

# D-Linker: Debloating Shared Libraries by Relinking From Object Files

Jiatai He<sup>1b</sup>, Pengpeng Hou<sup>1b</sup>, Jiageng Yu, Ji Qi<sup>1b</sup>, Ying Sun<sup>1b</sup>, Lijuan Li, Ruilin Zhao<sup>1b</sup>, and Yanjun Wu<sup>1b</sup>

**Abstract**—Shared libraries are widely used in software development to execute third-party functions. However, the size and complexity of shared libraries tend to increase with the need to support more features, resulting in bloated shared libraries. This leads to resource waste and security issues as a significant amount of generic functionality is included unnecessarily in most scenarios, especially in embedded systems. To address this issue, previous works attempt to debloat shared libraries through binary rewriting or recompilation. However, these works face a tradeoff between flexibility in usage (needs recompilation and runtime support) and the effectiveness of debloating (binary rewriting achieves insufficient file size reduction). We propose D-Linker, a tool that debloats shared libraries by reducing both code and data sections in link-time at the object level without recompilation. Our key insight is that object-level shared library debloating is especially suitable for embedded systems because it strikes a balance of flexibility and efficiency. D-Linker identifies the required ELF object files of the shared libraries in an application and relinks them to produce a debloated shared library with better-debloating effectiveness by avoiding the data reference analysis. Our approach achieves over 70% of gadgets reduction as a security benefit and an average size reduction of 49.6% for a stripped libc of coreutils. The results also indicate that D-Linker improves debloating effectiveness by approximately 30% compared to binary-level shared library debloating and incurs a 5% decrease in code gadgets reduction compared to source-code-level shared library debloating.

**Index Terms**—Binary debloating, embedded system, shared library.

## I. INTRODUCTION

SHARED libraries play a crucial role in modern software development by enabling code reusability and modularity [1], [2], [3]. As opposed to static libraries, which are susceptible to space wastage and complicated updates, shared libraries provide several benefits: 1) sharing a single module

Manuscript received 12 August 2024; accepted 12 August 2024. This work was supported in part by the National Key Research and Development Program of China under Grant 2023YFB4503905, and in part by the HKU-CAS Joint Laboratory for Intelligent System Software. This article was presented at the International Conference on Embedded Software (EMSOFT) 2024 and appeared as part of the ESWEK-TCAD special issue. This article was recommended by Associate Editor S. Dailey. (Jiatai He and Pengpeng Hou contributed equally to this work.) (Corresponding author: Ji Qi.)

Jiatai He and Ruilin Zhao are with the Institute of Software, Chinese Academy of Sciences, Beijing 100190, China, and also with the University of Chinese Academy of Sciences, Beijing 101408, China (e-mail: jiatai21@iscas.ac.cn; zhaoruilin22@mails.ucas.ac.cn).

Pengpeng Hou, Jiageng Yu, Ji Qi, Ying Sun, Lijuan Li, and Yanjun Wu are with the Institute of Software, Chinese Academy of Sciences, Beijing 100190, China (e-mail: pengpeng@iscas.ac.cn; jiageng08@iscas.ac.cn; qiji@iscas.ac.cn; sunying@nj.iscas.ac.cn; lilijuan@iscas.ac.cn; yanjun@iscas.ac.cn).

Digital Object Identifier 10.1109/TCAD.2024.3446712

among multiple programs eliminates the need to embed the module in each program, thus minimizing space wastage and 2) shared libraries only require an interface, without the need to recompile the program during program updating, which simplifies the process and increases module independence.

However, the proliferation of shared libraries in diverse applications [4] increases their size and complexity, because of supporting for new features while maintaining the old ones. For example, the shared library referenced by the CUDA toolkit has grown almost five times in size from 2012 to 2021 [5]. These bloated portions lead to shared library bloating and cause considerable waste of memory and disk space and security risks, as loading a library involves mapping the entire library into memory, including numerous unused functions and data. A study conducted on a diverse set of 2016 programs across various domains demonstrated that 95% of glibc code is never used [6].

Shared library bloating is a major concern, particularly in scenarios with limited memory and disk storage capabilities like embedded systems [7] or containers. For instance, the size of the Alpine lightweight image that employs musllibc is only 5.59MB, while shared libraries account for 3.9 MB [8]. The bloating of shared libraries not only leads to the waste of storage but also to the inefficient use of memory. Lee et al. [9] demonstrated that loading an entire library, including its unused components, into memory can lead to 10% RAM memory wastage and 35% flush memory wastage in non-MMU embedded systems due to the single address space. Even systems equipped with an MMU can face memory waste due to the 4KB page size and the read-ahead [10] mechanism, which often results in the entire shared library being loaded into physical memory. From the security perspective, failure to address code bloat in shared libraries also frequently leads to return-oriented programming (ROP) gadgets [11] that can be stitched together by the attacker to create malicious attacks. Certain critical vulnerabilities remain exploitable by attackers despite the broadly deployed defenses [12]. To mitigate these issues, the debloating of shared libraries of resource-constrained scenarios like embedded systems and container images is essential.

In resource-constrained embedded systems, different systems support diverse functionalities, while each specific system, such as those in the automotive or avionics industries, typically requires only a subset of the functions provided by shared libraries [13]. This necessitates the debloating of shared libraries to address two primary considerations: 1) due to the limited set of functions required by each specific

83 system, redundancy in shared libraries becomes significant,  
 84 necessitating shared library debloating and 2) due to the  
 85 diverse requirements of different types of embedded systems,  
 86 the effort and computational cost associated with customizing  
 87 and debloating code for each specific system are substantial.  
 88 This necessitates efficient mechanisms for binary debloating.

89 Existing work on shared library debloating can be catego-  
 90 rized into two types: 1) source-code-level debloating [6], [14],  
 91 [15], [16] and 2) binary debloating [13], [17], [18].

92 The source-code-level debloating involves shared library  
 93 debloating based on the source code, offering precise debloat-  
 94 ing but incurring the need of source code and associated costs  
 95 of compilation stage (e.g., lack of flexibility, need source code  
 96 and runtime support). For example, piecewise compilation and  
 97 loading (PCL) [6] embeds call graph information extracted  
 98 from source code in binaries and uses a custom loader  
 99 to rewrite unnecessary functions with invalid instructions at  
 100 runtime, whereas BlankIt [14] and Trimmer [15] leverage  
 101 additional analysis during compilation and load only necessary  
 102 functions into the program's memory at runtime based on the  
 103 current execution point.

104 The binary-level debloating does not need the source  
 105 code, achieving high flexibility but at the cost of debloating  
 106 effectiveness. These works use a variety of transforma-  
 107 tion approaches, including creating specialized libraries via  
 108 rewriting, to remove or blank unused functions with no-ops  
 109 or invalid instructions [13], [17], fragmenting libraries [19],  
 110 relocating symbols to statically link library functions [20],  
 111 and replacing unnecessary library functions [21]. However,  
 112 these works often face limitations in control flow analysis  
 113 at the binary level [22], [23], [24], leading to the need of  
 114 additional information (e.g., debug information [17] or specific  
 115 ISAs [18]) and the insufficient reduction in file size.

116 In this study, our key insight is that object-level debloating  
 117 is especially suitable for embedded system shared library  
 118 debloating, because it strikes a balance of flexibility and  
 119 debloating effectiveness.

120 First, object-level debloating offers higher flexibility  
 121 compared to source-code-level debloating. In embedded envi-  
 122 ronments, there is often a large amount of invalid or outdated  
 123 source code that cannot support additional compilation and  
 124 runtime mechanisms. In such cases, object-level debloating  
 125 does not face these issues: 1) static library files are often  
 126 more readily available than source code, allowing object-  
 127 level debloating to use static libraries to generate debloated  
 128 shared libraries and 2) object-level debloating is typically more  
 129 efficient, as the cost of relinking is much lower than the cost  
 130 of recompilation.

