CARVE: Practical Security-Focused Software Debloating Using Simple Feature Set Mappings

Michael D. Brown

School of Computer Science
Georgia Institute of Technology
Atlanta GA, USA
mbrown337@gatech.edu

Santosh Pande

School of Computer Science
Georgia Institute of Technology
Atlanta GA, USA
santosh.pande@gatech.edu

ABSTRACT

Software debloating is an emerging field of study aimed at improving the security and performance of software by removing excess library code and features that are not needed by the end user (called bloat). Software bloat is pervasive, and several debloating techniques have been proposed to address this problem. While these techniques are effective at removing bloat, they are not practical for the average end user, risk creating unsound programs and introducing new vulnerabilities, and are not well suited for debloating complex software such as network protocol implementations.

In this paper, we propose CARVE, a simple yet effective security-focused debloating technique that addresses these shortcomings. CARVE employs static source code annotation to map software features to source code, eliminating the need for advanced software analysis during debloating and reducing the overall level of technical sophistication required by the end user. CARVE also introduces the concept of *debloating with replacement*, which is capable of removing software features while preserving software interoperability and mitigating the risk of creating an unsound program or introducing a vulnerability.

We evaluate CARVE in 12 debloating scenarios and present our results demonstrating security and performance improvements that meet or exceed those of existing techniques.

CCS CONCEPTS

• Security and privacy → Software security engineering;

KEYWORDS

Software debloating; Protocol debloating; Security;

ACM Reference format:

Michael D. Brown and Santosh Pande. 2019. CARVE: Practical Security-Focused Software Debloating Using Simple Feature Set Mappings. In *The 3rd Workshop on Forming an Ecosystem Around Software Transformation (FEAST '19)*, November 15, 2019, London, United Kingdom. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3338502.3359764>

1. Introduction

Software bloat affects almost all software, negatively impacting security and performance [1, 6, 7]. The primary source of bloat is the result of the tendency of software packages to accumulate new features over time which are unlikely to be used by most end users [1, 8]. Unnecessary features are a security risk as they may contain reachable software vulnerabilities or be useful to an attacker crafting a code reuse exploit. Bloat also results from software engineering practices that favor reuse of modular software libraries. Programs typically use a small subset of the functionality provided by each external library [1], yet the entire library is loaded into the program's memory space at runtime. As is the case with unnecessary features, excess library code may be useful for crafting code reuse attacks and also increases memory consumption at runtime.

Several debloating techniques [2-5, 9, 10, 18] have been proposed that promise to improve software security and performance by eliminating unnecessary features and excess library code. While these approaches have been shown to be effective at reducing bloat, they have serious limitations that diminish their potential for use in real-world software engineering projects and security-focused debloating.

Motivation. Existing feature-based debloating techniques require complex specifications and advanced software analysis techniques which have a high technical barrier to entry for the typical end user. For example, CHISEL [3] requires a test script that compiles the source code, executes it with inputs corresponding to desired and undesired features, and tests the output of each run to determine if the resulting program is correct [14]. Even for small programs, CHISEL requires long, complex scripts that test a large number of concrete inputs to produce debloated variants. Other approaches such as TOSS [2] require sophisticated tools such as angr and TEMU [12-13] to debloat software.

© 2019 Association for Computing Machinery. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the United States Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

FEAST'19, November 15, 2019, London, United Kingdom.

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6834-6/19/11...\$15.00

<https://doi.org/10.1145/3338502.3359764>

Figure 1: Overview of CARVE.

Despite their complexity, existing techniques are not capable of preserving desirable software properties such as graceful error handling, user experience, specification compliance, and interoperability. Because these techniques are only capable of removing code, the debloated variants they produce cannot respond in a meaningful way to attempts to invoke debloated functionality. For example, invoking debloated behavior in CHISEL benchmark programs [14] results in the process terminating with no meaningful output. Network protocol implementations are particularly affected in terms of interoperability. Specification-compliant requests for removed functionality made to server software debloated by TOSS cannot respond in a protocol compliant manner, as there is no longer code to handle such requests. The client will likely misinterpret this behavior as a connection failure rather than an unimplemented feature, for which there is typically a well-formed response (HTTP code 501, FTP code 502, etc.).

