)	2450 (3.23)	2448 (3.16)	2456 (3.50)	2387 (0.57)	2472 (4.14)	2430 (2.37)	2431 (2.43)	2456 (3.48)	2429 (2.33)
P _{2-gram}	2349 (0.02)	2353 (0.18)	2350 (0.04)	2350 (0.06)	2440 (3.90)	2426 (3.28)	2447 (4.18)	2442 (3.96)	2450 (4.32)	2386 (1.58)	2461 (4.79)	2431 (3.49)	2414 (2.77)	2452 (4.40)	2427 (3.31)
P _{4-gram}	2428 (0.00)	2428 (0.01)	2428 (0.00)	2428 (0.01)	2461 (1.34)	2455 (1.43)	2463 (1.43)	2461 (1.35)	2471 (1.75)	2433 (0.20)	2491 (2.58)	2456 (1.14)	2452 (0.99)	2463 (1.42)	2450 (0.92)
P _{8-gram}	2368 (0.01)	2369 (0.02)	2368 (0.01)	2369 (0.04)	2441 (3.09)	2432 (2.72)	2446 (3.28)	2443 (3.16)	2451 (3.52)	2374 (0.25)	2465 (4.11)	2429 (2.56)	2427 (2.48)	2450 (3.48)	2413 (1.90)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(AG) *woff2* (context-insensitive).

B	2305 (2.79)	2346 (4.62)	2309 (2.98)	2413 (7.60)	3099 (38.21)	3056 (36.29)	3106 (38.51)	3097 (38.13)	3112 (38.78)	2597 (15.82)	3126 (39.43)	3054 (36.20)	3033 (35.27)	3125 (39.39)	3054 (36.20)
DF _{A+S/A}	2346 (2.28)	2365 (3.12)	2347 (2.31)	2423 (5.65)	3098 (35.06)	3056 (33.27)	3103 (35.29)	3095 (34.95)	3110 (35.58)	2597 (13.24)	3127 (36.34)	3051 (33.04)	3033 (32.25)	3122 (36.14)	3052 (33.07)
DF _{A+S/O}	2309 (4.67)	2347 (6.37)	2292 (3.91)	2413 (9.36)	3099 (40.46)	3056 (38.51)	3105 (40.73)	3095 (40.30)	3110 (40.99)	2601 (17.89)	3125 (41.68)	3053 (38.39)	3032 (37.45)	3124 (41.60)	3054 (38.43)
DF _{A+S/V}	2413 (1.20)	2423 (1.65)	2413 (1.20)	2448 (2.70)	3100 (30.06)	3058 (28.29)	3107 (30.33)	3100 (30.02)	3114 (30.61)	2613 (9.61)	3130 (31.29)	3057 (28.25)	3035 (27.32)	3126 (31.12)	3057 (28.22)
E _{1-ctx}	3099 (0.33)	3098 (0.29)	3099 (0.32)	3100 (0.38)	3157 (2.20)	3156 (2.17)	3166 (2.51)	3165 (2.45)	3177 (2.85)	3113 (0.79)	3215 (4.08)	3149 (1.93)	3145 (1.80)	3170 (2.64)	3147 (1.89)
E _{2-ctx}	3056 (0.35)	3056 (0.36)	3056 (0.34)	3058 (0.42)	3156 (3.63)	3137 (3.00)	3159 (3.75)	3155 (3.60)	3170 (4.08)	3091 (1.51)	3202 (5.13)	3140 (3.12)	3128 (2.71)	3164 (3.88)	3137 (3.01)
E _{3-ctx}	3106 (0.49)	3103 (0.40)	3105 (0.46)	3107 (0.53)	3166 (2.45)	3159 (2.23)	3162 (2.32)	3168 (2.49)	3181 (2.94)	3118 (0.88)	3219 (4.17)	3152 (1.97)	3151 (1.94)	3170 (2.57)	3152 (1.99)
E _{DDG}	3097 (0.47)	3095 (0.40)	3095 (0.46)	3100 (0.58)	3165 (2.66)	3155 (2.35)	3168 (2.76)	3157 (2.43)	3178 (3.10)	3119 (1.19)	3216 (4.34)	3152 (2.24)	3144 (2.00)	3168 (2.84)	3150 (2.18)
E _{LTO}	3112 (0.51)	3110 (0.44)	3110 (0.46)	3114 (0.58)	3177 (2.61)	3170 (2.39)	3181 (2.76)	3178 (2.66)	3180 (2.71)	3130 (1.09)	3217 (3.91)	3164 (2.22)	3159 (2.04)	3184 (2.84)	3163 (2.18)
E _{MA}	2597 (0.41)	2597 (0.42)	2601 (0.55)	2613 (1.03)	3113 (20.38)	3091 (19.52)	3118 (20.55)	3119 (20.60)	3130 (21.01)	2628 (1.60)	3154 (21.94)	3073 (18.81)	3075 (18.88)	3134 (21.16)	3063 (18.44)
E _{SANCOV}	3126 (0.42)	3127 (0.44)	3125 (0.39)	3130 (0.53)	3215 (3.26)	3202 (2.84)	3219 (3.41)	3216 (3.31)	3217 (3.33)	3154 (1.30)	3218 (3.35)	3199 (2.74)	3188 (2.39)	3226 (3.61)	3195 (2.61)
E _{coarse}	3054 (0.58)	3051 (0.50)	3053 (0.55)	3057 (0.69)	3149 (3.70)	3140 (3.43)	3152 (3.80)	3152 (3.80)	3164 (4.23)	3073 (1.21)	3199 (5.35)	3121 (2.79)	3125 (2.93)	3158 (4.00)	3127 (2.97)
P _{2-gram}	3033 (0.54)	3033 (0.55)	3032 (0.51)	3035 (0.62)	3145 (4.24)	3128 (3.69)	3151 (4.44)	3144 (4.23)	3159 (4.72)	3075 (1.92)	3188 (5.67)	3125 (3.59)	3106 (2.96)	3154 (4.57)	3124 (3.57)
P _{4-gram}	3125 (0.44)	3122 (0.34)	3124 (0.38)	3126 (0.45)	3170 (1.88)	3164 (1.66)	3170 (1.87)	3168 (1.82)	3184 (2.31)	3134 (0.70)	3226 (3.66)	3158 (1.47)	3154 (1.37)	3167 (1.78)	3156 (1.43)
P _{8-gram}	3054 (0.38)	3052 (0.31)	3054 (0.37)	3057 (0.47)	3147 (3.45)	3137 (3.11)	3152 (3.61)	3150 (3.53)	3163 (3.97)	3063 (0.68)	3195 (5.00)	3127 (2.76)	3124 (2.70)	3156 (3.74)	3112 (2.28)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(AH) *woff2* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

B	152 (7.46)	171 (20.52)	166 (17.43)	194 (36.76)	542 (282.98)	545 (284.82)	547 (286.02)	542 (282.84)	543 (283.33)	492 (247.25)	546 (285.24)	544 (284.25)	544 (284.18)	546 (285.73)	552 (289.62)
DF _{A+S/A}	171 (3.37)	184 (11.55)	183 (10.58)	203 (22.86)	542 (228.47)	545 (230.04)	547 (231.07)	542 (228.35)	543 (228.77)	492 (197.82)	546 (230.41)	544 (229.56)	544 (229.50)	546 (230.83)	552 (234.16)
DF _{A+S/O}	166 (6.11)	183 (16.51)	178 (13.36)	202 (29.00)	542 (246.08)	545 (247.73)	547 (248.82)	542 (245.95)	543 (246.39)	492 (213.78)	546 (248.12)	544 (247.22)	544 (247.16)	546 (248.56)	552 (252.07)
DF _{A+S/V}	194 (1.33)	203 (6.15)	202 (5.78)	214 (12.05)	542 (183.78)	545 (185.14)	547 (186.03)	542 (183.67)	543 (184.04)	492 (157.30)	546 (185.45)	544 (184.72)	544 (184.67)	546 (185.82)	552 (188.70)
E _{1-ctx}	542 (0.00)	542 (0.00)	542 (0.00)	542 (0.00)	544 (0.40)	547 (0.80)	547 (0.93)	545 (0.52)	545 (0.57)	542 (0.00)	547 (0.81)	546 (0.76)	546 (0.68)	549 (1.15)	554 (2.10)
E _{2-ctx}	545 (0.00)	545 (0.00)	545 (0.00)	545 (0.00)	547 (0.32)	548 (0.64)	549 (0.75)	547 (0.43)	547 (0.46)	545 (0.00)	548 (0.64)	548 (0.61)	548 (0.54)	550 (0.94)	554 (1.74)
E _{3-ctx}	547 (0.00)	547 (0.00)	547 (0.00)	547 (0.00)	547 (0.14)	549 (0.43)	549 (0.50)	548 (0.25)	548 (0.27)	547 (0.00)	549 (0.42)	549 (0.41)	548 (0.32)	550 (0.69)	555 (1.49)
E _{DDG}	542 (0.00)	542 (0.00)	542 (0.00)	542 (0.00)	545 (0.56)	547 (0.94)	548 (1.08)	546 (0.63)	546 (0.72)	542 (0.00)	547 (0.97)	547 (0.87)	547 (0.82)	549 (1.27)	554 (2.17)
E _{LTO}	543 (0.00)	543 (0.00)	543 (0.00)	543 (0.00)	545 (0.47)	547 (0.85)	548 (0.97)	546 (0.59)	546 (0.61)	543 (0.00)	547 (0.85)	547 (0.82)	547 (0.74)	549 (1.19)	554 (2.09)
E _{MA}	492 (0.00)	492 (0.00)	492 (0.00)	492 (0.00)	542 (10.29)	545 (10.82)	547 (11.17)	542 (10.25)	543 (10.39)	497 (0.98)	546 (10.94)	544 (10.66)	544 (10.64)	546 (11.08)	552 (12.20)
E _{SANCOV}	546 (0.00)	546 (0.00)	546 (0.00)	546 (0.00)	547 (0.22)	548 (0.53)	549 (0.62)	547 (0.34)	547 (0.35)	546 (0.00)	548 (0.51)	548 (0.51)	548 (0.43)	550 (0.83)	554 (1.64)
E _{coarse}	544 (0.00)	544 (0.00)	544 (0.00)	544 (0.00)	546 (0.43)	548 (0.75)	549 (0.87)	547 (0.50)	547 (0.58)	544 (0.00)	548 (0.77)	548 (0.65)	548 (0.63)	550 (1.04)	554 (1.78)
P _{2-gram}	544 (0.00)	544 (0.00)	544 (0.00)	544 (0.00)	546 (0.36)	548 (0.70)	548 (0.80)	547 (0.47)	547 (0.52)	544 (0.00)	548 (0.71)	548 (0.65)	548 (0.56)	549 (0.97)	554 (1.83)
P _{4-gram}	546 (0.00)	546 (0.00)	546 (0.00)	546 (0.00)	549 (0.43)	550 (0.70)	550 (0.77)	549 (0.51)	549 (0.56)	546 (0.00)	550 (0.70)	550 (0.65)	549 (0.57)	551 (0.87)	555 (1.63)
P _{8-gram}	552 (0.00)	552 (0.00)	552 (0.00)	552 (0.00)	554 (0.36)	554 (0.49)	555 (0.55)	554 (0.40)	554 (0.44)	552 (0.00)	554 (0.50)	554 (0.38)	554 (0.41)	555 (0.61)	556 (0.82)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(A1) *zlib* (context-insensitive).