131 Second, object-level debloating can erase more unused  
 132 content compared to binary-level debloating, resulting in  
 133 better-debloating effectiveness. As binary rewriting cannot  
 134 resolve data references [22], [23], [24], binary-level debloat-  
 135 ing only erases the functions in the code section (`.text`  
 136 section) and cannot resolve the data reference for erasing the  
 137 unused data in the data section [18]. Moreover, binary over-  
 138 writing erases the code section by memory page granularity,  
 139 so code erasing is limited to 4 KiB page size, which would  
 140 leave some unused code in the library and affect the file

size of the debloated library [13]. Compared to binary-level 141  
 (e.g., Nibbler [17], ELFtailor [13],  $\mu$ Trimmer [18]), debloating 142  
 shared library by object-level granularity removes unused ELF 143  
 object files rather than erases binary code, which reducing not 144  
 only the code sections [25] (e.g., `.text`) but also the data 145  
 sections (e.g., `.data`) and other sections (e.g., `.rela.dyn`), 146  
 resulting in further size reduction effectiveness improvement. 147

148 Finally, in practical scenarios, when faced with version 148  
 modifications or functional updates, object-level debloat- 149  
 ing presents several advantages over both binary-level 150  
 and source-code-level debloating. Unlike source-code-level 151  
 debloating, object-level debloating does not require source 152  
 code or additional runtime mechanisms when updating system 153  
 functionalities, significantly reducing the cost of updates. 154  
 Additionally, binary-level debloating needs modifications to 155  
 the code sections and ELF segment structures, which compro- 156  
 mise the maintainability of the shared library by increasing 157  
 the difficulty in debugging and introducing uncertainties in the 158  
 ELF format. In contrast, shared libraries produced by object- 159  
 level debloating are indistinguishable in format from native 160  
 shared libraries, minimizing additional maintenance costs. 161

162 Our insight leads to D-Linker, a tool which debloats shared 162  
 libraries in link-time at object-level by relinking the ELF 163  
 object files to the shared libraries. The debloating process 164  
 of D-Linker has two modes: 1) normal debloating mode and 165  
 2) in-depth debloating mode. In normal debloating mode, D- 166  
 Linker performs binary analysis on the target binary file to 167  
 generate its function call graph (FCG). Based on the FCG, 168  
 D-Linker identifies the used object files and then relinks 169  
 them to create the debloated shared library file with both 170  
 unused code and data sections eliminated. At the same time, 171  
 if certain fixed functionalities of the usage scenario are 172  
 determined, a better-debloating effectiveness can be achieved. 173  
 Therefore, we propose an in-depth debloating approach, which 174  
 involves dynamic tracking of shared libraries based on a 175  
 specific set of user-defined functional tests to achieve a 176  
 deeper debloating and ultimately obtain a better-debloating 177  
 effect. 178

179 D-Linker faces two challenges: 1) how to handle data 179  
 dependencies and 2) the completeness of symbol dependency. 180  
 These issues are typically resolved in normal debloating 181  
 through inherent symbol dependencies. However, in the case 182  
 of in-depth debloating, due to the inability to track data 183  
 symbols dynamically and the incompleteness of symbols 184  
 during linking, additional mechanisms are required to address 185  
 these challenges. We have adopted an object data dependency 186  
 analysis method to handle the data dependency issue in 187  
 in-depth debloating and use binary modification to address 188  
 symbol incompleteness in in-depth debloating. 189

190 We summarize our major contributions as follows. 190

- 191 1) We introduce D-Linker, an innovative technique for 191  
 debloating shared libraries by relinking from ELF object 192  
 files. Our approach offers the ability to reduce the size 193  
 of not only the code section but also the data and other 194  
 sections, achieving high-reduction effectiveness. 195
- 196 2) D-Linker efficiently leverages static libraries to extract 196  
 object files during the debloating process of most shared 197  
 libraries, resulting in no need of compilation and source 198

code (with a few exceptions, as elaborated in Section V), achieving high flexibility.

- 3) We evaluated D-Linker and prior methods for debloating shared libraries, focusing on key aspects, such as debloating effectiveness (i.e., file size reduction), runtime overhead, and security analysis. The results indicate that D-Linker improved debloating effectiveness by approximately 30% compared to the binary-level debloating and incurred only a 5% decrease of code gadgets reduction compared to the source-code-level debloating.

## II. BACKGROUND AND RELATED WORK

Binary debloating has been a hot topic in security and embedded systems. Mulliner and Neugschwandtner [26] first proposed the specialization of shared libraries for runtime environments, with the core idea of keeping the functions needed in the shared libraries for different scenarios and removing the functions not needed in the shared libraries at load time. Existing works can be categorized into the following two approaches.

*Source-Code-Level Debloating:* Quach et al. [6] proposed piecewise using llvm compiler for a compiler and loader-assisted mechanism to reduce executable code loaded by a process. They analyze the indirect points reference by interprocedure static value-flow analysis and saving a full-program dependency graph. Then they load the functions that are present in the dependency graph at load time. This approach worked well in removing useless code, but masking functions at load time incurs additional performance overhead, requiring source code and recompiling. Similar works, BlankIt proposed by Porter et al. [14] and Trimmer by Sharif et al. [15] BlankIt feed the functions that the user needs to call into a linked library using a decision tree predictor to reduce the exposed surface of the code, which reduces the code exposure surface by an average of 96%, but also increases the average runtime overhead by 18%.

*Binary-Level Debloating:* The benefit of debloating shared libraries by binary analysis and erasing is to remove unused code without source code and recompilation, which can be used in more scenarios than debloating during compilation. Ziegler et al. [13] proposed ELFtailor, which debloats shared libraries in two steps: 1) running binary static analysis using capstone [27] and dynamic analysis using uprobes [28] to obtain the FCG of the application and 2) overwriting and erasing the functions not in the FCG and reorganizing the shared libraries. A similar work is proposed by Ioannis et al. Nibbler [17], which obtains FCG of shared library unstripped statically. Then, Zhang et al. [18] proposed an MIPS-based shared library debloating tool  $\mu$ Trimmer. Due to the MIPS instruction set, indirect pointer calls to functions can be thoroughly analyzed statically in shared libraries stripped, which is a big step forward in binary analysis.  $\mu$ Trimmer almost reaches the upper limit of the code debloating of shared libraries. At the same time, these works also mention that although the code section is debloated, they still cannot resolve

data reference, which means binary rewriting cannot debloat the data section of the shared library.

## III. OVERVIEW OF D-LINKER

### A. Overview

We present D-Linker, a tool that aims to debloat shared libraries by relinking the ELF object files.

*Debloating Modes:* Our debloating process has two modes: 1) normal debloating mode and 2) in-depth debloating mode. In the normal debloating mode, the set of object files that are used is selected based on the dependencies between them, and then relinked to create the debloated shared library with both unused code and data sections eliminated. However, during the debloating process, there may still be unused object files in the set due to the presence of redundant dependencies. In order to further improve the debloating effect, we introduce the in-depth debloating mode. This mode involves leveraging the test suite of the target program and employing binary rewriting techniques to identify and eliminate unused object files. By employing this approach, we aim to achieve more effective debloating effectiveness. We will conduct a more detailed analysis of this in following sections.

*Challenges:* To achieve the above design, we encountered two challenges, which will be discussed in Section IV-B.

*Challenge 1 (Data Reference):* In the process of in-depth debloating, we need to erase the references to the unused object files identified by dynamic analysis. However, dynamic analysis is inherently limited to detecting called functions and does not adequately address data references. This limitation presents a significant challenge: How can we precisely identify object files that contain critical data references and therefore cannot be removed during in-depth debloating? To address this challenge, we employ a method to handle data reference when removing object files.

*Challenge 2 (The Completeness of Symbol Dependency):* Another challenge is to address the issue of symbol dependency completeness, which primarily needs to be resolved in in-depth debloating. Relinking the object files tracked by dynamic analysis requires breaking dependencies between some object files to achieve a better-debloating effectiveness. To address this issue, we employ binary rewriting to modify symbol attributes, thereby breaking dependencies between certain objects.