Existing approaches also have several limitations with respect to debloating for security. Techniques that use dynamic analysis to guide debloating such as CHISEL and TOSS can produce unsound programs and introduce new vulnerabilities as a result of the imprecise manner in which feature related code is identified [18]. Also, these techniques permanently alter the program’s representation, which has been shown to introduce new code reuse gadgets in unpredictable ways [22]. Gadget introduction can negatively impact security, and multiple debloating and analysis iterations may be required to achieve a desirable result. Techniques such as CHISEL are not well suited for this task, as debloating operations can take hours to complete [3].

Contributions. We present CARVE, a simple yet effective security-focused debloating technique designed to overcome the limitations of existing approaches. Using CARVE (Figure 1), developers statically map features directly to the source code that implements them using simple and flexible comment-based mappings similar to those commonly used by the software engineering community. To debloat software, end users need only specify a list of the fine-grained features they wish to debloat and invoke CARVE on the feature mapped source code.

CARVE scans the source code and intelligently removes code associated with the undesired features, optionally replacing debloated code with replacement code that preserves high level program properties or prevents the introduction of vulnerabilities. Debloating is fast, efficient, and sound producing a reduced version of the source code containing only desired features that is compatible with the original build chain. We discuss CARVE’s design and implementation in Section 2.

We evaluated CARVE by using it to debloat features from four network protocol implementations at varying levels of aggressiveness. Debloating this class of software is challenging due to its input / output complexity, compliance to specification, interoperability across platforms, and stateful nature. We analyzed the resulting debloated variants and observed security and performance benefits including vulnerability elimination, reduction in code reuse gadget set utility, reduction in code size, and external dependency elimination. We present the results of our evaluation in Section 3.

2. Design

Feature-based debloating techniques must perform two tasks: mapping software features to the source code that implements them and removing feature code according to the end user’s needs. Existing techniques [2, 3] employ dynamic analysis methods to accomplish both tasks in tandem when the deblocator is invoked. Feature mapping is accomplished by executing the program with concrete inputs corresponding to features, and analyzing the resulting execution paths or program output. These methods are technically complex, requiring the end user to provide in depth test scripts or use advanced software analysis tools to generate this mapping. Scalability is also a concern, since a large number of executions may be required to ensure that the generated mapping is correct. As a consequence, the cost of these approaches grows with the input space and execution time of the target program, which may prove cost-prohibitive for large, feature-rich programs. Using the dynamically generated feature mapping data, these techniques remove feature-related code from the program’s source code or binary representation to complete the debloating operation.

CARVE overcomes these limitations by decoupling the task of mapping features from the task of specifying and debloating feature code. This enables a division of labor between developers, who statically embed feature mappings directly in source code and end users, who need only specify a list of features they wish to debloat and invoke CARVE's debloater on the feature-mapped source code. This contrasts existing approaches (e.g. CHISEL and TOSS) which require the end user to perform technically complex tasks, however our approach requires developers complete the one-time task of generating feature mappings. If necessary, end users can also perform this task (e.g. debloating software that is no longer supported), however the effort required in this scenario is significant.

Embedding the feature mapping as comment-based metadata is familiar task for software developers; this method is used in popular automated software documentation tools such as Doxygen and Sphinx [24, 25] and functions similarly to C/C++ preprocessor directives. The primary disadvantage to this approach is the manual effort required to annotate the source code with feature mappings. Depending on the size of the code base and complexity of the features to be marked as debloatable, the manual effort required to generate feature mappings may be significant. However, we found that this was not the case in practice. Modern software engineering places heavy emphasis on code modularity, refactoring, and separation of concerns, which allows for one to one mapping of the majority of software features to their associated source code.

CARVE further mitigates this disadvantage by incorporating several design elements that reduce the effort required for developers to produce feature mappings. Since feature mappings are static and persist within source code, the cost to produce the initial mapping need only be paid once. Additionally, software features can be nested within a feature hierarchy, simplifying the task of mapping a segment of code to multiple related features. CARVE also supports implicit feature mappings that trigger syntax-aware debloating for common code constructs, which automatically handle control-flow implications of code removal. Implicit mappings provide rich debloating support that goes far beyond that of preprocessor directive schemes.