B	164 (6.88)	183 (19.05)	178 (16.20)	206 (34.02)	556 (262.17)	559 (263.87)	561 (264.97)	556 (262.04)	557 (262.50)	506 (229.23)	560 (264.26)	558 (263.35)	558 (263.28)	560 (264.71)	566 (268.29)
DF _{A+S/A}	183 (3.14)	197 (10.86)	195 (9.95)	215 (21.39)	556 (213.76)	559 (215.23)	561 (216.19)	556 (213.65)	557 (214.04)	506 (185.22)	560 (215.57)	558 (214.78)	558 (214.72)	560 (215.96)	566 (219.06)
DF _{A+S/O}	178 (5.67)	195 (15.42)	190 (12.50)	215 (27.02)	556 (229.37)	559 (230.91)	561 (231.91)	556 (229.25)	557 (229.66)	506 (199.41)	560 (231.26)	558 (230.43)	558 (230.37)	560 (231.68)	566 (234.93)
DF _{A+S/V}	206 (1.25)	215 (5.87)	215 (5.52)	227 (11.42)	556 (173.64)	559 (174.91)	561 (175.75)	556 (173.54)	557 (173.88)	506 (148.75)	560 (175.21)	558 (174.52)	558 (174.47)	560 (175.55)	566 (178.26)
E _{1-ctx}	556 (0.00)	556 (0.00)	556 (0.00)	556 (0.00)	558 (0.39)	561 (0.78)	561 (0.91)	559 (0.51)	559 (0.55)	556 (0.00)	561 (0.79)	560 (0.74)	560 (0.66)	563 (1.12)	568 (2.05)
E _{2-ctx}	559 (0.00)	559 (0.00)	559 (0.00)	559 (0.00)	561 (0.31)	562 (0.62)	563 (0.73)	561 (0.42)	561 (0.45)	559 (0.00)	562 (0.63)	562 (0.59)	562 (0.52)	564 (0.92)	568 (1.70)
E _{3-ctx}	561 (0.00)	561 (0.00)	561 (0.00)	561 (0.00)	561 (0.13)	563 (0.42)	563 (0.49)	562 (0.25)	562 (0.27)	561 (0.00)	563 (0.41)	563 (0.40)	562 (0.31)	564 (0.67)	569 (1.45)
E _{DDG}	556 (0.00)	556 (0.00)	556 (0.00)	556 (0.00)	559 (0.54)	561 (0.92)	562 (1.06)	560 (0.61)	560 (0.70)	556 (0.00)	561 (0.94)	561 (0.85)	561 (0.80)	563 (1.24)	568 (2.12)
E _{LTO}	557 (0.00)	557 (0.00)	557 (0.00)	557 (0.00)	559 (0.46)	561 (0.83)	562 (0.95)	560 (0.57)	560 (0.60)	557 (0.00)	561 (0.83)	561 (0.80)	561 (0.72)	563 (1.16)	568 (2.03)
E _{MA}	506 (0.00)	506 (0.00)	506 (0.00)	506 (0.00)	556 (10.01)	559 (10.52)	561 (10.86)	556 (9.97)	557 (10.10)	511 (0.95)	560 (10.64)	558 (10.36)	558 (10.34)	560 (10.78)	566 (11.86)
E _{SANCOV}	560 (0.00)	560 (0.00)	560 (0.00)	560 (0.00)	561 (0.21)	562 (0.52)	563 (0.61)	561 (0.33)	561 (0.34)	560 (0.00)	562 (0.50)	562 (0.50)	562 (0.42)	564 (0.81)	568 (1.60)
E _{coarse}	558 (0.00)	558 (0.00)	558 (0.00)	558 (0.00)	560 (0.42)	562 (0.73)	563 (0.85)	561 (0.49)	561 (0.56)	558 (0.00)	562 (0.75)	562 (0.64)	562 (0.62)	564 (1.01)	568 (1.74)
P _{2-gram}	558 (0.00)	558 (0.00)	558 (0.00)	558 (0.00)	560 (0.35)	562 (0.68)	562 (0.78)	561 (0.46)	561 (0.51)	558 (0.00)	562 (0.69)	562 (0.63)	561 (0.55)	563 (0.95)	568 (1.79)
P _{4-gram}	560 (0.00)	560 (0.00)	560 (0.00)	560 (0.00)	563 (0.42)	564 (0.69)	564 (0.75)	563 (0.50)	563 (0.54)	560 (0.00)	564 (0.69)	564 (0.63)	563 (0.55)	565 (0.85)	569 (1.59)
P _{8-gram}	566 (0.00)	566 (0.00)	566 (0.00)	566 (0.00)	568 (0.36)	568 (0.48)	569 (0.54)	568 (0.39)	568 (0.43)	566 (0.00)	568 (0.49)	568 (0.37)	568 (0.40)	569 (0.60)	570 (0.80)
	B	DF _{A+S/A}	DF _{A+S/O}	DF _{A+S/V}	E _{1-ctx}	E _{2-ctx}	E _{3-ctx}	E _{DDG}	E _{LTO}	E _{MA}	E _{SANCOV}	E _{coarse}	P _{2-gram}	P _{4-gram}	P _{8-gram}

(A2) *zlib* (context-sensitive).

FIGURE 7.5: Predicted coverage improvements (continued).

The evaluation performed by Wang et al. [214] was inconclusive when comparing the performance of individual coverage metrics: for CGC “*there is no single [coverage metric] winner that beats everyone else*”, while the use of LAVA-M was “*not very suitable for our goal [of comparing the performance of coverage metrics]*”. Finally, the real-world targets report crashes rather than bugs, and these bugs were not thoroughly duplicated (e.g., it is unrealistic to assume up to 1,400 unique bugs were found in the *info2cap* target). Despite this, Wang et al. [214] found that combining multiple coverage metrics in an ensemble configuration (where each fuzzer ran in parallel and shared seeds) “*outperformed the baseline [i.e., fuzzing with a single coverage metric] by wide margins*”.

Salls et al. [183] similarly studied the performance of different coverage metrics, however, their study focused on formalizing coverage metrics by borrowing concepts from static analysis. Again, their study involved a relatively small evaluation of targets sourced from the CGC benchmark and two real-world programs. Salls et al. [183] showed that fuzzing outcomes were maximized when multiple coverage metrics were used in an ensemble configuration (which we did not investigate). When restricted to only a single metric, the “*edge + return loc*”—or context-sensitive edge coverage here—metric achieved the highest code coverage.

In contrast to these two studies, ours is the first large-scale comparison of coverage metrics that takes into account data-flow-based metrics. Moreover, we focus on the coverage achieved by different coverage metrics (as opposed to crash and bug counts). Once again, we found that more-sensitive coverage metrics did not necessarily lead to improvements in fuzzer performance, due to the “queue explosion” problem [141]. While Salls et al. [183] recommend combining multiple coverage metrics to maximize fuzzer performance, we do not recommend combining highly sensitive metrics (e.g., memory accesses, more sensitive forms of *def-use* chains) unless aggressive filtering occurs when sharing seeds between fuzzing nodes (assuming an ensemble fuzzing strategy is used). Moreover, our results in Section 7.3.3 show that there is little reason to look past traditional edge coverage metrics: either the E_{LTO} or E_{SANCOV} varieties will suffice for maximizing a fuzzer’s state space search. If an ensemble configuration with multiple metrics is used, a combination of lossless edge coverage, context-sensitive edge coverage, edge coverage augmented with a DDG, and n -gram coverage should be preferred.

7.5 Future Work

We see a number of promising directions for future work, which we discuss here.

Visualizing coverage. As discussed in Section 7.2, a fuzzer’s state-space search is typically visualized with a “(control-flow) coverage over time” plot. However, *is this the best approach for visualizing coverage metrics across multiple dimensions?* In particular, how can we incorporate the data-flow coverage approaches discussed throughout this dissertation (e.g., here and in Section 5.5.4) with traditional control-flow coverage? If we think of control- and data-flow as different axes of a target’s state space (as illustrated in Fig. 3.2), then a method for visualizing high-dimensional data—e.g., Principal Component Analysis (PCA) or t-distributed stochastic neighbor embedding (t-SNE) [136]—may be appropriate. We leave improvements to the visualization of fuzzers’ state-space search as future work.

Querying state space search. While effective, fuzzers often benefit from human-in-the-loop interaction. For example, IJON [7] uses human-provided annotations to guide the fuzzer. Similarly, Yan et al. [226] observed that many security professionals optimize fuzzing by assisting and guiding the testing process. Unsurprisingly, improving fuzzing via human interaction requires a thorough understanding of state space coverage (e.g., to understand fuzz blockers [70]). While tools such as Fuzz Introspector [32] and InFuzz [226] improve human understanding of a fuzzer’s state space search (e.g., via a web-based user interface), they do not provide the user with a *convenient querying interface*. This lack of a convenient querying interface hinders security professionals—who often have relevant domain expertise—from better understanding fuzzer performance (e.g., beyond aggregate coverage statistics that are typically available). We propose using a logic-based query language—e.g., Datalog—for analyzing and querying the coverage achieved during a fuzzing campaign. Datalog has a long history of being used to build static program analyses [15, 26, 64, 66, 81, 179]. Developing a query tool in Datalog—rather than a traditional graph database (e.g., Neo4j [218])—would allow one to combine static and dynamic analysis data in a single query tool. This combination would enable a better understanding of fuzzer roadblocks and a mechanism for quantifying data flow coverage. We leave building such a query tool as future work.

Ensemble fuzzing. Our evaluation focused on understanding the performance of fuzzer coverage metrics individually. However, as discussed in Sections 7.3.3 and 7.4, different coverage metrics can be combined using an ensemble fuzzing configuration [44, 78, 171, 183, 214]. In this configuration, multiple coverage metrics are run in parallel and seeds are shared across fuzzing instances. However, care must be taken to prevent queue explosion, which can hinder (a) an individual fuzzer’s seed scheduler, and (b) the effective sharing of seeds across fuzzers. Moreover, a diversity of fuzzers may also be beneficial (e.g., non-AFL-based fuzzers). While we

used the CUPID [78] approach for predicting the performance of combining multiple coverage metrics, we leave the experimentation, optimization, and validation of these predictions as future work.

7.6 Chapter Summary

This final chapter discussed the difficulties in measuring and quantifying a fuzzer's state space search. To address this, we proposed using context-sensitive static and dynamic analyses to quantify the control-flow elements of a fuzzer's state space search. We use these analyses to compare the full range of control- and data-flow-based coverage metrics discussed throughout this dissertation. Our results once again demonstrated the superiority of traditional edge coverage metrics. Unfortunately, our results also show that even state-of-the-art static analyses are far from accurate on typical fuzzer targets; they either failed to complete their analysis or produced unrealistic results. We hope these results and ideas for future work inspire others to improve the accuracy and scalability of static analyses for real-world code bases.

Next, we summarize this dissertation and provide some final reflections.

Chapter 8

Conclusions and Reflections

While fuzzing has repeatedly demonstrated its effectiveness at finding bugs in real-world software, it is not a panacea for society’s software security woes. For example, most fuzzing applications remain focused on exposing memory safety violations (despite other bug classes—e.g., privilege escalation, type confusion—enabling exploitation), and “deep” bugs continue to evade long-running fuzzing campaigns (e.g., with the fuzzer unable to satisfy complex control- and data-flow dependencies). It is this last point that we focus on in this dissertation. In particular, those deep bugs lurking in the corners of a program’s state space.