*Overview:* Our methodology consists of three main steps. First, we generate the FCG by conducting static and dynamic binary analysis, with dynamic analysis performed for in-depth debloating mode. Second, we identify the used object files and their dependencies. Finally, we modify and relink the necessary object files to generate the debloated library. Fig. 1 provides a visual representation of our approach, which comprises three components: 1) FCG reconstruction; 2) identifying object files; and 3) relinking. Together, these components involve 7 steps, with steps 3 and 6 only applicable in cases where in-depth debloating is necessary.

*FCG Reconstruction—Steps 1–4:* To reconstruct the FCG of the application and shared library, we employ a four-step process. First, we analyze the dependency between shared

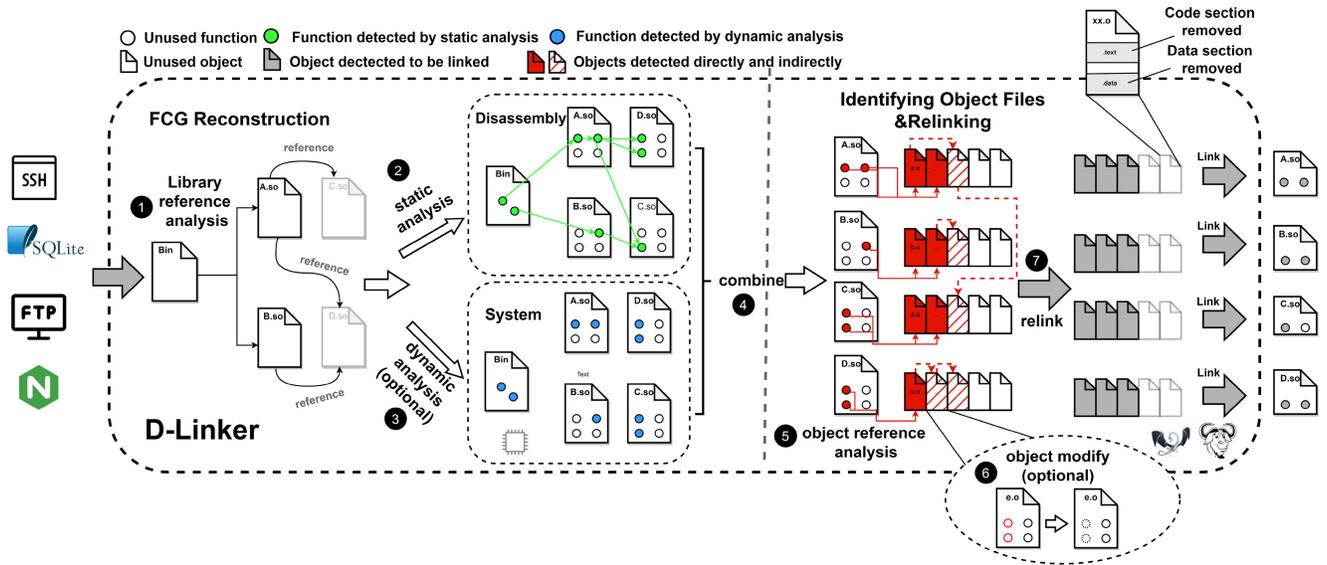


Fig. 1. Overview of D-Linker.

310 libraries by identifying the required shared libraries and their  
 311 dependencies based on the `.dynamic` section of binary  
 312 (step 1, Fig. 1 ①). Second, we perform static binary anal-  
 313 ysis and disassemble the application and shared libraries to  
 314 generate the FCG of the shared libraries (step 2, Fig. 1 ②).  
 315 Third, if in-depth debloating is required by user, we validate  
 316 the used functions of the shared libraries through dynamic  
 317 analysis using uprobes (step 3, Fig. 1 ③). Finally, we merge  
 318 the results from Step 2 and Step 3 (if applicable) to obtain a  
 319 comprehensive FCG (step 4, Fig. 1 ④).

320 *Identifying Object Files—Step 5:* After establishing the FCG  
 321 of the shared libraries in Step 4, the subsequent step is to  
 322 identify the used object files based on the guaranteed FCG. In  
 323 cases where in-depth debloating is required, we also pinpoint  
 324 the redundant symbol dependency between object files (step  
 325 5, Fig. 1 ⑤).

326 *Relinking—Steps 6 and 7:* The relinking process consists of  
 327 two steps. First, if in-depth debloating is necessary, we modify  
 328 the object files identified in Step 5 and eliminate any redundant  
 329 symbol dependencies (step 6, Fig. 1 ⑥). Finally, we relink  
 330 the used object files to generate debloated shared libraries,  
 331 discarding unused object files that contain both unused data  
 332 and code (step 7, Fig. 1 ⑦).

333 *Debloating Shared Libraries of the Total System:* Debloating  
 334 shared libraries within an total system (e.g., IoT devices,  
 335 containers) is a common use case of D-Linker, particularly in  
 336 the context of embedded systems that encompass numerous  
 337 binaries. D-Linker takes all the binaries of the embedded  
 338 system as input, following a procedure analogous to the  
 339 debloating of a single binary. The resultant debloated shared  
 340 libraries are relinked using the union of the object files  
 341 requisite for all binaries. In Section VI, we demonstrate the  
 342 efficacy of D-Linker in such scenarios by debloating the  
 343 shared libraries of an Alpine Linux system running `vsftpd`,  
 344 showcasing its significant effectiveness improvements in this  
 345 context.

## IV. DETAILED DESIGN OF D-LINKER

### A. FCG Reconstruction

346  
 347  
 348 The reconstruction of the FCG involves three steps. First, a  
 349 shared libraries dependency analysis is performed to discover  
 350 all required shared libraries and their dependencies relevant  
 351 to the target application. This step ensures a comprehensive  
 352 understanding of the libraries required. Second, a static binary  
 353 analysis is conducted using disassembly tools to establish the  
 354 FCG of the shared libraries. However, it should be noted  
 355 that static analysis tools may not be able to detect all used  
 356 interfaces. Thus, further runtime dynamic analysis is necessary  
 357 when in-depth debloating is required and the result of dynamic  
 358 analysis would be combined with the FCG established by static  
 359 analysis.

360 *Shared Libraries Dependency Analysis:* The method used  
 361 for shared libraries dependency analysis involves two steps.  
 362 First, after loading the application, D-Linker detects the  
 363 dependent shared libraries based on the information stored in  
 364 `DT_NEEDED`, which is an item in the `.dynamic` section.  
 365 This step ensures that the directly dependent shared libraries  
 366 are identified. Second, the shared libraries identified in the  
 367 previous step recursively search for their own dependent  
 368 shared libraries using the same method. This recursive process  
 369 continues until a complete shared library dependency graph is  
 370 established, encompassing all the dependencies.

371 Once the shared library dependency graph is obtained, D-  
 372 Linker proceeds to detect the undefined symbols in both  
 373 the shared libraries and the application. This step involves  
 374 identifying symbols that are referenced but not defined within  
 375 the code. By identifying these undefined symbols, D-Linker  
 376 is able to gather detailed reference information about how the  
 377 shared libraries rely on each other. This information helps in  
 378 understanding the relationships and interactions between the  
 379 shared libraries, which is crucial for the debloating process.