Feature Mapping Anatomy. Figure 2 contains sample code with embedded feature mappings. Feature mappings consist of three parts. The first part is a configurable tag (`///`), that distinguishes code comments from feature mappings. Immediately following the tag, one or more feature names are specified, each enclosed within a pair of square braces (`[ ]`). This feature list marks the mapped code as debloatable if the end user specifies all of the features for removal. The final part is the optional operator, which is used to differentiate between explicit file mappings, explicit segment mappings, and implicit mappings.

Explicit Feature Mappings. CARVE supports two types of explicit feature mappings (**bold blue** in Figure 2). File mappings are indicated by the `!` operator, and instruct the

```

0:  ///[FeatureGroup_A][Feature_G]!
1:
2:  #include <stdio.h>
3:
4:  ///[Feature_Y]~
5:  int a = 0;
6:  int b = 1;
7:  ///~
8:
9:  ///[Feature_Z]
10: char msg[10];
11:
12: int func_y (int i, int j){
13:     ///[Feature_Y]~
14:     ///^
15:     ///return 0;
16:     ///^
17:     int ret = i + a;
18:     int temp = b - j;
19:     return ret * temp;
20:     ///~
21: }
22: ///[Feature_Z]
23: void func_z(int k){ printf(msg); }
24:
25: ///[Feature_G]
26: int func_g (int i){
27:     int j = func_y(i, 3);
28:     ///[Feature_Z]
29:     if(j<0){ func_z(j); }
30:     else{ j = j + i; }
31:     return j;
32: }
```

Figure 2: Example Code with Feature Mappings.

debloater to remove all code in the file if the specified features are all undesired. Segment explicit mappings are indicated by the `~` operator, and instruct the debloater to remove code between the mapping and the next occurring termination tag (`///~`). Replacement code segments can also be specified between the two replacement tags (`///^`) for segment explicit mappings. For example, the debloating with replacement mapping on line 13 instructs the debloater to remove the code on lines 17-19, and replace it with the code on line 15.

Implicit Feature Mappings. Implicit feature mappings (**bold green** in Figure 2) reduce the effort required to generate feature mappings by offloading two key tasks to the debloater. First, implicit mappings do not require a termination tag to mark the end of the code segment associated with a mapping. Implicit mappings instruct the debloater to analyze the code following the mapping to determine its structure, and by extension what code should be removed. Implicit mappings also obviate the need to handle predictable control-flow implications associated with certain types of code constructs by using explicit mappings with replacement code. For example, mappings preceding switch cases or conditional branches instruct the debloater to analyze the control-flow implications of removing the mapped code, and debloat the code in way that does produce unintended control-flow for the resulting debloated program. Examples of implicit mappings are shown in Figure 2. When processing the mapping on line 9, the debloater will recognize that the code following the

Before Debloating [Feature_B]	After Debloating [Feature_B]
<pre>0: // Switch Case Example ... 13: else { handle_generic(input); }</pre>	<pre>0: // Switch Case Example ... 13: </pre>

Figure 3: CARVE Transformations for Code with Control-Flow Implications.

mapping is a single statement, and remove only line 10. The mappings on lines 22 and 25 target function definitions; in these cases, the debloater will scan for the closing brace associated with the function definition to determine the which lines to remove. The debloater will remove line 23 and lines 26-32 for these mappings, respectively.

Figure 3 shows examples of debloating code constructs with control-flow implications. The code on the right results from debloating `Feature_B` from the code on the left. During syntax-aware analysis of the mapped `switch` case on line 4, the debloater will determine that the mapped code on lines 6 and 7 are reachable by the previous `case` label due to the absence of a `break` statement. In this case, the debloater cannot remove these lines and can only remove the `case` label. Similarly, the debloater will analyze the conditional branch structure on lines 12 and 13 and determine that the presence of an `else` block in the branching construct prevents complete removal of line 12. To maintain the correct branching behavior for the `else` block, only the body of the `else if` block can be debloated.