This dissertation attempts to understand better and improve fuzzers’ state space search. To this end, we present several techniques, frameworks, and methodologies for measuring and enhancing a fuzzer’s ability to explore a target program’s state space.

Chapters 2 and 3 start with an overview of coverage-guided greybox fuzzing, exploring and systematizing the different techniques fuzzers use to measure a target’s state space. We find that most prior work relies on control-flow-based coverage metrics. While performant, these metrics only consider one dimension of a program’s state space (i.e., features based on control flow), disregarding data-flow features.

Chapter 4 empirically demonstrates the importance of carefully curating a fuzzer’s initial seed set, bootstrapping the fuzzer’s state space search. Our results show that fuzzing outcomes (i.e., bug discovery and state space exploration) vary significantly depending on this initial seed set. This led to the design of a new fuzzing corpus minimizer, OPTIMIN. Unlike prior work, OPTIMIN can derive an optimal corpus by encoding the corpus minimization problem as a boolean satisfiability (SAT) problem. OPTIMIN has since been adopted by the broader fuzzing community and incorporated into AFL++¹ and LIBAFL², two state-of-the-art coverage-guided greybox

¹<https://github.com/AFLplusplus/AFLplusplus/commit/62f1bfed99b82bc073c138a00ff9a30bb596d09d>

²<https://github.com/AFLplusplus/LibAFL/pull/739>

fuzzers.

Chapter 5 promotes data-flow coverage to a first-class citizen (rather than using it to complement control-flow coverage, or ignoring it altogether). We design and implement a new fuzzer, DATAFLOW, that abstracts a target’s state space based on data-flow features (rather than traditional control-flow features); specifically, *def-use* chains. We show that, while not as performant as control-flow-driven fuzzers, DATAFLOW can discover bugs traditional fuzzers miss.

Chapter 6 investigates the difficulties of measuring fuzzers’ state space search. We evaluate the use of several static analysis frameworks for quantifying a fuzzer’s state space search, with our results revealing the limitations of these frameworks.

Finally, in Chapter 7 we implement the full range of coverage metrics discussed throughout this dissertation, comparing their ability to explore a target’s state space. Our implementation of *n*-gram coverage has since been adopted by the broader fuzzing community and incorporated into AFL++.³

Each of the chapters that make up this dissertation supports our thesis: enhancing state space search leads to improved fuzzing outcomes. Without a strong foundation, a fuzzer’s state space exploration is limited. Some bugs only manifest in program states not explicitly visible in control-flow-based abstractions. Security professionals deploying fuzzers require rigorous techniques to reason about their fuzzer’s state space search. Ultimately, these insights lead to better fuzzing outcomes and, ideally, more secure software.

8.1 Reflections

This dissertation opens the door to new approaches for improving fuzzers’ state space exploration. We present a summary of reflections here; detailed future work ideas are presented at the end of each chapter.

Chapter 4 focused on characterizing initial seed sets using traditional control-flow features (e.g., edge coverage). However, as this dissertation shows, control-flow features cover only one dimension of a target’s state space. Do fuzzers perform better when bootstrapped with seeds that also maximize coverage across data-flow features?

In Chapter 5, we concluded that one of the main impediments to DATAFLOW’s success was its high run-time cost (reducing fuzzer iteration rates). Does improving

³<https://github.com/AFLplusplus/AFLplusplus/commit/5a74cffa0f22b4e3b3dbc829dfb1c8f7c7a6fb76>

and optimizing the performance of our *def-use* chain coverage implementation lead to better fuzzing outcomes?

The static analyses presented in Chapter 6 rely on access to source code. This simplifies analysis, particularly compared to binary code (where control flow is harder to recover, and types are nonexistent). Can similar analyses be used by those measuring and comparing fuzzer performance on targets where source code is not available (e.g., proprietary software, embedded firmware)?

Bibliography

- [1] Humberto Abdelnur et al. *Spectral Fuzzing: Evaluation & Feedback*. Research Report RR-7193. Inria, 2010. URL: <https://hal.inria.fr/inria-00452015>.
- [2] Alif Ahmed et al. “BigMap: Future-proofing Fuzzers with Efficient Large Maps”. In: *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. DSN. IEEE, 2021, pp. 531–542. DOI: [10.1109/DSN48987.2021.00062](https://doi.org/10.1109/DSN48987.2021.00062).
- [3] Mike Aizatsky et al. *Announcing OSS-Fuzz: Continuous fuzzing for open source software*. 2016. URL: <https://opensource.googleblog.com/2016/12/announcing-oss-fuzz-continuous-fuzzing.html> (visited on 05/11/2022).
- [4] Lars Ole Andersen. “Program analysis and specialization for the C programming language”. PhD thesis. University of Copenhagen, 1994.
- [5] Andrea Arcuri and Lionel Briand. “A Practical Guide for Using Statistical Tests to Assess Randomized Algorithms in Software Engineering”. In: *International Conference on Software Engineering*. ICSE. 2011, pp. 1–10. DOI: [10.1145/1985793.1985795](https://doi.org/10.1145/1985793.1985795).
- [6] Cornelius Aschermann. *On Measuring and Visualizing Fuzzer Performance*. 2020. URL: <https://hexgolems.com/2020/08/on-measuring-and-visualizing-fuzzer-performance/> (visited on 12/03/2022).
- [7] Cornelius Aschermann et al. “Ijon: Exploring Deep State Spaces via Fuzzing”. In: *Security and Privacy*. S&P. IEEE, 2020, pp. 1597–1612. DOI: [10.1109/SP40000.2020.00117](https://doi.org/10.1109/SP40000.2020.00117).
- [8] Cornelius Aschermann et al. “NAUTILUS: Fishing for Deep Bugs with Grammars”. In: *Network and Distributed System Security*. NDSS. The Internet Society, 2019. DOI: [10.14722/ndss.2019.23412](https://doi.org/10.14722/ndss.2019.23412).
- [9] Cornelius Aschermann et al. “REDQUEEN: Fuzzing with Input-to-State Correspondence”. In: *Network and Distributed System Security*. NDSS. The Internet Society, 2019. DOI: [10.14722/ndss.2019.23371](https://doi.org/10.14722/ndss.2019.23371).
- [10] Florent Avellaneda. “A short description of the solver EvalMaxSAT”. In: *MaxSAT Evaluations*. MSE. 2020, pp. 8–9.
- [11] Nathaniel Ayewah et al. “Using Static Analysis to Find Bugs”. In: *IEEE Software* 25.5 (2008), pp. 22–29. DOI: [10.1109/MS.2008.130](https://doi.org/10.1109/MS.2008.130).

- [12] Fahiem Bacchus et al. *MaxSAT Evaluation 2021 - Rules*. 2021. URL: <https://maxsat-evaluations.github.io/2021/rules.html#input> (visited on 06/06/2022).
- [13] George Balatsouras and Yannis Smaragdakis. "Structure-Sensitive Points-To Analysis for C and C++". In: *Static Analysis Symposium*. SAS. Springer, 2016, pp. 84–104. DOI: [10.1007/978-3-662-53413-7_5](https://doi.org/10.1007/978-3-662-53413-7_5).
- [14] Roberto Baldoni et al. "A Survey of Symbolic Execution Techniques". In: *Computing Surveys*. CSUR 51.3 (2018), pp. 1–39. DOI: [10.1145/3182657](https://doi.org/10.1145/3182657).
- [15] Langston Barrett. *treeedb*. 2023. URL: <https://github.com/langston-barrett/treeedb> (visited on 10/01/2023).
- [16] Tim Blazytko et al. "GRIMOIRE: Synthesizing Structure While Fuzzing". In: *USENIX Security*. SEC. USENIX, 2019, pp. 1985–2002.
- [17] Marcel Böhme. "STADS: Software Testing as Species Discovery". In: *Transactions on Software Engineering Methodology*. TOSEM 27.2 (2018). DOI: [10.1145/3210309](https://doi.org/10.1145/3210309).
- [18] Marcel Böhme, Cristian Cadar, and Abhik Roychoudhury. "Fuzzing: Challenges and Reflections". In: *IEEE Software* 38.3 (2021), pp. 79–86. DOI: [10.1109/MS.2020.3016773](https://doi.org/10.1109/MS.2020.3016773).
- [19] Marcel Böhme and Brandon Falk. "Fuzzing: On the Exponential Cost of Vulnerability Discovery". In: *European Software Engineering Conference/Foundations of Software Engineering*. ESEC/FSE. 2020, pp. 713–724. DOI: [10.1145/3368089.3409729](https://doi.org/10.1145/3368089.3409729).
- [20] Marcel Böhme, Danushka Liyanage, and Valentin Wüstholtz. "Estimating Residual Risk in Greybox Fuzzing". In: *European Software Engineering Conference and Foundations of Software Engineering*. ESEC/FSE. 2021, pp. 230–241. DOI: [10.1145/3468264.3468570](https://doi.org/10.1145/3468264.3468570).
- [21] Marcel Böhme, Valentin J. M. Manès, and Sang Kil Cha. "Boosting Fuzzer Efficiency: An Information Theoretic Perspective". In: *European Software Engineering Conference/Foundations of Software Engineering*. ESEC/FSE. 2020, pp. 678–689. DOI: [10.1145/3368089.3409748](https://doi.org/10.1145/3368089.3409748).
- [22] Marcel Böhme, Van-Thuan Pham, and Abhik Roychoudhury. "Coverage-Based Greybox Fuzzing as Markov Chain". In: *Computer and Communications Security*. CCS. ACM, 2016, pp. 1032–1043. DOI: [10.1145/2976749.2978428](https://doi.org/10.1145/2976749.2978428).
- [23] Marcel Böhme, László Szekeres, and Jonathan Metzman. "On the Reliability of Coverage-Based Fuzzer Benchmarking". In: *International Conference on Software Engineering*. ICSE. 2022. DOI: [10.1145/3510003.3510230](https://doi.org/10.1145/3510003.3510230).
- [24] Marcel Böhme et al. "Directed Greybox Fuzzing". In: *Computer and Communications Security*. CCS. ACM, 2017, pp. 2329–2344. DOI: [10.1145/3133956.3134020](https://doi.org/10.1145/3133956.3134020).