380 *Static Binary Analysis:* During static binary analysis, the  
 381 initial step involves disassembling both the applications and

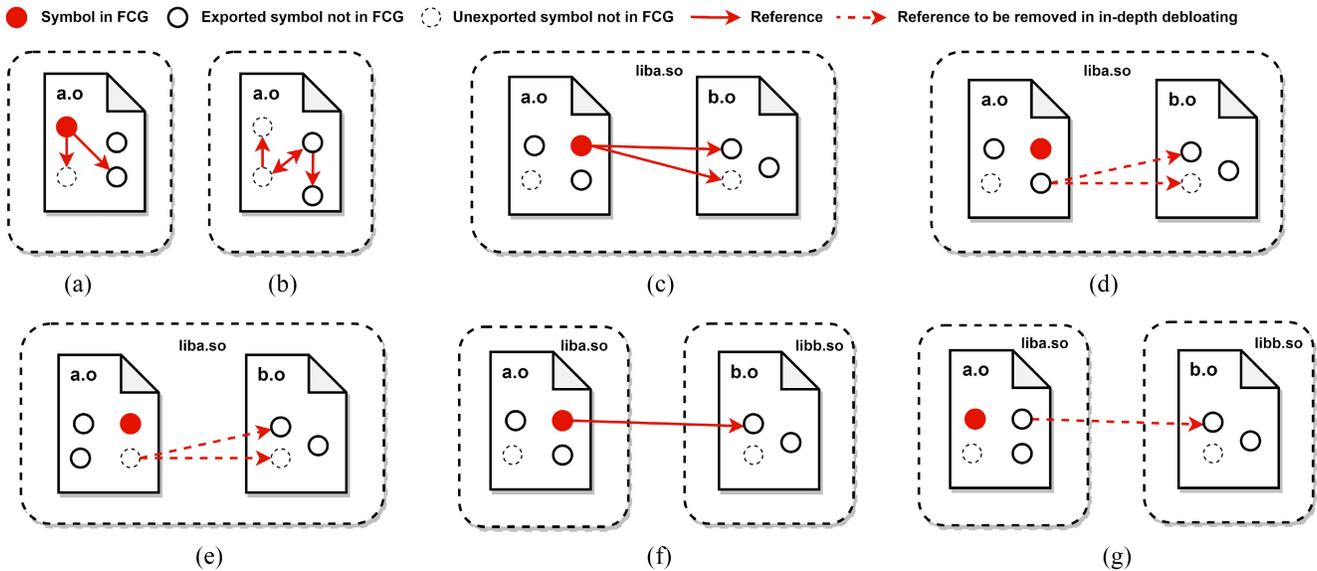


Fig. 2. Reference of object files. (a) Symbols in FCG require symbols in same object. (b) Symbols not in FCG require symbols in same object. (c) Symbols in FCG require symbols in another object but in the same shared library. (d) Default global symbol not in FCG require symbols in another object but in the same shared library. (e) Hidden global symbol not in FCG require symbols in another object but in the same shared library. (f) Symbols in FCG require symbols in another shared library. (g) Symbols not in FCG require symbols in another shared library.

shared libraries. Subsequently, the call relationships within the execution path of the binary file are determined by examining its disassembly code and symbolic information. These findings are then utilized to reconstruct the FCG. It is important to note that the FCG solely encompasses global symbols from the shared libraries, as the handling of local symbol references is performed by our object-level debloating technique.

The FCG described here is adequate for generating necessary object files to produce normal debloating shared libraries, thanks to the symbolic dependencies. However, in cases where in-depth debloating is needed, dynamic analysis becomes necessary.

*Dynamic Binary Analysis:* Dynamic binary analysis serves as a complementary counterpart to static binary analysis by detecting all used functions through runtime dynamic analysis using probes. To detect functions called by function pointers and virtual functions in C++ that cannot be found through static analysis, probes are deployed to all the global symbols of shared libraries required by the application. The probes are triggered upon executing the application's test suite, resulting in a list of called functions. This list is then merged with the FCG obtained through static binary analysis to yield the final FCG for the shared libraries. To obtain a precise representation of the target use case, it is necessary for the dynamic analysis test suite to cover all functionalities essential in the final system. This goal can be achieved by adhering to strictly defined procedures for interacting with the system and utilizing extensive test suites, as required for certification purposes. This requirement is particularly relevant in industries, such as automotive engineering.

Dynamic analysis is employed to facilitate in-depth debloating, which aims to eliminate unnecessary runtime dependencies and improve the effectiveness of debloating outcomes. However, it is important to note that dynamic

analysis relies on the coverage provided by the test suite. Consequently, it cannot guarantee the proper functioning of functions that are not covered by the tests. This limitation can result in certain stability compromises.

## B. Relinking

After obtaining the FCG in shared libraries, we select the used object files based on the FCG to relink the debloated shared libraries. In this section, we will discuss how D-Linker handles the reference between object files especially the data reference (Challenge 1) and relinks the object files with incomplete symbol dependency (Challenge 2).

*Symbols in Object Files:* Symbols within object files are classified into two types: 1) symbols included in the FCG and 2) symbols not included in the FCG, as shown in Fig. 2.

*Resolve Challenge 1:* We employ a method to generate a set of used object files with safe data reference. The method for generating a set of used object files is outlined as follows:

First, we categorize the dependencies into two types: 1) data dependencies, which refer to references to data type symbols and 2) function dependencies, which refer to references to function type symbols.

Second, we generate a full dependency set, encompassing both function and data dependencies, and a data dependency set for each object file. This process is described in Algorithm 1 as follows: for every object file, we add its dependent object files to the dependency set, continuing this process for all object files in the set until there are no more references to external object files. When constructing the full dependency set, we consider both function and data dependencies. However, when forming the data dependency set, we focus exclusively on data dependencies.

---

**Algorithm 1** Get Full Dependency Set and Data Dependency Set
 

---

```

Require: Object files of a shared library  $O$ ;
Ensure: full dependency set  $F$  and data dependency set  $D$ ;
1: function DFSFINDDEPENDENCY( $o$ ,  $dependency$ )
2:   if  $dependency == full$  then
3:      $F(o).add(o)$ 
4:     /*  $x$  is a object file referenced by  $o$  */
5:     for  $x \in fullreference(o)$  do
6:       if  $x \notin F(o)$  then
7:         DFSFINDDEPENDENCY( $x$ ,  $full$ )
8:          $F(o).add(F(x))$ 
9:       end if
10:    end for
11:  else if  $dependency == data$  then
12:     $D(o).add(o)$ 
13:    /*  $x$  is a object file data referenced by  $o$  */
14:    for  $x \in datareference(o)$  do
15:      if  $x \notin D(o)$  then
16:        DFSFINDDEPENDENCY( $x$ ,  $full$ )
17:         $D(o).add(D(x))$ 
18:      end if
19:    end for
20:  end if
21: end function
22: for  $o$  in  $O$  do
23:   DFSFINDDEPENDENCY( $o$ ,  $full$ )
24:   DFSFINDDEPENDENCY( $o$ ,  $data$ )
25: end for

```

---

448 Finally, the used object set in normal debloating consists  
 449 of the object files that contain the symbols appearing in the  
 450 FCG and their associated full dependency set. In the context of  
 451 in-depth debloating, the used object set comprises the object  
 452 files for all symbols present in the FCG, obtained from both  
 453 static and dynamic analysis, as well as their corresponding  
 454 data dependency sets.

455 It is significant to highlight the dependencies that were elim-  
 456 inated during the debloating process. We specifically consider  
 457 single-step references between object files, as recursive single-  
 458 step references can resolve the overall dependency. In Fig. 2,  
 459 the references to symbols within object files are depicted.  
 460 Fig. 2(a) and (b) correspond to references within the same  
 461 object file, but these do not need to be taken into account due  
 462 to the object-level granularity of debloating. Fig. 2(c) contains  
 463 two types of symbol references in the FCG: 1) references  
 464 to symbols exported by the shared library in other objects  
 465 and 2) references to symbols not exported by the shared  
 466 library in other objects. These references will be retained after  
 467 debloating. Additionally, Fig. 2(d) and (e) illustrate additional  
 468 references of symbols that are not present in the FCG but exist  
 469 within the same object file as the symbols in the FCG. These  
 470 references arise due to the object-level debloating granularity  
 471 and will be kept in normal debloating. In cases where in-depth  
 472 debloating is required, these references will be eliminated  
 473 through binary rewriting. When dealing with cross-shared  
 474 libraries, it is necessary to distinguish between the references