Processing implicit mappings in the debloater requires syntax-aware analysis that is specific to different programming languages. For this reason, CARVE’s design supports the creation of custom, language-specific debloating modules for future expansion. The resource debloater for C/C++ supports implicit mappings for individual statements, function and struct definitions, switch cases, and conditional branches.

Expressive Power. We found that implicit feature mappings are sufficiently expressive to correctly map code to features in the vast majority of cases. Due to the flexibility of high-level programming languages, there are situations in which implicit mappings are insufficient. In these cases, segment explicit mappings should be used (with replacement code if necessary). These mappings are capable of arbitrary removal and replacement of code, and are sufficiently expressive for any debloating operation. Across all four of our benchmarks, use of explicit mappings in this manner was required for only 5 of 668 total mappings.

Debloater Operation. The static feature mapping scheme used in CARVE decouples the task of mapping features to code from the task of debloating the code. As a result, debloating is a straightforward and inexpensive operation. CARVE takes as input a single configuration file specifying the location of the feature mapped source code, the programming language the code is written in (used to select the correct syntax-aware debloating module), the hierarchy of all debloatable features, and the list of features the end user wishes to debloat. When executed, the debloater makes a copy of the specified source code directories, and begins scanning the feature mapped source code. When feature mappings are found, the debloater compares them to the features specified for debloating. If the feature mapping is a subset of or equal to the features specified for debloating, the debloater processes the mapping. Otherwise, the debloater continues until all files have been scanned and processed.

Soundness. Assuming that the feature mappings (including replacement code) are correct, CARVE is a sound debloating technique. The individual transformations performed are sound with respect to program syntax and control-flow. The static feature mapping scheme is fine-grained and directly maps features to the source code that implements them; therefore, a correct mapping does not miss code associated with features or mis-associate code with a feature it does not implement. The same cannot be said for dynamic approaches. CHISEL has been shown to introduce vulnerabilities and cause unexpected program behavior in the author’s benchmarks due to limitations inherent to dynamic analysis [18]. The feature identification mechanism used by TOSS [2] cannot guarantee that all feature code is identified. Post-debloating fuzzing of both desired and undesired functionality is required to test for unsound variants.

Debloating source code is challenging due to its highly expressive nature. There are instances where it is not possible to remove code without producing an unsound program or introducing a vulnerability (e.g. high-level security policies). Consider the example in Figure 4. Debloating either case from the switch statement does not violate language semantics, but does create a reachable path in which the execution of the protected function without authorizing credentials is possible.

```
switch(auth_type){
    case 1: perform_auth_1(creds); break;
    case 2: perform_auth_2(creds); break;
}
protected_function();
```

Figure 4: Vulnerability Introduction via Code Removal.

Using debloating with replacement, the developer can properly handle the implications of removing source code. In our example, the switch case can be replaced during debloating with exception handling code that traps execution if a debloated `auth_type` is specified at runtime.

As is the case with writing source code, creating a feature mapping can be an error prone process. In practice, errors made during the creation of a feature mapping can be identified using the same tools used to ensure that the source code itself is sound. Mapping errors that result in violations of language semantics will be caught by the compiler when building the debloated code. Mapping errors present in code that successfully compiles can be identified with common testing techniques such as unit and integration testing. Security concerns, such as inadvertent removal of bounds checks, can be identified using common static analysis tools like Coverity [26].

3. Evaluation

To evaluate CARVE, we selected four software packages implementing network protocols that vary in size, feature density, and complexity. The software packages we selected are:

- `libmodbus` v3.1.4, an industrial protocol library.
- `Bftpd` v4.9, an FTP server utility program.
- `libcurl` v7.61.0, a data transfer utility library.
- `mongoose` v6.8, an embedded web server library.

For each package, we manually enriched the source code with fine-grained feature mappings and created three different debloating configurations. These configurations were selected to reflect reasonable real-world use cases corresponding to aggressive, moderate, and conservative debloating scenarios. We define these levels in the following common language:

- **Conservative (C):** Some peripheral features in the package are targeted for debloating.
- **Moderate (M):** Some peripheral features and some core features are targeted for debloating.
- **Aggressive (A):** All debloatable features except for a small set of core features are targeted for debloating.