- [25] Michael D. Bond and Kathryn S. McKinley. “Probabilistic Calling Context”. In: *Object-Oriented Programming Systems, Languages and Applications*. OOPSLA. 2007, pp. 97–112. DOI: [10.1145/1297027.1297035](https://doi.org/10.1145/1297027.1297035).
- [26] Martin Bravenboer and Yannis Smaragdakis. “Strictly Declarative Specification of Sophisticated Points-to Analyses”. In: *Object Oriented Programming Systems Languages and Applications*. OOPSLA. ACM, 2009, pp. 243–262. DOI: [10.1145/1640089.1640108](https://doi.org/10.1145/1640089.1640108).
- [27] Cristian Cadar, Daniel Dunbar, and Dawson Engler. “KLEE: Unassisted and Automatic Generation of High-Coverage Tests for Complex Systems Programs”. In: *Operating Systems Design and Implementation*. OSDI. 2008, pp. 209–224.
- [28] Miguel Castro, Manuel Costa, and Tim Harris. “Securing Software by Enforcing Data-Flow Integrity”. In: *Operating Systems Design and Implementation*. OSDI. 2006, pp. 147–160.
- [29] Marcos Lordello Chaim et al. “Efficiently Finding Data Flow Subsumptions”. In: *Software Testing, Verification and Validation*. ICST. 2021, pp. 94–104. DOI: [10.1109/ICST49551.2021.00021](https://doi.org/10.1109/ICST49551.2021.00021).
- [30] Madhurima Chakraborty et al. “Automatic Root Cause Quantification for Missing Edges in JavaScript Call Graphs”. In: *European Conference on Object-Oriented Programming*. Vol. 222. ECOOP. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 3:1–3:28. DOI: [10.4230/LIPICS.ECOOP.2022.3](https://doi.org/10.4230/LIPICS.ECOOP.2022.3).
- [31] Oliver Chang. *Taking the next step: OSS-Fuzz in 2023*. 2023. URL: <https://security.googleblog.com/2023/02/taking-next-step-oss-fuzz-in-2023.html> (visited on 02/03/2023).
- [32] Oliver Chang et al. *Introducing Fuzz Introspector, an OpenSSF Tool to Improve Fuzzing Coverage*. 2022. URL: <https://openssf.org/blog/2022/06/09/introducing-fuzz-introspector-an-openssf-tool-to-improve-fuzzing-coverage/> (visited on 03/10/2023).
- [33] Oliver Chang et al. *OSS-Fuzz: Five months later, and rewarding projects*. 2017. URL: <https://opensource.googleblog.com/2017/05/oss-fuzz-five-months-later-and.html> (visited on 05/11/2022).
- [34] Hongxu Chen et al. “Hawkeye: Towards a Desired Directed Grey-Box Fuzzer”. In: *Computer and Communications Security*. CCS. ACM, 2018, pp. 2095–2108. DOI: [10.1145/3243734.3243849](https://doi.org/10.1145/3243734.3243849).
- [35] Hongxu Chen et al. “MUZZ: Thread-Aware Grey-Box Fuzzing for Effective Bug Hunting in Multithreaded Programs”. In: *USENIX Security*. SEC. USENIX, 2020, pp. 2325–2342.

- [36] Ju Chen et al. "JIGSAW: Efficient and Scalable Path Constraints Fuzzing". In: *Security and Privacy*. S&P. IEEE, 2022, pp. 1531–1531. DOI: [10.1109/SP46214.2022.00102](https://doi.org/10.1109/SP46214.2022.00102).
- [37] Peng Chen and Hao Chen. "Angora: Efficient Fuzzing by Principled Search". In: *Security and Privacy*. S&P. IEEE, 2018, pp. 711–725. DOI: [10.1109/SP.2018.00046](https://doi.org/10.1109/SP.2018.00046).
- [38] Xi Chen, Asia Slowinska, and Herbert Bos. "On the Detection of Custom Memory Allocators in C Binaries". In: *Empirical Softw. Engg.* 21.3 (2016), pp. 753–777. DOI: [10.1007/s10664-015-9362-z](https://doi.org/10.1007/s10664-015-9362-z).
- [39] Yaohui Chen et al. "MEUZZ: Smart Seed Scheduling for Hybrid Fuzzing". In: *Research in Attacks, Intrusions and Defenses*. RAID. 2020, pp. 77–92.
- [40] Yaohui Chen et al. "PTrix: Efficient Hardware-Assisted Fuzzing for COTS Binary". In: *Asia Computer and Communications Security*. AsiaCCS. ACM, 2019, pp. 633–645. DOI: [10.1145/3321705.3329828](https://doi.org/10.1145/3321705.3329828).
- [41] Yaohui Chen et al. "SAVIOR: Towards Bug-Driven Hybrid Testing". In: *Security and Privacy*. S&P. IEEE, 2020, pp. 1580–1596. DOI: [10.1109/SP40000.2020.00002](https://doi.org/10.1109/SP40000.2020.00002).
- [42] Yiqun T. Chen et al. "Revisiting the Relationship Between Fault Detection, Test Adequacy Criteria, and Test Set Size". In: *Automated Software Engineering*. ASE. ACM, 2021, pp. 237–249. DOI: [10.1145/3324884.3416667](https://doi.org/10.1145/3324884.3416667).
- [43] Yongheng Chen et al. "One Engine to Fuzz 'em All: Generic Language Processor Testing with Semantic Validation". In: *Security and Privacy*. S&P. IEEE, 2021, pp. 642–658. DOI: [10.1109/SP40001.2021.00071](https://doi.org/10.1109/SP40001.2021.00071).
- [44] Yuanliang Chen et al. "EnFuzz: Ensemble Fuzzing with Seed Synchronization among Diverse Fuzzers". In: *USENIX Security*. USENIX, 2019, pp. 1967–1983.
- [45] B. Chess and G. McGraw. "Static analysis for security". In: *Security and Privacy*. S&P 2.6 (2004), pp. 76–79. DOI: [10.1109/MSP.2004.111](https://doi.org/10.1109/MSP.2004.111).
- [46] Jaeseung Choi et al. "Grey-Box Concolic Testing on Binary Code". In: *International Conference on Software Engineering*. ICSE. 2019, pp. 736–747. DOI: [10.1109/ICSE.2019.00082](https://doi.org/10.1109/ICSE.2019.00082).
- [47] Zheng Leong Chua et al. "One Engine To Serve 'em All: Inferring Taint Rules Without Architectural Semantics". In: *Network and Distributed System Security*. NDSS. The Internet Society, 2019. DOI: [10.14722/ndss.2019.23339](https://doi.org/10.14722/ndss.2019.23339).
- [48] Edmund M. Clarke, Orna Grumberg, and David E. Long. "Model Checking and Abstraction". In: *Transactions on Programming Languages and Systems*. TOPLAS 16.5 (1994), pp. 1512–1542. DOI: [10.1145/186025.186051](https://doi.org/10.1145/186025.186051).
- [49] Nicolas Coppik, Oliver Schwahn, and Neeraj Suri. "MemFuzz: Using Memory Accesses to Guide Fuzzing". In: *Software Testing, Validation and Verification*. ICST. 2019, pp. 48–58. DOI: [10.1109/ICST.2019.00015](https://doi.org/10.1109/ICST.2019.00015).

- [50] Baozeng Ding et al. “Baggy Bounds with Accurate Checking”. In: *International Symposium on Software Reliability Engineering Workshops*. ISSREW. 2012, pp. 195–200. DOI: [10.1109/ISSREW.2012.24](https://doi.org/10.1109/ISSREW.2012.24).
- [51] Ren Ding et al. “Hardware Support to Improve Fuzzing Performance and Precision”. In: *Computer and Communications Security*. CCS. ACM, 2021, pp. 2214–2228. DOI: [10.1145/3460120.3484573](https://doi.org/10.1145/3460120.3484573).
- [52] Zhen Yu Ding and Claire Le Goues. “An Empirical Study of OSS-Fuzz Bugs”. In: *Mining Software Repositories*. MSR. 2021, pp. 131–142. DOI: [10.1109/MSR52588.2021.00026](https://doi.org/10.1109/MSR52588.2021.00026).
- [53] Sung Ta Dinh et al. “Favocado: Fuzzing the Binding Code of JavaScript Engines Using Semantically Correct Test Cases”. In: *Network and Distributed System Security*. NDSS. The Internet Society, 2021. DOI: [10.14722/ndss.2021.24224](https://doi.org/10.14722/ndss.2021.24224).
- [54] Dino Distefano et al. “Scaling Static Analyses at Facebook”. In: *Communications of the ACM*. CACM 62.8 (2019), pp. 62–70. DOI: [10.1145/3338112](https://doi.org/10.1145/3338112).
- [55] Brendan Dolan-Gavitt et al. “LAVA: Large-Scale Automated Vulnerability Addition”. In: *Security and Privacy*. S&P. IEEE, 2016, pp. 110–121. DOI: [10.1109/SP.2016.15](https://doi.org/10.1109/SP.2016.15).
- [56] Thomas Dullien. “Weird Machines, Exploitability, and Provable Unexploitability”. In: *Transactions on Emerging Topics in Computing*. TETC 8.2 (2020), pp. 391–403. DOI: [10.1109/TETC.2017.2785299](https://doi.org/10.1109/TETC.2017.2785299).
- [57] Bradley Efron. “Bootstrap Methods: Another Look at the Jackknife”. In: *The Annals of Statistics* 7.1 (1979), pp. 1–26. DOI: [10.1214/aos/1176344552](https://doi.org/10.1214/aos/1176344552).
- [58] Facebook. HHVM. 2022. URL: <https://github.com/facebook/hhvm> (visited on 05/20/2022).
- [59] Brandon Falk. *Fuzzing: Corpus Minimization*. 2021. URL: <https://youtu.be/947b01gyvJs> (visited on 05/24/2022).
- [60] Jeanne Ferrante, Karl J. Ottenstein, and Joe D. Warren. “The Program Dependence Graph and Its Use in Optimization”. In: *Transactions on Programming Languages and Systems*. TOPLAS 9.3 (1987), pp. 319–349. DOI: [10.1145/24039.24041](https://doi.org/10.1145/24039.24041).
- [61] Andrea Fioraldi, Daniele Cono D’Elia, and Davide Balzarotti. “The Use of Likely Invariants as Feedback for Fuzzers”. In: *USENIX Security*. SEC. USENIX, 2021, pp. 2829–2846.
- [62] Andrea Fioraldi, Daniele Cono D’Elia, and Emilio Coppa. “WEIZZ: Automatic Grey-Box Fuzzing for Structured Binary Formats”. In: *International Symposium on Software Testing and Analysis*. ISSTA. 2020, pp. 1–13. DOI: [10.1145/3395363.3397372](https://doi.org/10.1145/3395363.3397372).