TABLE I  
REDUCTION OF LIBRARY TOTAL SIZE

Program	File size		
	Baseline	Tailored	Reduction
<b>Nginx</b>			
libc	612KB	169.8KB	-72.3%
libcrypto	1.8MB	265.7KB	-85.3%
libpcre	153.4KB	120.6KB	-21.4%
libz	117KB	59.2KB	-49.5%
<b>Coreutils</b>			
libc	612KB	325KB	-46.9%
<b>SqLite</b>			
libc	612KB	170.2KB	-72.1%
libz	117KB	88KB	-25.8%
libreadline	379KB	370.6KB	-2.3%
libncurses	398KB	131.1KB	-66.9%
<b>Openssh</b>			
libc	612KB	198.9KB	-67.5%
libcrypto	1.8MB	1.7MB	-5.6%
libz	117KB	96.8KB	-17.3%
<b>Vsftpd</b>			
libc	612KB	166.1KB	-72.9%
libcrypto	1.8MB	1.5MB	-16.7%
libssl	406KB	375.2KB	-7.6%
<b>Vsftpd(in-depth)</b>			
libc	612KB	166.1KB	-72.9%
libcrypto	1.8MB	1.1MB	-38.9%
libssl	406KB	288.2KB	-29.1%
<b>Alpine(in-depth)</b>			
libc	612KB	358.6KB	-41.5%
libcrypto	1.8MB	1.1MB	-38.9%
libssl	406KB	288.2KB	-29.1%

from symbols in the FCG to other shared libraries (Fig. 2(f) 475  
 should not be removed) and the references from symbols 476  
 not in the FCG to other shared libraries (Fig. 2(g) should 477  
 be removed). In normal debloating, Fig. 2(g) will also be 478  
 retained, but this kind of reference will be removed through 479  
 binary rewriting when in-depth debloating is required. 480

*Resolve Challenge 2:* To remove the references in in-depth 481  
 debloating, we need to perform binary rewriting to change 482  
 the attribute of undefined symbols which reference symbols 483  
 in removed object files from “undefined” to “weak”. 484  
 This is necessary to ensure that we avoid causing “symbol 485  
 undefined” errors when we load the library linked from 486  
 the object files into memory. After performing the binary 487  
 rewriting, we can proceed to relink the object files and obtain 488  
 the debloated library. 489

### C. Data Reference (Compared to Binary-Level Debloating) 490

One notable advantage of D-Linker over previous binary- 491  
 level debloating methods is that we utilize object files to 492  
 debloat various sections in binary besides code section, such 493  
 as data sections and relocation tables. In this section, we will 494  
 introduce how D-Linker reduces the size of the data sections in 495  
 shared libraries. 496

*Limitation to Binary Analysis Debloating Tools:* Prior 497  
 debloating tools that rely on binary analysis and rewriting 498  
 are unable to reduce data sections due to several challenges. 499  
 One of the primary difficulties lies in the fact that runtime 500  
 operations often involve references to data and updates to 501  
 pointers. Existing binary analysis techniques have difficulty in 502  
 providing an accurate analysis of the indirect control flow of 503  
 data access, making it challenging to reduce data sections [18], 504

[22], [23], [24]. Additionally, conventional methods, such as deploying probes, that set breakpoint instructions at the function entry point are insufficient to detect access to data blocks. These limitations make it challenging to resolve data references using symbol tables and disassembly code with existing techniques. Hence, debloating data sections with such tools has been a significant challenge.

*Why D-Linker can Debloat Data Sections?* D-Linker employs an innovative method to debloat shared libraries through relinking from object files. This allows direct debloating of data sections and circumvents the complexities of analyzing control flows of data access. In object files, functions access data either via symbolic tables from other object files or within the object itself. Upon linking a shared library, these data blocks consolidate, leading to runtime calculations and updated pointers, complicating reference resolution via binary analysis.

D-Linker addresses this by utilizing symbol reference analysis between object files, eliminating the need to examine consolidated data blocks within shared libraries. Additionally, by debloating at the object level, D-Linker mitigates concerns about data access within object file. Consequently, by selecting necessary data blocks and discarding redundant object files before linking, D-Linker effectively reduces data section sizes without delving into indirect control flow analysis.

## V. IMPLEMENTATION

In this section, we provide details on the implementation of D-Linker.

*Binary Analysis:* We use static and dynamic binary analysis to generate the complete FCG: the static analysis part uses capstone [27] and objdump for binary disassembly; the dynamic analysis part is done by using uprobes to get the triggered functions. We use python and nm [29] to read the symbol table of ELF files and resolve dependencies between object files. For the shared library files opened by `dlopen()`, D-Linker needs to know in advance the shared libraries and their object file locations that need to be opened. Static analysis cannot guarantee the correctness of its FCG, and dynamic analysis is required to ensure its correctness.

*Obtain the Object Files:* Due to the fact that the majority of shared libraries have corresponding static libraries, D-Linker can extract object files from the static libraries, which enables D-Linker to debloat shared libraries without source code in most cases. In our evaluation, aside from `libc`, which required additional object files for dynamic linking (e.g., `dlstart.o`), all other shared libraries were debloated by relinking the object files extracted from the static libraries without the source code.

## VI. EVALUATION

We conducted an evaluation of D-Linker on Ubuntu 20.04, which is a Linux operating system running on x86-64 architecture. To test the effectiveness, we executed various applications, including Nginx (v1.9.2), Coreutils (v9.1), Sqlite (v3.40.0), Openssh (v7.3), and Vsftpd (v3.03), to evaluate their effectiveness in reducing bloat. All of the aforementioned

applications were implemented using musllibc (v1.2.3) [30] as the C library.

For in-depth debloating, we selected Vsftpd, which is a lightweight and secure FTP server for Unix-like systems, used in debloating evaluation of prior work [13] and a complete Alpine [8] system which is often used in docker and IoT as our evaluation targets. We developed a testing script for Vsftpd (both normal and in-depth debloating), encompassing various tasks, including user login, logout, file upload, and download, with SSL transfers facilitated by Vsftpd. During the dynamic analysis phase, this script was deployed on both a standalone Vsftpd application and a whole Alpine system equipped with Vsftpd. Following the debloating procedure, we persistently employed this script as a test case to authenticate the accuracy of our in-depth debloating methodology.

Following is the structure of this section. In Section VI-A, we analyze the file size of debloated shared libraries. In Section VI-B, we discuss the effectiveness of debloating object files and functions, evaluating the dependencies of the object files. In Section VI-C, we analyze the debloating of data sections. In Section VI-D, we carried out an investigation into the dependencies among object files within multiple shared libraries to conduct further analysis of our debloating results. In Sections VI-E and VI-F, we evaluate the security benefits of D-Linker by analyzing reduction of gadgets and compare D-Linker with related works. In Section VI-F, we quantitatively compared the code debloating capabilities of D-Linker and ELFTailor [13], as well as the security benefits observed with D-Linker and piecewise. Finally, we evaluated D-Linker's effectiveness against other debloating techniques in terms of file size reduction, runtime availability, and security analysis, demonstrating D-Linker's superior advantages in embedded scenarios.

### A. Reduction of File Size

*Reduction of Normal Debloating:* Table I presents the file sizes of shared libraries for both debloated and baseline (stripped) versions across various target applications (the unannotated applications are normal debloating, and the same applies to the following tables). The Program column denotes the shared libraries associated with each target application case, noting that Vsftpd does not debloat `libgcc_s`. D-Linker achieved the highest reduction in Nginx, minimizing 77.4% of its shared library size, whereas Openssh had the lowest reduction at 21.2%. The variation in debloating efficacy among different applications primarily stems from the disparate effectiveness in debloating their dependent libraries, which we will elaborate on. Specifically, D-Linker consistently performed best with `libc` across all cases, achieving a reduction of 72.3% for Nginx and 49.6% for Coreutils. This efficacy is attributed to musllibc's extensive set of functions, including multithreading and mathematical libraries, which are largely unused and thus removable. Other notable results include an 86.4% reduction in `libcrypto` for Nginx. In contrast, debloating is less effective for libraries closely aligned with application requirements; for example, `libreadline` used in Sqlite was only reduced by 3%. This limited reduction is due to the significant

TABLE II  
REDUCED OBJECT FILES OF THE LIBRARY

Program	Number of object files		
	Baseline	Tailored	Reduction
<b>Nginx</b>			
libc	1341	366	-72.7%
libcrypto	556	25	-95.5%
libpcrc	22	9	-59.1%
libz	16	5	-68.8%
<b>Coreutils</b>			
libc	1341	464	-65.4%
<b>Sqlite</b>			
libc	1341	377	-71.9%
libz	16	10	-37.5%
libreadline	35	33	-5.7%
libncurses	149	31	-79.2%
<b>Openssh</b>			
libc	1341	456	-66.0%
libcrypto	556	404	-27.3%
libz	16	8	-50.0%
<b>Vsftpd</b>			
libc	1341	349	-74.0%
libcrypto	556	390	-29.9%
libssl	46	38	-17.4%
<b>Vsftpd(In-depth)</b>			
libc	1341	349	-74.0%
libcrypto	556	233	-58.1%
libssl	46	24	-47.8%
<b>Alpine(In-depth)</b>			
libc	1341	508	-62.1%
libcrypto	556	233	-58.1%
libssl	46	24	-47.8%

functional overlap between the library and application, coupled with the dense distribution of functions and data blocks within the library's object files, leading to a retention of many superfluous functions.