In each scenario, we debloated the package with CARVE and built the resulting variant using the default package build configuration and GCC v7.3.0. All software packages and their debloated variants were built on the same platform. We tested all variants using a combination of provided test scripts and custom testing scripts to ensure that they were correct with respect to kept features, and that inputs invoking debloated features did not cause the variant to crash.

Security Improvements. The most compelling security benefit of debloating software features is the potential to eliminate software vulnerabilities contained within these features. To demonstrate CARVE's potential for vulnerability elimination, we searched the CVE database [19] for known vulnerabilities affecting our selected software packages. There were no listed CVEs for `libmodbus` and `Bftpd`, however there are six known vulnerabilities affecting `libcurl` and 17 affecting `mongoose`. We analyzed the debloated versions of our selected packages to determine how many vulnerabilities were actually eliminated in our scenarios. Our results are shown in column two of Table 1.

Another commonly cited security benefit of debloating compiled software is increased resistance against gadget-based code reuse attacks. In gadget-based code reuse methods such as ROP, JOP, and COP [15, 16, 17], execution is redirected to a chain of short instruction sequences (i.e. gadgets) resident in the program to construct a malicious payload without injecting code. Debloating unnecessary software features removes gadgets found in the code implementing these features. Measuring the security benefit of debloating against gadget-based code reuse attacks is complicated because code-removing debloaters like CARVE, TRIMMER, and CHISEL may introduce new gadgets as a side effect [22]. We measure CARVE's effectiveness against gadget-based attacks using metrics proposed by Brown and Pande [22], which account for introduced gadgets. Our results are shown in the grouped columns of Table 1.

In general, CARVE is effective at reducing the computational expressivity of gadget sets with respect to three different levels of expressivity [20, 21]. In seven of the twelve scenarios, debloating reduced the overall gadget set expressivity. However, debloating had no effect in three scenarios and negative effects in two scenarios. CARVE was also generally successful at removing special purpose gadgets, achieving the highest reduction for `libcurl`. As was the case with gadget set expressivity, negative side effects were observed in two scenarios. The negative side effects observed in our results are consistent with the negative side effects observed in CHISEL [22] and TRIMMER [4]. Since CARVE is easy to reconfigure and debloating is fast, it is well suited for mitigating these negative side effects using iterative debloating [22].

Code Size Reduction. CARVE is also capable of significantly reducing the size of a software package. The average LOC, function count, and binary size reduction achieved at each level aggressiveness is displayed in Table 2. Interestingly, we observed slightly higher binary reduction than anticipated, which we attribute to improved performance by downstream compiler optimizations operating on simplified debloated code.

In some cases, debloating eliminates dependencies on external libraries. This lowers runtime memory consumption as these libraries are no longer loaded during execution. Eliminating these libraries may also have positive security impacts if they contain code and gadgets useful in a code reuse

Table 1. Summary of Security Improvements. Values Indicate Count (Number Reduced) of Expressive Classes and Gadgets.