- [63] Andrea Fioraldi et al. “AFL++: Combining Incremental Steps of Fuzzing Research”. In: *Workshop on Offensive Technologies*. WOOT. 2020.
- [64] Antonio Flores-Montoya and Eric Schulte. “Datalog Disassembly”. In: *USENIX Security*. SEC. USENIX, 2020, pp. 1075–1092.
- [65] Phyllis G. Frankl and Stewart N. Weiss. “An experimental comparison of the effectiveness of branch testing and data flow testing”. In: *Transactions on Software Engineering*. TSE 19.8 (1993), pp. 774–787. DOI: [10.1109/32.238581](https://doi.org/10.1109/32.238581).
- [66] Galois Inc. *clyzer++*. 2022. URL: <https://galoisinc.github.io/clyzerpp/> (visited on 04/09/2023).
- [67] Shuitao Gan et al. “CollAFL: Path Sensitive Fuzzing”. In: *Security and Privacy*. S&P. IEEE, 2018, pp. 679–696. DOI: [10.1109/SP.2018.00040](https://doi.org/10.1109/SP.2018.00040).
- [68] Shuitao Gan et al. “GREYONE: Data Flow Sensitive Fuzzing”. In: *USENIX Security*. SEC. USENIX, 2020, pp. 2577–2594.
- [69] Vijay Ganesh, Tim Leek, and Martin Rinard. “Taint-Based Directed White-box Fuzzing”. In: *International Conference on Software Engineering*. ICSE. 2009, pp. 474–484. DOI: [10.1109/ICSE.2009.5070546](https://doi.org/10.1109/ICSE.2009.5070546).
- [70] Wentao Gao et al. “Beyond the Coverage Plateau: A Comprehensive Study of Fuzz Blockers (Registered Report)”. In: *Fuzzing Workshop*. FUZZING. ACM, 2023, pp. 47–55. DOI: [10.1145/3605157.3605177](https://doi.org/10.1145/3605157.3605177).
- [71] GitLab. *GitLab Protocol Fuzzer Community Edition*. 2022. URL: <https://gitlab.com/gitlab-org/security-products/protocol-fuzzer-ce> (visited on 06/23/2022).
- [72] GNU Binutils. 2023. URL: <https://www.gnu.org/software/binutils/> (visited on 02/26/2023).
- [73] Patrice Godefroid, Michael Y. Levin, and David Molnar. “SAGE: Whitebox Fuzzing for Security Testing”. In: *Queue* 10.1 (2012), pp. 20–27. DOI: [10.1145/2090147.2094081](https://doi.org/10.1145/2090147.2094081).
- [74] Philipp Goerz et al. “Systematic Assessment of Fuzzers using Mutation Analysis”. In: *USENIX Security*. SEC. USENIX, 2023.
- [75] Google. *Fuzzer Test Suite*. 2021. URL: <https://github.com/google/fuzzer-test-suite> (visited on 05/10/2022).
- [76] Rahul Gopinath, Carlos Jensen, and Alex Groce. “Code Coverage for Suite Evaluation by Developers”. In: *International Conference on Software Engineering*. ICSE. 2014, pp. 72–82. DOI: [10.1145/2568225.2568278](https://doi.org/10.1145/2568225.2568278).
- [77] Gustavo Grieco, Martín Ceresa, and Pablo Buiras. “QuickFuzz: An Automatic Random Fuzzer for Common File Formats”. In: *Haskell*. HASKELL. 2016, pp. 13–20. DOI: [10.1145/2976002.2976017](https://doi.org/10.1145/2976002.2976017).

- [78] Emre Güler et al. “Cupid: Automatic Fuzzer Selection for Collaborative Fuzzing”. In: *Annual Computer Security Applications Conference*. ACSAC. ACM, 2020, pp. 360–372. DOI: [10.1145/3427228.3427266](https://doi.org/10.1145/3427228.3427266).
- [79] Arie Gurfinkel and Jorge A. Navas. “Abstract Interpretation of LLVM with a Region-Based Memory Model”. In: *Software Verification*. VSTTE. Springer International Publishing, 2021, pp. 122–144. DOI: [10.1007/978-3-030-95561-8_8](https://doi.org/10.1007/978-3-030-95561-8_8).
- [80] Arie Gurfinkel et al. “The SeaHorn Verification Framework”. In: *Computer Aided Verification*. CAV. Springer International Publishing, 2015, pp. 343–361. DOI: [10.1007/978-3-319-21690-4_20](https://doi.org/10.1007/978-3-319-21690-4_20).
- [81] Elnar Hajiyev, Mathieu Verbaere, and Oege de Moor. “CodeQuest: Scalable Source Code Queries with Datalog”. In: *European Conference on Object-Oriented Programming*. ECOOP. Springer Berlin Heidelberg, 2006, pp. 2–27. DOI: [10.1007/11785477_2](https://doi.org/10.1007/11785477_2).
- [82] HyungSeok Han, DongHyeon Oh, and Sang Kil Cha. “CodeAlchemist: Semantics-Aware Code Generation to Find Vulnerabilities in JavaScript Engines”. In: *Network and Distributed System Security*. NDSS. The Internet Society, 2019. DOI: [10.14722/ndss.2019.23263](https://doi.org/10.14722/ndss.2019.23263).
- [83] Nikolas Havrikov and Andreas Zeller. “Systematically Covering Input Structure”. In: *Automated Software Engineering*. ASE. IEEE, 2020, pp. 189–199. DOI: [10.1109/ASE.2019.00027](https://doi.org/10.1109/ASE.2019.00027).
- [84] Ahmad Hazimeh, Adrian Herrera, and Mathias Payer. “Magma: A Ground-Truth Fuzzing Benchmark”. In: *Measurement and Analysis of Computing Systems*. POMACS 4.3 (2020). DOI: [10.1145/3428334](https://doi.org/10.1145/3428334).
- [85] Xiaoyu He et al. “SoFi: Reflection-Augmented Fuzzing for JavaScript Engines”. In: *Computer and Communications Security*. CCS. ACM, 2021, pp. 2229–2242. DOI: [10.1145/3460120.3484823](https://doi.org/10.1145/3460120.3484823).
- [86] Mats P. E. Heimdahl and Devaraj George. “Test-Suite Reduction for Model Based Tests: Effects on Test Quality and Implications for Testing”. In: *Automated Software Engineering*. ASE. IEEE, 2004, pp. 176–185. DOI: [10.1109/ASE.2004.1342735](https://doi.org/10.1109/ASE.2004.1342735).
- [87] Aki Helin. *Radamsa*. 2022. URL: <https://gitlab.com/akihe/radamsa> (visited on 06/17/2022).
- [88] Hadi Hemmati. “How Effective Are Code Coverage Criteria?” In: *Software Quality, Reliability and Security*. QRS. 2015, pp. 151–156. DOI: [10.1109/QRS.2015.30](https://doi.org/10.1109/QRS.2015.30).
- [89] John L. Henning. “SPEC CPU2006 Benchmark Descriptions”. In: *Computer Architecture News* 34.4 (2006), pp. 1–17. DOI: [10.1145/1186736.1186737](https://doi.org/10.1145/1186736.1186737).

- [90] Adrian Herrera, Mathias Payer, and Antony L. Hosking. “datAFLOW: Toward a Data-Flow-Guided Fuzzer”. In: *Transactions on Software Engineering Methodology*. TOSEM 1.1 (2023). DOI: [10.1145/3587156](https://doi.org/10.1145/3587156).
- [91] Adrian Herrera, Mathias Payer, and Antony L. Hosking. “datAFLOW: Towards a Data-Flow-Guided Fuzzer”. In: *Fuzzing Workshop*. FUZZING. The Internet Society, 2022. DOI: [10.14722/fuzzing.2022.23001](https://doi.org/10.14722/fuzzing.2022.23001).
- [92] Adrian Herrera et al. “Seed Selection for Successful Fuzzing”. In: *International Symposium on Software Testing and Analysis*. ISSTA. 2021, pp. 230–243. DOI: [10.1145/3460319.3464795](https://doi.org/10.1145/3460319.3464795).
- [93] Joseph R. Horgan and Saul London. “Data Flow Coverage and the C Language”. In: *Testing, Analysis, and Verification*. TAV. 1991, pp. 87–97. DOI: [10.1145/120807.120815](https://doi.org/10.1145/120807.120815).
- [94] Joseph R. Horgan, Saul London, and Michael R. Lyu. “Achieving software quality with testing coverage measures”. In: *Computer* 27.9 (1994), pp. 60–69. DOI: [10.1109/2.312032](https://doi.org/10.1109/2.312032).
- [95] Katherine Hough and Jonathan Bell. “A Practical Approach for Dynamic Taint Tracking with Control-Flow Relationships”. In: *Transactions on Software Engineering Methodology*. TOSEM 31.2 (2021). DOI: [10.1145/3485464](https://doi.org/10.1145/3485464).
- [96] *How SQLite Is Tested*. 2022. URL: <https://www.sqlite.org/testing.html> (visited on 02/03/2023).
- [97] Chin-Chia Hsu et al. “InsTrim: Lightweight Instrumentation for Coverage-guided Fuzzing”. In: *Binary Analysis Research*. BAR. The Internet Society, 2018. DOI: [10.14722/bar.2018.23014](https://doi.org/10.14722/bar.2018.23014).
- [98] Hwa-You Hsu and Alessandro Orso. “MINTS: A General Framework and Tool for Supporting Test-Suite Minimization”. In: *International Conference on Software Engineering*. ICSE. 2009, pp. 419–429. DOI: [10.1109/ICSE.2009.5070541](https://doi.org/10.1109/ICSE.2009.5070541).
- [99] Heqing Huang et al. “BEACON: Directed Grey-Box Fuzzing with Provable Path Pruning”. In: *Security and Privacy*. S&P. IEEE, 2022, pp. 104–118. DOI: [10.1109/SP46214.2022.00007](https://doi.org/10.1109/SP46214.2022.00007).
- [100] Heqing Huang et al. “Pangolin: Incremental Hybrid Fuzzing with Polyhedral Path Abstraction”. In: *Security and Privacy*. S&P. IEEE, 2020, pp. 1613–1627. DOI: [10.1109/SP40000.2020.00063](https://doi.org/10.1109/SP40000.2020.00063).
- [101] Monica Hutchins et al. “Experiments on the Effectiveness of Dataflow- and Controlflow-Based Test Adequacy Criteria”. In: *International Conference on Software Engineering*. ICSE. 1994, pp. 191–200. DOI: [10.1109/ICSE.1994.296778](https://doi.org/10.1109/ICSE.1994.296778).

- [102] Laura Inozemtseva and Reid Holmes. “Coverage is Not Strongly Correlated with Test Suite Effectiveness”. In: *International Conference on Software Engineering*. ICSE. 2014, pp. 435–445. DOI: [10.1145/2568225.2568271](https://doi.org/10.1145/2568225.2568271).
- [103] Kyriakos K. Ispoglou et al. “FuzzGen: Automatic Fuzzer Generation”. In: *USENIX Security*. SEC. USENIX, 2020, pp. 2271–2287.
- [104] Yuseok Jeon et al. “FuZZan: Efficient Sanitizer Metadata Design for Fuzzing”. In: *Annual Technical Conference*. ATC. 2020, pp. 249–263.
- [105] Yuseok Jeon et al. “HexType: Efficient Detection of Type Confusion Errors for C++”. In: *Computer and Communications Security*. CCS. ACM, 2017, pp. 2373–2387. DOI: [10.1145/3133956.3134062](https://doi.org/10.1145/3133956.3134062).
- [106] Ranjit Jhala and Rupak Majumdar. “Software Model Checking”. In: *Computing Surveys*. CSUR 41.4 (2009), pp. 1–54. DOI: [10.1145/1592434.1592438](https://doi.org/10.1145/1592434.1592438).
- [107] Zhiyuan Jiang et al. “Evocatio: Conjuring Bug Capabilities from a Single PoC”. In: *Computer and Communications Security*. CCS. ACM, 2022, pp. 1599–1613. DOI: [10.1145/3548606.3560575](https://doi.org/10.1145/3548606.3560575).
- [108] Zhiyuan Jiang et al. “Igor: Crash Deduplication Through Root-Cause Clustering”. In: *Computer and Communications Security*. CCS. ACM, 2021, pp. 3318–3336. DOI: [10.1145/3460120.3485364](https://doi.org/10.1145/3460120.3485364).
- [109] Brittany Johnson et al. “Why Don’t Software Developers Use Static Analysis Tools to Find Bugs?” In: *International Conference on Software Engineering*. ICSE. IEEE, 2013, pp. 672–681. DOI: [10.5555/2486788.2486877](https://doi.org/10.5555/2486788.2486877).
- [110] Edward L. Kaplan and Paul Meier. “Nonparametric Estimation from Incomplete Observations”. In: *Journal of the American Statistical Association* 53.282 (1958), pp. 457–481. ISSN: 01621459. DOI: [10.2307/2281868](https://doi.org/10.2307/2281868).
- [111] Richard M. Karp. “Reducibility among Combinatorial Problems”. In: *Complexity of Computer Computations*. 1972, pp. 85–103. DOI: [10.1007/978-1-4684-2001-2_9](https://doi.org/10.1007/978-1-4684-2001-2_9).
- [112] Kenney. *Kenney*. 2022. URL: <https://www.kenney.nl/> (visited on 05/20/2022).
- [113] Yonit Kesten and Amir Pnueli. “Control and data abstraction: the cornerstones of practical formal verification”. In: *Software Tools for Technology Transfer*. STTT 2.4 (2000), pp. 328–342. DOI: [10.1007/s100090050040](https://doi.org/10.1007/s100090050040).
- [114] Se-Won Kim, Xavier Rival, and Sukyoung Ryu. “A Theoretical Foundation of Sensitivity in an Abstract Interpretation Framework”. In: *Transactions on Programming Languages and Systems*. TOPLAS 40.3 (2018). DOI: [10.1145/3230624](https://doi.org/10.1145/3230624).
- [115] George Klees et al. “Evaluating Fuzz Testing”. In: *Computer and Communications Security*. CCS. ACM, 2018, pp. 2123–2138. DOI: [10.1145/3243734.3243804](https://doi.org/10.1145/3243734.3243804).