*Reduction of Indepth Debloating:* The results of the in-depth debloating experiment demonstrate a significant improvement in the overall reduction effect of vsftpd, by approximately 23.8%, compared to normal debloating. Regarding the reduction of libraries, the libcrypto library was identified as the most significant, resulting in an additional reduction of approximately 38%. The libssl library also achieved a reduction of above 29.1%, which achieved an implementation above 20% compared to normal debloating.

Overall, D-Linker's normal debloating and in-depth debloating achieved satisfactory results in most cases, demonstrating D-Linker's effectiveness in file size reduction. However, certain shared libraries (e.g., libssl) do not exhibit similar effectiveness in debloating. This issue will be discussed in detail in Section VI-D.

### B. Reduction of Object Files and Functions

*Reduction of Normal Debloating:* Tables II and III summarize the normal debloating results for object files and function symbols across five applications. Our method showed significant reductions, particularly with libc, which was reduced by 66%-74% in object files and approximately 70% in functions. Similarly, libcrypto usage in Nginx was reduced by 95.5%, requiring only 25 out of 556 object files. In contrast, libssl and libreadline demonstrated minimal reduction, with libssl in Vsftpd and libreadline in Sqlite having only 8 and 9 object

files debloated, respectively, and this will be discussed in Section VI-D.

*Reduction of Indepth Debloating:* In the case of in-depth debloating, the reduction of function symbols achieved a 42.2% improvement and the reduction of object files achieved a 20.2% improvement compared to normal debloating. Furthermore, it was evident that the reduction of libcrypto surpasses that of libssl. This will be discussed in Section VI-D.

Overall, we can observe that D-Linker removed plenty of redundant object files, thereby eliminating redundant functions. This indicates that, in most scenarios, some redundant object files are linked into the shared libraries. Therefore, for resource-constrained environments, object-level debloating proved to be an efficient approach.

### C. Reduction of Data Sections

A clear advantage of D-Linker over previous works is that D-Linker can debloat data sections (e.g., .data and .rodata sections) in a shared library, bypassing the accurate analysis of the control flow of data access.

Tables III and IV summarize the number of data symbols and the size of the data section debloated by D-Linker. The data symbols were counted from the symbol table of the shared library; the size of the data section was read from the .data and .rodata sections (in addition to the code and data sections, other sections, such as .rela.dyn, were also debloated, and here, we are only concerned with the debloating of the data part). In normal debloating, nginx achieved the most efficient reduction above 70% in data size and 78% in data symbol.

In the case of in-depth debloating, the in-depth debloating technique achieved an improvement of 20.6% in data symbol reduction and 11.3% in size reduction compared to normal debloating. In summary, the effectiveness in the libc library was excellent both in terms of number (about 70% on average) and size (about 80% on average), while the effectiveness in the libssl (6% on number and 10% on size) and libreadline (almost no debloating) were not noticeable. The reason of this will also be discussed in Section VI-D.

Overall, we can observe that, through object-level debloating, D-Linker effectively debloated the data sections in shared libraries.

### D. Analysis of Object File Dependency

As the analysis from above sections, the effectiveness of D-Linker on the different shared libraries was different. The main reason is the different organization of object file dependence in shared libraries, and this section analyzes the object file dependencies in different libraries to evaluate the unequal effectiveness of debloating is related to the object file dependence in shared libraries.

Tables I-III show that shared libraries, such as libc, had a more excellent debloating effectiveness. In contrast, libraries such as libssl and libreadline were poor, mainly related to the dependencies of object files in shared libraries. The reasons are as follows.

TABLE III  
REDUCED SYMBOLS OF THE LIBRARY

Program	Export func			Local func			Data		
	Baseline	Tailored	Reduction	Baseline	Tailored	Reduction	Baseline	Tailored	Reduction
<b>Nginx</b>									
libc	1670	403	-75.9%	435	245	-43.7%	498	239	-52.0%
libcrypto	3563	320	-91.0%	823	62	-92.5%	1318	97	-92.6%
libpcre	27	9	-66.7%	50	42	-16.0%	51	47	-7.8%
libz	88	32	-63.6%	38	23	-39.5%	66	29	-56.1%
<b>Coreutils</b>									
libc	1670	617	-63.1%	435	232	-46.7%	498	239	-52.0%
<b>Sqlite</b>									
libc	1670	424	-74.6%	435	207	-52.4%	498	198	-60.2%
libz	88	55	-37.5%	38	26	-31.6%	66	38	-42.4%
libreadline	490	482	-1.6%	227	222	-2.2%	700	700	0.0%
libncurses	464	99	-78.7%	192	37	-80.7%	194	166	-14.4%
<b>Openssh</b>									
libc	1670	548	-67.2%	435	261	-40.0%	498	289	-42.0%
libcrypto	3563	2873	-19.4%	823	690	-16.2%	1318	1173	-11.0%
libz	88	50	-43.2%	38	26	-31.6%	66	38	-42.4%
<b>Vsftpd</b>									
libc	1670	388	-76.8%	435	252	-42.1%	498	258	-48.2%
libcrypto	3563	2878	-19.2%	823	673	-18.2%	1318	1133	-14.0%
libssl	483	453	-6.2%	63	51	-19.0%	74	70	-5.4%
<b>Vsftpd(In-depth)</b>									
libc	1670	388	-76.8%	435	252	-42.1%	498	258	-48.2%
libcrypto	3563	1837	-48.4%	823	541	-34.3%	1318	846	-35.8%
libssl	483	395	-18.2%	63	34	-46.0%	74	56	-24.3%
<b>Alpine(In-depth)</b>									
libc	1670	666	-60.1%	435	266	-38.9%	498	296	-40.6%
libcrypto	3563	1837	-48.4%	823	541	-34.3%	1318	846	-35.8%
libssl	483	395	-18.2%	63	34	-46.0%	74	56	-24.3%

TABLE IV  
REDUCED DATA OF THE LIBRARY

Program	Data size		
	Baseline	Tailored	Reduction
<b>Nginx</b>			
libc	218432	34607	-84.2%
libcrypto	188424	41232	-78.1%
libpcre	13872	11480	-17.2%
libz	18426	12248	-33.5%
<b>Coreutils</b>			
libc	218432	166710	-23.7%
<b>Sqlite</b>			
libc	218432	39501	-81.9%
libz	18426	15432	-16.2%
libreadline	28254	28254	0.0%
libncurses	46915	14049	-70.1%
<b>Openssh</b>			
libc	218432	30189	-86.2%
libcrypto	188424	171604	-8.9%
libz	18426	15427	-16.3%
<b>Vsftpd</b>			
libc	218432	23868	-89.1%
libcrypto	188424	158400	-15.9%
libssl	53045	37732	-28.9%
<b>Vsftpd(In-depth)</b>			
libc	218432	23868	-89.1%
libcrypto	188424	133610	-29.1%
libssl	53045	37732	-28.9%
<b>Alpine(In-depth)</b>			
libc	218432	182788	-16.3%
libcrypto	188424	133610	-29.1%
libssl	53045	37732	-28.9%

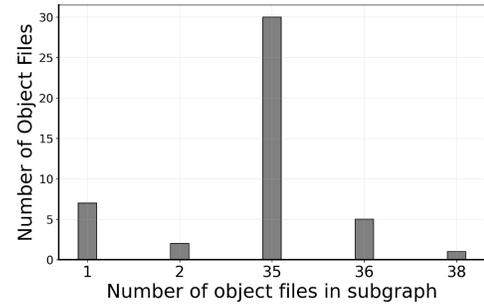


Fig. 3. Distribution of symbols in Libssl and Libc.

with the most object files (1341), has 1670 exported functions, whereas libcrypto has 556 object files but 3563 functions, and libssl has 46 object files with 489 functions. Consequently, libc's sparser distribution of functions results in superior debloating effectiveness. Detailed analysis (Fig. 3) shows that libc's object files uniformly contain fewer symbols (within 5), while libssl has object files with many symbols, such as `ssl_lib.o` with 148 symbols. This dependency on heavily populated object files renders libssl essential to the shared library, leading to poor debloating effectiveness.