Debloated Variant	CVEs Removed	Gadget Set Expressivity			Special Purpose Gadget Availability					
		Practical ROP Exploits	ASLR-Proof ROP Exploits	Turing Complete Exploits	Syscall Gadgets	JOP Dispatcher Gadgets	JOP Data Loader Gadgets	JOP Trampoline Gadgets	COP Dispatcher Gadgets	COP Intra Stack Pivot Gadgets
libmodbus (C)	0 of 0	6 (0)	10 (3)	6 (1)	0 (0)	0 (0)	3 (-2)	0 (0)	0 (0)	0 (0)
libmodbus (M)	0 of 0	6 (0)	13 (0)	7 (0)	0 (0)	0 (0)	1 (0)	0 (0)	0 (0)	0 (0)
libmodbus (A)	0 of 0	6 (0)	13 (0)	7 (0)	0 (0)	0 (0)	1 (0)	0 (0)	0 (0)	0 (0)
Bftpd (C)	0 of 0	7 (-1)	12 (3)	6 (1)	0 (0)	0 (0)	9 (1)	0 (0)	0 (0)	0 (0)
Bftpd (M)	0 of 0	7 (-1)	17 (-2)	6 (1)	0 (0)	0 (0)	1 (9)	0 (0)	0 (0)	0 (0)
Bftpd (A)	0 of 0	6 (0)	11 (4)	5 (2)	0 (0)	0 (0)	1 (9)	0 (0)	0 (0)	0 (0)
libcurl (C)	1 of 6	9 (0)	26 (-1)	11 (-1)	5 (-1)	7 (1)	50 (1)	0 (1)	304 (15)	4 (0)
libcurl (M)	2 of 6	9 (0)	24 (1)	10 (0)	0 (4)	4 (4)	41 (10)	0 (1)	287 (32)	3 (1)
libcurl (A)	1 of 6	10 (-1)	24 (1)	10 (0)	4 (0)	3 (5)	14 (37)	1 (0)	170 (149)	2 (2)
mongoose (C)	5 of 17	7 (0)	10 (6)	8 (0)	0 (0)	1 (-1)	6 (0)	0 (0)	0 (0)	0 (0)
mongoose (M)	1 of 17	7 (0)	10 (6)	8 (0)	0 (0)	0 (0)	6 (0)	0 (0)	0 (0)	0 (0)
mongoose (A)	8 of 17	7 (0)	10 (6)	8 (0)	0 (0)	0 (0)	6 (0)	0 (0)	0 (0)	0 (0)

exploit. We did not observe a reduction in required libraries for `libmodbus`, `Bftpd`, and `mongoose`. While conservatively debloating `libcurl` only reduced the number of required libraries from 42 to 41, we observed a *two-thirds* reduction in required libraries for moderate and aggressive variants of `libcurl` (14 and 15 required libraries after debloating, respectively).

Overhead Costs. CARVE’s decoupled design minimizes per-debloat overhead cost. Debloating is fast and efficient, taking less than 5 seconds for each scenario. Debloating also scales effectively with increasing code size and density of feature mappings. The vast majority of cost associated with this technique occurs during feature mapping. The amount of effort required to create and maintain feature mappings depends on a number of variables, including developer familiarity with the software, feature density, and complexity of the target software. In practice, we found that creating the initial feature mapping was not cost-prohibitive in situations where a developer has limited familiarity with the source code. Software engineering practices favor modularization and separation of concerns within software packages, which has the effect of reducing the number of feature mappings required to produce a correct mapping. Producing a mapping for our benchmark packages required 61.5 mappings per thousand LOC (m/KLOC) for `libmodbus`, 24.2 m/KLOC for `Bftpd`, 13.4 m/KLOC for `mongoose`, and 4.9 m/KLOC for `libcurl`. Additionally, these mappings did not require a significant number of complex code-replacing mappings; `Bftpd` and `mongoose` required no such mappings, `libmodbus` requires 17, and `libcurl` required 20. This data suggests that the one-time costs of creating feature mappings are reasonable.

Table 2. Average Code Size Reduction.

Level of Aggressiveness	LOC Reduction	Function Count Reduction	Binary Size Reduction
Conservative (C)	10.7%	11.1%	10.8%
Moderate (M)	16.1%	13.7%	18.5%
Aggressive (A)	31.3%	27.1%	33.0%

Comparison to Other Techniques. Aggressively debloating simple programs (e.g. Linux Core-Utils) with CHISEL and TOSS results in average code size reductions of 86% and 43.8% respectively [2, 3]. Aggressively debloating more complex network-based software (e.g. cUrl) with TRIMMER reduces code size by 21% on average [4]. CARVE outperforms TRIMMER on similar benchmarks, and performs comparably to CHISEL and TOSS when accounting for benchmark selection. CARVE also performs comparably to CHISEL with respect to CVE elimination [3]. Finally, CARVE reduces gadget counts and introduces new gadgets at rates comparable to CHISEL and TRIMMER [22, 4].

Limitations. CARVE operates solely on source code and cannot be used to debloat binaries directly; it is not a suitable for closed-source or legacy software. While CARVE is technology independent and can be used on any software that is text-based and supports comments, at this time it implements only one language-specific debloating module (C/C++).