- [116] David G. Kleinbaum and Mitchel Klein. *Survival Analysis*. 2012. DOI: [10.1007/978-1-4419-6646-9](https://doi.org/10.1007/978-1-4419-6646-9).
- [117] Pavneet Singh Kochhar, Ferdian Thung, and David Lo. “Code coverage and test suite effectiveness: Empirical study with real bugs in large systems”. In: *Software Analysis, Evolution, and Reengineering*. SANER. 2015, pp. 560–564. DOI: [10.1109/SANER.2015.7081877](https://doi.org/10.1109/SANER.2015.7081877).
- [118] Erik van der Kouwe et al. “SoK: Benchmarking Flaws in Systems Security”. In: *European Security and Privacy*. EuroS&P. IEEE, 2019, pp. 310–325. DOI: [10.1109/EuroSP.2019.00031](https://doi.org/10.1109/EuroSP.2019.00031).
- [119] Jakub Kuderski, Jorge A. Navas, and Arie Gurfinkel. “Unification-based Pointer Analysis without Oversharing”. In: *Formal Methods in Computer Aided Design*. FMCAD. IEEE, 2019, pp. 37–45. DOI: [10.23919/FMCAD.2019.8894275](https://doi.org/10.23919/FMCAD.2019.8894275).
- [120] Chris Lattner and Vikram Adve. “LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation”. In: *Code Generation and Optimization*. CGO. 2004, pp. 75–86. DOI: [10.1109/CGO.2004.1281665](https://doi.org/10.1109/CGO.2004.1281665).
- [121] Gwangmu Lee, Woochul Shim, and Byoungyoung Lee. “Constraint-guided Directed Greybox Fuzzing”. In: *USENIX Security*. SEC. USENIX, 2021, pp. 3559–3576.
- [122] Yuekang Li et al. “Cerebro: Context-Aware Adaptive Fuzzing for Effective Vulnerability Detection”. In: *European Software Engineering Conference/Foundations of Software Engineering*. ESEC/FSE. 2019, pp. 533–544. DOI: [10.1145/3338906.3338975](https://doi.org/10.1145/3338906.3338975).
- [123] Yuekang Li et al. “Steelix: Program-State Based Binary Fuzzing”. In: *European Software Engineering Conference/Foundations of Software Engineering*. ESEC/FSE. 2017, pp. 627–637. DOI: [10.1145/3106237.3106295](https://doi.org/10.1145/3106237.3106295).
- [124] Yuwei Li et al. “UNIFUZZ: A Holistic and Pragmatic Metrics-Driven Platform for Evaluating Fuzzers”. In: *USENIX Security*. SEC. USENIX, 2021, pp. 2777–2794.
- [125] Jie Liang et al. “PATA: Fuzzing with Path Aware Taint Analysis”. In: *Security and Privacy*. S&P. IEEE, 2022, pp. 154–170. DOI: [10.1109/SP46214.2022.00010](https://doi.org/10.1109/SP46214.2022.00010).
- [126] Jun-Wei Lin et al. “Nemo: Multi-Criteria Test-Suite Minimization with Integer Nonlinear Programming”. In: *International Conference on Software Engineering*. ICSE. 2018, pp. 1039–1049. DOI: [10.1145/3180155.3180174](https://doi.org/10.1145/3180155.3180174).
- [127] Stephan Lipp et al. “FuzzTastic: A Fine-Grained, Fuzzer-Agnostic Coverage Analyzer”. In: *International Conference on Software Engineering: Companion Proceedings*. ICSE. ACM, 2022, pp. 75–79. DOI: [10.1145/3510454.3516847](https://doi.org/10.1145/3510454.3516847).
- [128] Zhengyang Liu and John Criswell. “Flexible and Efficient Memory Object Metadata”. In: *International Symposium on Memory Management*. ISMM. 2017, pp. 36–46. DOI: [10.1145/3092255.3092268](https://doi.org/10.1145/3092255.3092268).

- [129] Danushka Liyanage et al. “Reachable Coverage: Estimating Saturation in Fuzzing”. In: *International Conference on Software Engineering*. ICSE. ACM, 2023.
- [130] LLVM. *libFuzzer - a library for coverage-guided fuzz testing*. 2022. URL: <https://llvm.org/docs/LibFuzzer.html> (visited on 06/23/2022).
- [131] Lunge Technology. *Cyber Grand Challenge Corpus*. 2017. URL: <http://www.lungetech.com/cgc-corporus/> (visited on 05/22/2022).
- [132] Simon Luo et al. “Make out like a (Multi-Armed) Bandit: Improving the Odds of Fuzzer Seed Scheduling with T-Scheduler”. In: *Asia Computer and Communications Security*. AsiaCCS. 2024.
- [133] Chenyang Lyu et al. “EMS: History-Driven Mutation for Coverage-based Fuzzing”. In: *Network and Distributed System Security*. NDSS. The Internet Society, 2022. DOI: [10.14722/ndss.2022.23162](https://doi.org/10.14722/ndss.2022.23162).
- [134] Chenyang Lyu et al. “MOPT: Optimized Mutation Scheduling for Fuzzers”. In: *USENIX Security*. SEC. USENIX, 2019, pp. 1949–1966.
- [135] Chenyang Lyu et al. “SmartSeed: Smart Seed Generation for Efficient Fuzzing”. In: *arXiv preprint* (2019). DOI: [10.48550/arXiv.1807.02606](https://doi.org/10.48550/arXiv.1807.02606).
- [136] Laurens van der Maaten and Geoffrey Hinton. “Visualizing Data using t-SNE”. In: *Journal of Machine Learning Research* 9.86 (2008), pp. 2579–2605.
- [137] Valentin J. M. Manès, Soomin Kim, and Sang Kil Cha. “Ankou: Guiding Grey-Box Fuzzing towards Combinatorial Difference”. In: *International Conference on Software Engineering*. ICSE. 2020, pp. 1024–1036. DOI: [10.1145/3377811.3380421](https://doi.org/10.1145/3377811.3380421).
- [138] Valentin J. M. Manès et al. “The Art, Science, and Engineering of Fuzzing: A Survey”. In: *Transactions on Software Engineering*. TSE 47.11 (2021), pp. 2312–2331. DOI: [10.1109/TSE.2019.2946563](https://doi.org/10.1109/TSE.2019.2946563).
- [139] Henry B. Mann and Donald R. Whitney. “On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other”. In: *The Annals of Mathematical Statistics* 18.1 (1947), pp. 50–60. DOI: [10.1214/aoms/1177730491](https://doi.org/10.1214/aoms/1177730491).
- [140] Nathan Mantel. “Evaluation of survival data and two new rank order statistics arising in its consideration”. In: *Cancer Chemotherapy Reports* 50.3 (1966), pp. 163–170.
- [141] Alessandro Mantovani, Andrea Fioraldi, and Davide Balzarotti. “Fuzzing with Data Dependency Information”. In: *European Security and Privacy*. EuroS&P. IEEE, 2022, pp. 286–302. DOI: [10.1109/EuroSP53844.2022.00026](https://doi.org/10.1109/EuroSP53844.2022.00026).
- [142] Björn Mathis et al. “Parser-Directed Fuzzing”. In: *Programming Language Design and Implementation*. PLDI. 2019, pp. 548–560. DOI: [10.1145/3314221.3314651](https://doi.org/10.1145/3314221.3314651).

- [143] Ruijie Meng et al. “Linear-time Temporal Logic guided Greybox Fuzzing”. In: *International Conference on Software Engineering*. ICSE. 2022. DOI: [10.1145/3510003.3510082](https://doi.org/10.1145/3510003.3510082).
- [144] Jonathan Metzman et al. “FuzzBench: An Open Fuzzer Benchmarking Platform and Service”. In: *European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. ESEC/FSE. 2021, pp. 1393–1403. DOI: [10.1145/3468264.3473932](https://doi.org/10.1145/3468264.3473932).
- [145] Xianya Mi et al. “LeanSym: Efficient Hybrid Fuzzing Through Conservative Constraint Debloating”. In: *Research in Attacks, Intrusions and Defenses*. RAID. 2021, pp. 62–77. DOI: [10.1145/3471621.3471852](https://doi.org/10.1145/3471621.3471852).
- [146] Bart Miller. *CS736 Projects*. 1988. URL: <https://pages.cs.wisc.edu/~bart/fuzz/CS736-Projects-f1988.pdf> (visited on 08/06/2022).
- [147] Matt Miller. *Trends and Challenges in the Vulnerability Mitigation Landscape*. 2019.
- [148] The Motion Monkey. *Free Retro Arcade Video Game Sound Effects*. 2017. URL: <https://www.themotionmonkey.co.uk/free-resources/retro-arcade-sounds/> (visited on 05/20/2022).
- [149] Mozilla. *Fuzzing—Test Samples*. 2020. URL: <https://firefox-source-docs.mozilla.org/tools/fuzzing/index.html> (visited on 05/11/2022).
- [150] Mozilla. *Introducing the ASan Nightly Project*. 2018. URL: <https://blog.mozilla.org/security/2018/07/19/introducing-the-asan-nightly-project/> (visited on 05/11/2022).
- [151] Stefan Nagy and Matthew Hicks. “Full-Speed Fuzzing: Reducing Fuzzing Overhead through Coverage-Guided Tracing”. In: *Security and Privacy*. S&P. IEEE, 2019, pp. 787–802. DOI: [10.1109/SP.2019.00069](https://doi.org/10.1109/SP.2019.00069).
- [152] Stefan Nagy et al. “Breaking Through Binaries: Compiler-quality Instrumentation for Better Binary-only Fuzzing”. In: *USENIX Security*. SEC. USENIX, 2021, pp. 1683–1700.
- [153] Stefan Nagy et al. “Same Coverage, Less Bloat: Accelerating Binary-Only Fuzzing with Coverage-Preserving Coverage-Guided Tracing”. In: *Computer and Communications Security*. CCS. ACM, 2021, pp. 351–365. DOI: [10.1145/3460120.3484787](https://doi.org/10.1145/3460120.3484787).
- [154] George C. Necula et al. “CCured: Type-Safe Retrofitting of Legacy Software”. In: *Transactions on Programming Languages and Systems*. TOPLAS 27.3 (2005), pp. 477–526. DOI: [10.1145/1065887.1065892](https://doi.org/10.1145/1065887.1065892).
- [155] Neelofar et al. “Instance Space Analysis of Search-Based Software Testing”. In: *IEEE Transactions on Software Engineering*. TSE (2022), pp. 1–20. DOI: [10.1109/TSE.2022.3228334](https://doi.org/10.1109/TSE.2022.3228334).