*Reference Graph of the Library:* We also analyzed the reference graph of libssl, as shown in Fig. 4, the reference graph between the object files in libssl; we can see that most of the nodes have a large degree of entry and exit, which causes the entire graph to be distributed more aggregated and thus the number of complete subgraphs (smallest subgraph without undefined symbols) is small. As shown in Fig. 5, in libssl, 80% of the object files are in the subgraph with more than

*Number of Symbols in Object Files:* The number of object files and functions in shared libraries is not proportional; libraries with fewer symbols per object file exhibit better debloating effectiveness. For instance, libc, the shared library



TABLE VI  
COMPARING WITH PIROR WORKS

	Piece-wise [6]	ELFtailor [13]	BlankIt [14]	$\mu$ Trimmer [18]	D-Linker (this study)
No Source Code Needed	No	Yes	No	Yes	Yes(need object files)
No Runtime Support	No	Yes	No	Yes	Yes
Debloating Granularity	Function	Function	Function	Code Block	Object file
Load Time Slowdown	20X	0%	10X	0%	0%
Runtime Slowdown	0%	0%	18X	0%	0%
Gadgets Reduce	70-80%	50-60%	90%	50-60%	50-60%
Data Debloating	No	No	No	No	Yes

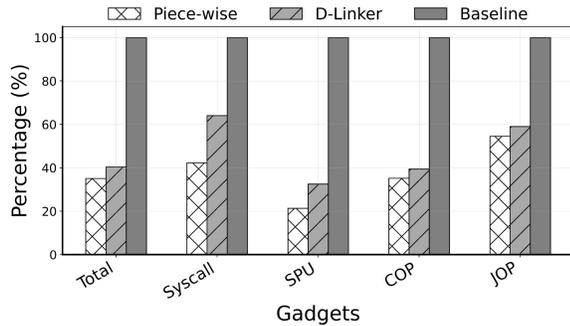


Fig. 7. Reduced gadgets comparison with piecewise, baseline is the original musllibc.

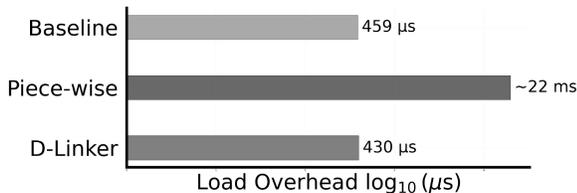


Fig. 8. Load time comparison with piecewise, baseline is the original musllibc.

767 reduction. Consequently, ELFtailor’s advantage in code size  
 768 reduction was not as prominent when compared to D-Linker.  
 769 However, ELFtailor’s finer-grained binary analysis enabled  
 770 it to replace unused code with NOP operations, offering  
 771 advantages above 10% over D-Linker in terms of unused  
 772 code elimination. In terms of memory usage, D-Linker’s  
 773 file size reduction can be effectively reflected in memory  
 774 utilization, leveraging the existing Linux memory management  
 775 mechanisms. Additionally, the debloating of the data section  
 776 translates to significant memory savings in multiprocess  
 777 scenarios, surpassing the capabilities of ELFtailor and other  
 778 related works.

779 *Comparison With Piecewise:* To evaluate the security  
 780 performance of D-Linker, we compared its effectiveness in  
 781 reducing gadgets with that of piecewise [6]. Piecewise opti-  
 782 mizes the compilation process at the source code level during  
 783 linking, thereby minimizing the amount of unused code loaded  
 784 into memory and enhancing security benefits. As illustrated in  
 785 Fig. 7, our analysis of various gadgets showed that although  
 786 piecewise, due to its source-level optimizations, had certain  
 787 advantages above 5% over D-Linker, and the requirement for  
 788 source code and the additional costs incurred during load  
 789 time (as shown in Fig. 8) make D-Linker highly attractive for  
 790 security optimizations in embedded scenarios.

791 *Comparison With Other Works:* Finally, we compared exist-  
 792 ing shared library pruning tools and evaluated their features  
 793 against those of D-Linker. As shown in Table VI, we assessed  
 794 each tool based on several criteria: requirement for source  
 795 code, need for additional runtime mechanism support, debloat-  
 796 ing granularity, introduction of runtime overhead, security  
 797 evaluation, and ability to debloat data sections. Approaches  
 798 we compared included piecewise [6] and ELFtailor [13],  
 799 as previously mentioned, along with BlankIt [14] and  
 800  $\mu$ Trimmer [18]. We observed that source-based optimizations  
 801 by piecewise and BlankIt offer apparent advantages in code  
 802 debloating, reflected in security assessments. However, these  
 803 come at the cost of increased runtime overhead. Additionally,  
 804 ELFtailor, while effective, is less efficient than D-Linker in  
 805 reducing shared library size due to the overhead introduced by  
 806 binary-level rewriting, as previously discussed. Both ELFtailor  
 807 and D-Linker employ dynamic analysis during the debloating  
 808 process, whereas  $\mu$ Trimmer avoids dynamic analysis but is  
 809 limited to the MIPS instruction set. Among all the tools eval-  
 810 uated, only D-Linker effectively debloats data sections. Based  
 811 on the above evaluations, we find that D-Linker’s object-file-  
 812 level debloating offers several advantages: reduced disk and  
 813 memory usage, elimination of the need for a compilation  
 814 process during linking, no introduction of runtime overhead,  
 815 and significant security benefits. These advantages make  
 816 D-Linker more suitable for embedded scenarios compared to  
 817 prior works.

## VII. DISCUSSION

818 In this section, we discuss D-Linker’s effectiveness, limita-  
 819 tions, and future works.

820 *Analysis of D-Linker’s Effectiveness:* The advantage of  
 821 D-Linker in terms of debloating, compared to other works,  
 822 mainly stems from its object file-based debloating granularity.  
 823 This approach ensures that many unnecessary sections beyond  
 824 the code segments, such as relocation information, symbol  
 825 information, and data section, are omitted from the binary. This  
 826 advantage becomes particularly evident in in-depth debloating  
 827 scenarios. However, the number of object files that can be  
 828 debloated is heavily dependent on the interdependencies  
 829 among object files within shared libraries. For instance, in  
 830 musllibc, most functions correspond to individual object files,  
 831 enabling D-Linker to achieve nearly function-level debloating  
 832 granularity, resulting in better outcomes. This distinction  
 833 accounts for the variation in debloating effectiveness across  
 834 different shared libraries and underscores why D-Linker, with  
 835

its object file-based debloating approach, achieves superior reduction in size compared to previous works.