4. Future Work

In future versions of CARVE, we plan to create new language-specific debloating modules for other common programming languages such as Java and Python. We also plan to develop static analysis tools to help developers quickly produce sound feature mappings. Such tools would analyze mappings as they are created to suggest new mappings and identify potential soundness issues.

5. Conclusion

In this paper, we introduced CARVE, a simple and powerful software debloating technique that overcomes the shortcomings of existing approaches. We demonstrated that effective debloating can be accomplished in a manner that preserves desirable software properties (e.g. interoperability, specification compliance) without the need for advanced software analysis or complicated test scripts. The results of our evaluation show that CARVE is well-suited for debloating complex software, improving the security and performance of four network protocol

implementations. Across 12 scenarios, CARVE eliminated vulnerabilities, reduced the utility of code reuse gadgets, and reduced overall code size on par with existing approaches using a simple approach that is accessible to average end users and developers.

ACKNOWLEDGMENTS

We would like to thank all of our reviewers for their helpful feedback, especially our shepherd Nick Nikiforakis. We also thank Andrei Homescu and his co-authors for making their microgadget scanner available for this project. Finally, we thank Joshua Kassab for his assistance with feature mapping source code.

REFERENCES

- [1] Anh Quach, Rukayat Erinfoami, David Demicco, and Aravind Prakash. 2017. A Multi-OS Cross-Layer Study of Bloating in User Programs, Kernel and Managed Execution Environments. In *Proceedings of the 2017 Workshop on Forming an Ecosystem Around Software Transformation (FEAST '17)*. ACM, New York, NY, USA, 65-70. DOI: <https://doi.org/10.1145/3141235.3141242>.
- [2] Yurong Chen, Shaowen Sun, Tian Lan, and Guru Venkataramani. 2018. TOSS: Tailoring Online Server Systems through Binary Feature Customization. In *Proceedings of the 2018 Workshop on Forming an Ecosystem Around Software Transformation (FEAST '18)*. ACM, New York, NY, USA, 1-7. DOI: <https://doi.org/10.1145/3273045.3273048>.
- [3] Kihong Heo, Woosuk Lee, Pardis Pashakhanloo, and Mayur Naik. 2018. Effective Program Debloating via Reinforcement Learning. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (CCS '18)*. ACM, New York, NY, USA, 380-394. DOI: <https://doi.org/10.1145/3243734.3243838>.
- [4] Hashim Sharif, Muhammad Abubakar, Ashish Gehani, and Fareed Zaffar. 2018. TRIMMER: application specialization for code debloating. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering (ASE 2018)*. ACM, New York, NY, USA, 329-339. DOI: <https://doi.org/10.1145/3238147.3238160>.
- [5] Anh Quach, Aravind Prakash, and Lok Yan. 2018. Debloating software through piece-wise compilation and loading. In *27th [USENIX] Security Symposium ([USENIX] Security 18)*. USENIX Association, Baltimore, MD. <https://www.usenix.org/conference/usenixsecurity18/presentation/quach>
- [6] Guoqing Xu, Nick Mitchell, Matthew Arnold, Atanas Rountev, and Gary Sevitsky. 2010. Software bloat analysis: finding, removing, and preventing performance problems in modern large-scale object-oriented applications. In *Proceedings of the FSE/SDP workshop on Future of software engineering research (FoSER '10)*. ACM, New York, NY, USA, 421-426. DOI: <http://dx.doi.org/10.1145/1882362.1882448>.
- [7] Suparna Bhattacharya, Karthick Rajamani, K. Gopinath, and Manish Gupta. 2011. The interplay of software bloat, hardware energy proportionality and system bottlenecks. In *Proceedings of the 4th Workshop on Power-Aware Computing and Systems (HotPower '11)*. ACM, New York, NY, USA, Article 1, 5 pages. DOI: <http://dx.doi.org/10.1145/2039252.2039253>.