- [156] Manh-Dung Nguyen et al. “Binary-level Directed Fuzzing for Use-After-Free Vulnerabilities”. In: *Research in Attacks, Intrusions and Defenses*. RAID. 2020, pp. 47–62.
- [157] Ivica Nikolić et al. “Refined Grey-Box Fuzzing with Sivo”. In: *Detection of Intrusions and Malware, and Vulnerability Assessment*. DIMVA. 2021, pp. 106–129. DOI: [10.1007/978-3-030-80825-9_6](https://doi.org/10.1007/978-3-030-80825-9_6).
- [158] NIST. CVE-2021-43527. 2021. URL: <https://nvd.nist.gov/vuln/detail/CVE-2021-43527> (visited on 02/13/2023).
- [159] Tim Nosco et al. “The Industrial Age of Hacking”. In: *USENIX Security*. SEC. USENIX, 2020, pp. 1129–1146.
- [160] Carlos Oliveira et al. “Mapping the Effectiveness of Automated Test Suite Generation Techniques”. In: *IEEE Transactions on Reliability*. TR 67.3 (2018), pp. 771–785. DOI: [10.1109/TR.2018.2832072](https://doi.org/10.1109/TR.2018.2832072).
- [161] *Open Game Art*. 2022. URL: <https://opengameart.org/> (visited on 05/20/2022).
- [162] Tavis Ormandy. *This shouldn't have happened: A vulnerability postmortem*. 2021. URL: <https://googleprojectzero.blogspot.com/2021/12/this-shouldnt-have-happened.html> (visited on 02/12/2023).
- [163] Sebastian Österlund et al. “ParmeSan: Sanitizer-Guided Greybox Fuzzing”. In: *USENIX Security*. SEC. USENIX, 2020, pp. 2289–2306.
- [164] Rohan Padhye et al. “FuzzFactory: Domain-Specific Fuzzing with Waypoints”. In: *Programming Languages*. OOPSLA 3 (2019). DOI: [10.1145/3360600](https://doi.org/10.1145/3360600).
- [165] Rohan Padhye et al. “Semantic Fuzzing with Zest”. In: *International Symposium on Software Testing and Analysis*. ISSTA. 2019, pp. 329–340. DOI: [10.1145/3293882.3330576](https://doi.org/10.1145/3293882.3330576).
- [166] Shankara Pailoor, Andrew Aday, and Suman Jana. “Moonshine: Optimizing OS Fuzzer Seed Selection with Trace Distillation”. In: *USENIX Security*. SEC. USENIX, 2018, pp. 729–743.
- [167] Weyu Pan et al. “Grammar-Agnostic Symbolic Execution by Token Symbolization”. In: *International Symposium on Software Testing and Analysis*. ISSTA. 2021, pp. 374–387. DOI: [10.1145/3460319.3464845](https://doi.org/10.1145/3460319.3464845).
- [168] Soyeon Park et al. “Fuzzing JavaScript Engines with Aspect-preserving Mutation”. In: *Security and Privacy*. S&P. IEEE, 2020, pp. 1629–1642. DOI: [10.1109/SP40000.2020.00067](https://doi.org/10.1109/SP40000.2020.00067).
- [169] Terence Parr and Kathleen Fisher. “LL(*): The Foundation of the ANTLR Parser Generator”. In: *Programming Language Design and Implementation*. PLDI. 2011, pp. 425–436. DOI: [10.1145/1993498.1993548](https://doi.org/10.1145/1993498.1993548).
- [170] Hui Peng, Yan Shoshitaishvili, and Mathias Payer. “T-Fuzz: Fuzzing by Program Transformation”. In: *Security and Privacy*. S&P. IEEE, 2018, pp. 697–710. DOI: [10.1109/SP.2018.00056](https://doi.org/10.1109/SP.2018.00056).

- [171] Van-Thuan Pham et al. "Towards Systematic and Dynamic Task Allocation for Collaborative Parallel Fuzzing". In: *Automated Software Engineering*. ASE. 2021, pp. 1337–1341. DOI: [10.1109/ASE51524.2021.9678810](https://doi.org/10.1109/ASE51524.2021.9678810).
- [172] Daniel Plohmann et al. "Malpedia: A Collaborative Effort to Inventorize the Malware Landscape". In: *Journal on Cybercrime & Digital Investigations* 3.1 (2018). DOI: [10.18464/cybin.v3i1.17](https://doi.org/10.18464/cybin.v3i1.17).
- [173] Sebastian Poeplau and Aurélien Francillon. "Symbolic execution with SymCC: Don't interpret, compile!" In: *USENIX Security*. SEC. USENIX, 2020, pp. 181–198.
- [174] Sebastian Poeplau and Aurélien Francillon. "SymQEMU: Compilation-based symbolic execution for binaries". In: *Network and Distributed System Security*. NDSS. The Internet Society, 2021. DOI: [10.14722/ndss.2021.23118](https://doi.org/10.14722/ndss.2021.23118).
- [175] Sandra Rapps and Elaine J. Weyuker. "Selecting Software Test Data Using Data Flow Information". In: *Transactions on Software Engineering*. TSE 11.4 (1985), pp. 367–375. DOI: [10.1109/TSE.1985.232226](https://doi.org/10.1109/TSE.1985.232226).
- [176] Sanjay Rawat et al. "VUzzer: Application-aware Evolutionary Fuzzing". In: *Network and Distributed System Security*. NDSS. The Internet Society, 2017. DOI: [10.14722/ndss.2017.23404](https://doi.org/10.14722/ndss.2017.23404).
- [177] Alexandre Rebert et al. "Optimizing Seed Selection for Fuzzing". In: *USENIX Security*. SEC. USENIX, 2014, pp. 861–875.
- [178] *RegExLib: Regular Expression Library*. 2022. URL: <https://regexlib.com/> (visited on 05/20/2022).
- [179] Thomas W. Reps and Raghu Ramakrishnan. "Demand Interprocedural Program Analysis Using Logic Databases". In: *Applications of Logic Databases*. Springer, 1995, pp. 163–196. DOI: [10.1007/978-1-4615-2207-2_8](https://doi.org/10.1007/978-1-4615-2207-2_8).
- [180] H. G. Rice. "Classes of Recursively Enumerable Sets and Their Decision Problems". In: *Transactions of the American Mathematical Society* 74.2 (1953), pp. 358–366. DOI: [10.1090/s0002-9947-1953-0053041-6](https://doi.org/10.1090/s0002-9947-1953-0053041-6).
- [181] Matt Ruhstaller, TPM, and Oliver Chang. *A New Chapter for OSS-Fuzz*. 2018. URL: <https://security.googleblog.com/2018/11/a-new-chapter-for-oss-fuzz.html> (visited on 05/11/2022).
- [182] Caitlin Sadowski et al. "Lessons from Building Static Analysis Tools at Google". In: *Communications of the ACM*. CACM 61.4 (2018), pp. 58–66. DOI: [10.1145/3188720](https://doi.org/10.1145/3188720).
- [183] Christopher Salls et al. "Exploring Abstraction Functions in Fuzzing". In: *Communications and Network Security*. CNS. 2020, pp. 1–9. DOI: [10.1109/CNS48642.2020.9162273](https://doi.org/10.1109/CNS48642.2020.9162273).
- [184] Christopher Salls et al. "Token-Level Fuzzing". In: *USENIX Security*. SEC. USENIX, 2021, pp. 2795–2809.

- [185] Philipp Dominik Schubert, Ben Hermann, and Eric Bodden. “PhASAR: An Inter-procedural Static Analysis Framework for C/C++”. In: *Tools and Algorithms for the Construction and Analysis of Systems*. TACAS. Springer International Publishing, 2019, pp. 393–410.
- [186] Sergej Schumilo et al. “Nyx: Greybox Hypervisor Fuzzing using Fast Snapshots and Affine Types”. In: *USENIX Security*. SEC. USENIX, 2021, pp. 2597–2614.
- [187] Konstantin Serebryany et al. “AddressSanitizer: A Fast Address Sanity Checker”. In: *Annual Technical Conference*. ATC. 2012, pp. 309–318.
- [188] Kosta Serebryany. “Continuous Fuzzing with libFuzzer and AddressSanitizer”. In: *Cybersecurity Development*. SecDev. 2016, pp. 157–157. DOI: [10.1109/SecDev.2016.043](https://doi.org/10.1109/SecDev.2016.043).
- [189] Kostya Serebryany. “OSS-Fuzz - Google’s continuous fuzzing service for open source software”. In: *USENIX Security*. SEC. USENIX, 2017.
- [190] Dongdong She, Abhishek Shah, and Suman Jana. “Effective Seed Scheduling for Fuzzing with Graph Centrality Analysis”. In: *Security and Privacy*. S&P. IEEE, 2022, pp. 1558–1558. DOI: [10.1109/SP46214.2022.00129](https://doi.org/10.1109/SP46214.2022.00129).
- [191] Dongdong She et al. “MTFuzz: Fuzzing with a Multi-Task Neural Network”. In: *European Software Engineering Conference and Foundations of Software Engineering*. ESEC/FSE. 2020, pp. 737–749. DOI: [10.1145/3368089.3409723](https://doi.org/10.1145/3368089.3409723).
- [192] Dongdong She et al. “Neutaint: Efficient Dynamic Taint Analysis with Neural Networks”. In: *Security and Privacy*. S&P. IEEE, 2022, pp. 1527–1543. DOI: [10.1109/SP40000.2020.00022](https://doi.org/10.1109/SP40000.2020.00022).
- [193] Dongdong She et al. “NEUZZ: Efficient Fuzzing with Neural Program Smoothing”. In: *Security and Privacy*. IEEE, 2019, pp. 803–817. DOI: [10.1109/SP.2019.00052](https://doi.org/10.1109/SP.2019.00052).
- [194] Yan Shoshitaishvili et al. “(State of) The Art of War: Offensive Techniques in Binary Analysis”. In: *Security and Privacy*. S&P. IEEE, 2016, pp. 138–157. DOI: [10.1109/SP.2016.17](https://doi.org/10.1109/SP.2016.17).
- [195] Laurent Simon and Akash Verma. “Improving Fuzzing through Controlled Compilation”. In: *European Security and Privacy*. EuroS&P. IEEE, 2020, pp. 34–52. DOI: [10.1109/EuroSP48549.2020.00011](https://doi.org/10.1109/EuroSP48549.2020.00011).
- [196] Software Engineering Institute. *CERT BFF*. 2016. URL: <https://www.cert.org/vulnerability-analysis/tools/bff.cfm> (visited on 06/17/2022).
- [197] Dokyung Song et al. “SoK: Sanitizing for Security”. In: *Security and Privacy*. S&P. IEEE, 2019, pp. 1275–1295. DOI: [10.1109/SP.2019.00010](https://doi.org/10.1109/SP.2019.00010).
- [198] Suhwan Song et al. “CrFuzz: Fuzzing Multi-Purpose Programs through Input Validation”. In: *European Software Engineering Conference and Foundations of*