*D-Linker's Limitations:* D-Linker debloats shared library by object-level granularity, which causes more reduction effectiveness than other tools [13], [17], [18] that only debloat the code section. However, D-Linker still has two limitations. First, due to the poor code design or other reasons, there are always software or shared libraries with high coupling between code modules, resulting in a tight dependency between the object files in the shared libraries, and more symbols are contained in the same object file, which can affect the effectiveness of D-Linker. Anyhow, for this case, it is still possible to use other binary-level debloating tools to debloat the shared library after being debloated by D-Linker, which is in principle better than any binary-level debloating tool. Second, as for the in-depth debloating, the tradeoff for its superior debloating effect over static linking is that it can only guarantee effectiveness within the scope of the test cases.

*Future Work:* D-Linker only relinks the object files without making further binary-level changes and does not use the reference information of applications and libraries completely. In future work, further analysis of dependencies and unused code between object files during the debloating process can reduce useless dependencies between objects by wiping out symbols, binary rewriting, etc., which will further improve the debloating effectiveness.

## VIII. CONCLUSION

In this article, we propose a object-level debloating approach to enhance debloating effectiveness while ensuring flexibility of usage. The object-level debloating offers higher flexibility compared to the source-code-level debloating, and is more precise compared to the binary-level debloating, resulting in a better-debloating effectiveness. Based on this approach, we also propose D-Linker, a tool that debloats shared libraries by reducing both code and data sections at object-level without recompilation. D-Linker effectively extract object files from the corresponding static library of shared libraries, thereby eliminating the need of compilation and source code. We have applied D-Linker to debloat shared libraries of vsftpd and our approach achieves an average reduction in size of 27.6% for shared libraries, a maximum reduction of vsftpd in size of 44.9% when specific features are identified. The results indicate that D-Linker improves debloating effectiveness by approximately 30% compared to binary-level shared library debloating. Additionally, in terms of security, it incurs a 5% decrease in code gadgets reduction compared to source-code-level shared library debloating.

## REFERENCES

[1] C. Y. Baldwin and K. B. Clark, *Design Rules: The Power of Modularity*. Cambridge, MA, USA: MIT Press, 2000.

[2] A. N. Habermann, L. Flon, and L. Coopriider, "Modularization and hierarchy in a family of operating systems," *Commun. ACM*, vol. 19, no. 5, pp. 266–272, 1976.

[3] B. Hardung, T. Kölzow, and A. Krüger, "Reuse of software in distributed embedded automotive systems," in *Proc. 4th ACM Int. Conf. Embed. Softw.*, 2004, pp. 203–210.

[4] J. Larus, "Spending moore's dividend," *Commun. ACM*, vol. 52, no. 5, pp. 62–69, 2009.

[5] "Shared library loading." Accessed: Aug. 30, 2022. [Online]. Available: <https://os.educg.net/2022CSCC?op=6>

[6] A. Quach, A. Prakash, and L. Yan, "Debloating software through piecewise compilation and loading," in *Proc. 27th USENIX Security Symp. (USENIX Security)*, 2018, pp. 869–886. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity18/presentation/quach>

[7] H. Lekatsas and W. Wolf, "Code compression for embedded systems," in *Proc. 35th Annu. Design Autom. Conf.*, 1998, pp. 516–521.

[8] "Alpine official image." DockerHub. 2022. Accessed: Aug. 30, 2022 from [Online]. Available: [https://hub.docker.com/\\_/alpine](https://hub.docker.com/_/alpine)

[9] J. Lee, J. Park, and S. Hong, "Memory footprint reduction with quasi-static shared libraries in MMU-less embedded systems," in *Proc. 12th IEEE Real-Time Embed. Technol. Appl. Symp. (RTAS)*, 2006, pp. 24–36.

[10] "readahead(2)—Linux manual page." 2022. Accessed: Aug. 30, 2022. [Online]. Available: <https://www.man7.org/linux/man-pages/man2/readahead.2.html>

[11] R. Roemer, E. Buchanan, H. Shacham, and S. Savage, "Return-oriented programming: Systems, languages, and applications," *ACM Trans. Inf. Syst. Security (TISSEC)*, vol. 15, no. 1, pp. 1–34, 2012.

[12] "BlueKeep vulnerability (CVE-2019-0708): National vulnerability database." NIST. 2019. [Online]. <https://nvd.nist.gov/vuln/detail/CVE-2019-0708>

[13] A. Ziegler, J. Geus, B. Heinloth, T. Hönig, and D. Lohmann, "Honey, I shrunk the ELFs: Lightweight binary tailoring of shared libraries," *ACM Trans. Embed. Comput. Syst.*, vol. 18, no. 5s, pp. 1–23, Oct. 2019. [Online]. Available: <https://doi.org/10.1145/3358222>

[14] C. Porter, G. Mururu, P. Barua, and S. Pande, "Blankit library debloating: Getting what you want instead of cutting what you don't," in *Proc. 41st ACM SIGPLAN Conf. Program. Lang. Design Implement.*, 2020, pp. 164–180.

[15] H. Sharif, M. Abubakar, A. Gehani, and F. Zaffar, "TRIMMER: Application specialization for code debloating," in *Proc. 33rd ACM/IEEE Int. Conf. Autom. Softw. Eng.*, 2018, pp. 329–339.

[16] B. Shteinfeld, "LibFilter: Debloating dynamically-linked libraries through binary recompilation," Undergraduate Honors Thesis, Brown Univ., Providence, RI, USA, 2019.

[17] I. Agadakos, D. Jin, D. Williams-King, V. P. Kemerlis, and G. Portokalidis, "Nibbler: Debloating binary shared libraries," in *Proc. 35th Annu. Comput. Security Appl. Conf.*, 2019, pp. 70–83.

[18] H. Zhang, M. Ren, Y. Lei, and J. Ming, "One size does not fit all: Security hardening of MIPS embedded systems via static binary debloating for shared libraries," in *Proc. 27th ACM Int. Conf. Archit. Support Program. Lang. Oper. Syst.*, 2022, pp. 255–270.

[19] H. Koo, S. Ghavamnia, and M. Polychronakis, "Configuration-driven software debloating," in *Proc. 12th Eur. Workshop Syst. Security*, 2019, pp. 1–6.

[20] I. Grammatech. "Binary reduce." 2022. [Online]. Available: <https://grammatech.github.io/prj/binary-reduce/>

[21] A. Turcotte, E. Arteca, A. Mishra, S. Alimadadi, and F. Tip, "Stubifier: Debloating dynamic server-side JavaScript applications," *Empir. Softw. Eng.*, vol. 27, no. 7, p. 161, 2022.

[22] A. Altinay et al., "BinRec: Dynamic binary lifting and recompilation," in *Proc. 15th Eur. Conf. Comput. Syst.*, 2020, pp. 1–16.

[23] G. J. Duck, X. Gao, and A. Roychoudhury, "Binary rewriting without control flow recovery," in *Proc. 41st ACM SIGPLAN Conf. Program. Lang. Design Implement.*, 2020, pp. 151–163.

[24] X. Meng and W. Liu, "Incremental CFG patching for binary rewriting," in *Proc. 26th ACM Int. Conf. Archit. Support Program. Lang. Oper. Syst.*, 2021, pp. 1020–1033.

[25] "Executable and linkable format." 2022. Accessed: Aug. 8, 2022. [Online]. Available: [https://en.wikipedia.org/wiki/Executable\\_and\\_Linkable\\_Format](https://en.wikipedia.org/wiki/Executable_and_Linkable_Format)

[26] C. Mulliner and M. Neugschwandtner, "Breaking payloads with runtime code stripping and image freezing," presented at Proc. Black Hat USA, 2015.

[27] A. Q. Nguyen. "Capstone: The ultimate disassembler." 2019. Accessed: Jul. 30, 2019. [Online]. Available: <https://www.capstone-engine.org/>

[28] J. Keniston, A. Mavinakayanahalli, P. Panchamukhi, and V. Prasad, "Ptrace, utrace, uprobes: Lightweight, dynamic tracing of user apps," in *Proc. Linux Symp.*, 2007, pp. 215–224.

[29] "nm." 2022. Accessed: Aug. 30, 2022. [Online]. Available: <https://linux.die.net/man/1/nm>

[30] "Musl libc." Accessed: Aug. 30, 2022. [Online]. Available: <https://www.musl.libc.org/>