- [8] Yufei Jiang, Can Zhang, Dinghao Wu, and Peng Liu. 2016. Feature-based software customization: Preliminary analysis, formalization, and methods. In *2016 IEEE 17th International Symposium on High Assurance Systems Engineering (HASE)*, pp. 122-131. IEEE. DOI: [10.1109/HASE.2016.27](https://doi.org/10.1109/HASE.2016.27).
- [9] Yufei Jiang, Dinghao Wu, and Peng Liu. 2016. Jred: Program customization and bloatware mitigation based on static analysis. In *2016 IEEE 40th Annual Computer Software and Applications Conference (COMPSAC)*, vol. 1, pp. 12-21. IEEE. DOI: [10.1109/COMPSAC.2016.146](https://doi.org/10.1109/COMPSAC.2016.146).
- [10] Hyungjoon Koo, Seyedhamed Ghavamnia, and Michalis Polychronakis. 2019. Configuration-Driven Software Debloating. In *Proceedings of the 12th European Workshop on Systems Security (EuroSec '19)*. ACM, New York, NY, USA, Article 9, 6 pages. DOI: <https://doi.org/10.1145/3301417.3312501>.
- [11] The LLVM Compiler Infrastructure. <https://llvm.org/>.
- [12] Angr: A python framework for analyzing binaries. <http://angr.io/>.
- [13] BitBlaze. <http://bitblaze.cs.berkeley.edu/>.
- [14] Aspire Project: ChiselBench. <https://github.com/aspire-project/ChiselBench>
- [15] Hovav Shacham. 2007. The geometry of innocent flesh on the bone: return-into-libc without function calls (on the x86). In *Proceedings of the 14th ACM conference on Computer and communications security (CCS '07)*. ACM, New York, NY, USA, 552-561. DOI: <https://doi.org/10.1145/1315245.1315313>.
- [16] Tyler Bletsch, Xuxian Jiang, Vince W. Freeh, and Zhenkai Liang. 2011. Jump-oriented programming: a new class of code-reuse attack. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security (ASIACCS '11)*. ACM, New York, NY, USA, 30-40. DOI: <https://doi.org/10.1145/1966913.1966919>.
- [17] AliAkbar Sadeghi, Salman Niksefat and Maryam Rostampour. 2018. Pure-Call Oriented Programming (PCOP): chaining the gadgets using call instructions. In *Computer Virology Hacking Techniques 14*: 139. Springer Paris. <https://doi.org/10.1007/s11416-017-0299-1>.
- [18] Chenxiong Qian, Hong Hu, Mansour Alharthi, Pak Ho Chung, Taesoo Kim, and Wenke Lee. 2019. RAZOR: A Framework for Post-deployment Software Debloating. In *28th [USENIX] Security Symposium ([USENIX] Security 19)*. USENIX Association, Baltimore, MD. <https://www.usenix.org/conference/usenixsecurity19/presentation/qian>
- [19] CVE - Common Vulnerabilities and Exposures. <https://cve.mitre.org/>.
- [20] Andrei Homescu, Michael Stewart, Per Larsen, Stefan Brunthaler, and Michael Franz. 2012. Microgadgets: size does matter in turing-complete return-oriented programming. In *Proceedings of the 6th USENIX conference on Offensive Technologies (WOOT'12)*. USENIX Association, Berkeley, CA, USA, 7-7.
- [21] Dino Zovi. 2010. Practical return-oriented programming. SOURCE Boston. <https://trailofbits.files.wordpress.com/2010/04/practical-rop.pdf>.
- [22] Michael D. Brown and Santosh Pande. 2019. Is Less Really More? Towards Better Metrics for Measuring Security Improvements Realized Through Software Debloating. In *12th [USENIX] Workshop on Cyber Security Experimentation and Test ([CSET] 19)*. USENIX Association, Baltimore, MD. <https://www.usenix.org/conference/cset19/presentation/brown>
- [23] Jonathan Salwan. 2011. ROPgadget: Gadgets finder and auto-roper. <http://shell-storm.org/project/ROPgadget/>.
- [24] Doxygen. <http://www.doxygen.nl/>
- [25] Sphinx. <http://www.sphinx-doc.org/en/master/>.
- [26] Coverity SAST. <https://www.synopsys.com/software-integrity/security-testing/static-analysis-sast>.