- Software Engineering*. ESEC/FSE. 2020, pp. 690–700. DOI: [10.1145/3368089.3409769](https://doi.org/10.1145/3368089.3409769).
- [199] Prashast Srivastava and Mathias Payer. “Gramatron: Effective Grammar-Aware Fuzzing”. In: *International Symposium on Software Testing and Analysis*. ISSTA. 2021, pp. 244–256. DOI: [10.1145/3460319.3464814](https://doi.org/10.1145/3460319.3464814).
- [200] JHU/APL Staff. *Assembled Labeled Library for Static Analysis Research (ALLSTAR) Dataset*. 2019. URL: <https://allstar.jhuapl.edu/> (visited on 05/19/2022).
- [201] Nick Stephens et al. “Driller: Augmenting Fuzzing Through Selective Symbolic Execution”. In: *Network and Distributed System Security*. NDSS. The Internet Society, 2016. DOI: [10.14722/ndss.2016.23368](https://doi.org/10.14722/ndss.2016.23368).
- [202] Ting Su et al. “A Survey on Data-Flow Testing”. In: *Computing Surveys*. CSUR 50.1 (2017). DOI: [10.1145/3020266](https://doi.org/10.1145/3020266).
- [203] Yulei Sui and Jingling Xue. “On-Demand Strong Update Analysis via Value-Flow Refinement”. In: *Foundations of Software Engineering*. FSE. ACM, 2016, pp. 460–473. DOI: [10.1145/2950290.2950296](https://doi.org/10.1145/2950290.2950296).
- [204] Yulei Sui and Jingling Xue. “SVF: Interprocedural Static Value-Flow Analysis in LLVM”. In: *Compiler Construction*. CC. 2016, pp. 265–266. DOI: [10.1145/2892208.2892235](https://doi.org/10.1145/2892208.2892235).
- [205] William N. Sumner and Xiangyu Zhang. “Identifying Execution Points for Dynamic Analyses”. In: *Automated Software Engineering*. ASE. IEEE, 2013, pp. 81–91. DOI: [10.1109/ASE.2013.6693069](https://doi.org/10.1109/ASE.2013.6693069).
- [206] William N. Sumner et al. “Precise Calling Context Encoding”. In: *International Conference on Software Engineering*. ICSE. 2010, pp. 525–534. DOI: [10.1145/1806799.1806875](https://doi.org/10.1145/1806799.1806875).
- [207] Robert Swiecki. *honggfuzz*. 2016. URL: <https://honggfuzz.dev/> (visited on 05/19/2022).
- [208] The Clang Team. *SanitizerCoverage*. 2022. URL: <https://clang.llvm.org/docs/SanitizerCoverage.html> (visited on 06/14/2022).
- [209] The Clang Team. *Source-Based Code Coverage*. 2022. URL: <https://clang.llvm.org/docs/SourceBasedCodeCoverage.html> (visited on 05/22/2022).
- [210] *UndefinedBehaviorSanitizer*. 2022. URL: <https://clang.llvm.org/docs/UndefinedBehaviorSanitizer.html> (visited on 06/20/2022).
- [211] Jonas Benedict Wagner. “Elastic Program Transformations Automatically Optimizing the Reliability/Performance Trade-off in Systems Software”. In: (2017). DOI: [10.5075/epfl-thesis-7745](https://doi.org/10.5075/epfl-thesis-7745).
- [212] Daimeng Wang et al. “SyzVegas: Beating Kernel Fuzzing Odds with Reinforcement Learning”. In: *USENIX Security*. SEC. USENIX, 2021, pp. 2741–2758.

- [213] Jinghan Wang, Chengyu Song, and Heng Yin. "Reinforcement Learning-based Hierarchical Seed Scheduling for Greybox Fuzzing". In: NDSS. The Internet Society, 2021. DOI: [10.14722/ndss.2021.24486](https://doi.org/10.14722/ndss.2021.24486).
- [214] Jinghan Wang et al. "Be Sensitive and Collaborative: Analyzing Impact of Coverage Metrics in Greybox Fuzzing". In: *Research in Attacks, Intrusions and Defenses*. RAID. 2019, pp. 1–15.
- [215] Junjie Wang et al. "Skyfire: Data-Driven Seed Generation for Fuzzing". In: *Security and Privacy*. S&P. IEEE, 2017, pp. 579–594. DOI: [10.1109/SP.2017.23](https://doi.org/10.1109/SP.2017.23).
- [216] Junjie Wang et al. "Superion: Grammar-Aware Greybox Fuzzing". In: *International Conference on Software Engineering*. ICSE. 2019, pp. 724–735. DOI: [10.1109/ICSE.2019.00081](https://doi.org/10.1109/ICSE.2019.00081).
- [217] Yanhao Wang et al. "Not All Coverage Measurements Are Equal: Fuzzing by Coverage Accounting for Input Prioritization". In: *Network and Distributed System Security*. NDSS. The Internet Society, 2020. DOI: [10.14722/ndss.2020.24422](https://doi.org/10.14722/ndss.2020.24422).
- [218] Jim Webber. "A Programmatic Introduction to Neo4j". In: *Systems, Programming, and Applications: Software for Humanity*. SPLASH. ACM, 2012, pp. 217–218. DOI: [10.1145/2384716.2384777](https://doi.org/10.1145/2384716.2384777).
- [219] Cheng Wen et al. "MemLock: Memory Usage Guided Fuzzing". In: *International Conference on Software Engineering*. ICSE. 2020, pp. 765–777. DOI: [10.1145/3377811.3380396](https://doi.org/10.1145/3377811.3380396).
- [220] W. Eric Wong et al. "Effect of test set minimization on fault detection effectiveness". In: *International Conference on Software Engineering*. ICSE. 1995, pp. 41–50. DOI: [10.1145/225014.225018](https://doi.org/10.1145/225014.225018).
- [221] WordPress. *WordPress*. 2022. URL: <https://github.com/WordPress/WordPress> (visited on 05/20/2022).
- [222] Mingyuan Wu et al. "Evaluating and Improving Neural Program-Smoothing-based Fuzzing". In: *International Conference on Software Engineering*. ICSE. 2022. DOI: [10.1145/3510003.3510089](https://doi.org/10.1145/3510003.3510089).
- [223] Mingyuan Wu et al. "One Fuzzing Strategy to Rule Them All". In: *International Conference on Software Engineering*. ICSE. DOI: [10.1145/3510003.3510174](https://doi.org/10.1145/3510003.3510174).
- [224] Bin Xin, William N. Sumner, and Xiangyu Zhang. "Efficient Program Execution Indexing". In: *Programming Language Design and Implementation*. PLDI. 2008, pp. 238–248. DOI: [10.1145/1375581.1375611](https://doi.org/10.1145/1375581.1375611).
- [225] Wen Xu et al. "Designing New Operating Primitives to Improve Fuzzing Performance". In: *Computer and Communications Security*. CCS. ACM, 2017, pp. 2313–2328. DOI: [10.1145/3133956.3134046](https://doi.org/10.1145/3133956.3134046).

- [226] Qian Yan et al. “InFuzz: An Interactive Tool for Enhancing Efficiency in Fuzzing through Visual Bottleneck Analysis (Registered Report)”. In: *Fuzzing Workshop*. FUZZING. ACM, 2023, pp. 56–61. DOI: [10.1145/3605157.3605847](https://doi.org/10.1145/3605157.3605847).
- [227] Shengbo Yan et al. “PathAFL: Path-Coverage Assisted Fuzzing”. In: *Asia Computer and Communications Security*. AsiaCCS. ACM, 2020, pp. 598–609. DOI: [10.1145/3320269.3384736](https://doi.org/10.1145/3320269.3384736).
- [228] Wei You et al. “SLF: Fuzzing without Valid Seed Inputs”. In: *International Conference on Software Engineering*. ICSE. 2019, pp. 712–723. DOI: [10.1109/ICSE.2019.00080](https://doi.org/10.1109/ICSE.2019.00080).
- [229] Tai Yue et al. “EcoFuzz: Adaptive Energy-Saving Greybox Fuzzing as a Variant of the Adversarial Multi-Armed Bandit”. In: *USENIX Security*. SEC. USENIX, 2020, pp. 2307–2324.
- [230] Insu Yun et al. “QSYM: A Practical Concolic Execution Engine Tailored for Hybrid Fuzzing”. In: *USENIX Security*. SEC. USENIX, 2018, pp. 745–761.
- [231] Michał Zalewski. *American Fuzzy Lop (AFL)*. 2015. URL: <http://lcamtuf.coredump.cx/afl/> (visited on 05/19/2022).
- [232] Gen Zhang et al. “MobFuz: Adaptive Multi-objective Optimization in Graybox Fuzzing”. In: *Network and Distributed System Security*. NDSS. The Internet Society, 2022. DOI: [10.14722/ndss.2022.24314](https://doi.org/10.14722/ndss.2022.24314).
- [233] Kunpeng Zhang et al. “Path Transitions Tell More: Optimizing Fuzzing Schedules via Runtime Program States”. In: *International Conference on Software Engineering*. ICSE. 2022. DOI: [10.1145/3510003.3510063](https://doi.org/10.1145/3510003.3510063).
- [234] Zenong Zhang et al. “FIXREVERTER: A Realistic Bug Injection Methodology for Benchmarking Fuzz Testing”. In: *USENIX Security*. SEC. USENIX, 2022, pp. 3699–3715.
- [235] Zenong Zhang et al. “Registered Report: Fuzzing Configurations of Program Options”. In: *Fuzzing Workshop*. FUZZING. The Internet Society, 2022. DOI: [10.14722/fuzzing.2022.23008](https://doi.org/10.14722/fuzzing.2022.23008).
- [236] Zhuo Zhang et al. “StochFuzz: Sound and Cost-effective Fuzzing of Stripped Binaries by Incremental and Stochastic Rewriting”. In: *Security and Privacy*. S&P. IEEE, 2021, pp. 659–676. DOI: [10.1109/SP40001.2021.00109](https://doi.org/10.1109/SP40001.2021.00109).
- [237] Chijin Zhou et al. “Zeror: Speed up Fuzzing with Coverage-Sensitive Tracing and Scheduling”. In: *Automated Software Engineering*. ASE. 2020, pp. 858–870. DOI: [10.1145/3324884.3416572](https://doi.org/10.1145/3324884.3416572).
- [238] Tong Zhou et al. “Valence: Variable Length Calling Context Encoding”. In: *Compiler Construction*. CC. ACM, 2019, pp. 147–158. DOI: [10.1145/3302516.3307351](https://doi.org/10.1145/3302516.3307351).

-
- [239] Xiaogang Zhu and Marcel Böhme. “Regression Greybox Fuzzing”. In: *Computer and Communications Security*. CCS. ACM, 2021, pp. 2169–2182. DOI: [10.1145/3460120.3484596](https://doi.org/10.1145/3460120.3484596).
- [240] Peiyuan Zong et al. “FuzzGuard: Filtering out Unreachable Inputs in Directed Grey-Box Fuzzing through Deep Learning”. In: *USENIX Security*. SEC. USENIX, 2020, pp. 2255–2269.
- [241] Zyte. *Scrapy*. 2022. URL: <https://scrapy.org/> (visited on 05/20/2022).