Received: Added at production | Revised: Added at production | Accepted: Added at production

DOI: xxx/xxxx

ARTICLE TYPE

Practical Dynamic Reconstruction of Control Flow Graphs

Andrei Rimsa[1] | José Nelson Amaral[2] | Fernando Magno Quintão Pereira[3]

[1]Department of Computing, CEFET-MG, Belo Horizonte, Brazil
[2]Department of Computing Science, University of Alberta, Edmonton, Canada
[3]Computer Science Department, UFMG, Belo Horizonte, Brazil

Correspondence
Av. Amazonas, 7675. Belo Horizonte, MG/Brazil. 30.510-000. Email: andrei@cefetmg.br

Summary

The automatic recovery of a program's high-level representation from its binary version is a well-studied problem in programming languages. However, most of the solutions to this problem are based on purely static approaches: techniques such as dataflow analyses or type inference are used to convert the bytes that constitute the executable code back into a control flow graph (CFG). This paper departs from such a modus operandi to show that a dynamic analysis can be effective and useful, both as a standalone technique, and as a way to enhance the precision of static approaches. The experimental results provide evidence that completeness, i.e., the ability to conclude that the entire CFG has been discovered, is achievable on many functions that are part of industry-strong benchmarks. Experiments also indicate that dynamic information greatly enhances the ability of DynInst, a state-of-the-art binary reconstructor, to deal with code stripped of debugging information. These results were obtained with CFGgrind, a new implementation of a dynamic code reconstructor, built on top of Valgrind. When applied to cBench, CFGgrind is 9% faster than callgrind, Valgrind's tool used to track targets of function calls; and 7% faster in Spec Cpu2017. CFGgrind recovers the complete CFG of 40% of all the procedures invoked during the standard execution of programs in Spec Cpu2017, and 37% in cBench. When combined with CFGgrind, DynInst finds 15% more CFGs for cBench, and 7% more CFGs for Spec Cpu2017. Finally, CFGgrind is more than 7 times faster than DCFG, a CFG reconstructor from Intel, and 1.30 times faster than bfTrace, a CFG reconstructor used in research. CFGgrind is also more precise than these two tools, handling operating system signals, shared code in functions, and unaligned instructions; besides supporting multi-threaded programs, exact profiling and incremental refinements.

KEYWORDS:
Control flow graph, Dynamic analysis, Code instrumentation

1 | INTRODUCTION

The Control Flow Graph (CFG) [1, p.525] is a fundamental structure supporting the analysis and optimization of programs. A CFG is a directed graph where the vertices represent *basic blocks*. A basic block is a maximal sequence of instructions without branches, except at the end. Edges in the CFG denote possible program flows. Since its introduction in the 70's, likely due to the work of Frances Allen [2], CFGs have emerged as a mandatory program representation adopted in compilers, virtual machines and

program verifiers. In program analyses based on source code, a CFG is produced either directly from that source code or from some high-level intermediate representation. However, there exists also much interest in recovering the CFG from the program's binary representation, as many researchers have demonstrated throughout the 90's[3,4,5,6]. However, while the construction of a CFG from source has a trivial solution, and is routinely performed by compilers, the reconstruction of a CFG from the binary representation is undecidable. Undecidability is easy to see: indirect branches, plus a simple extension of Rice's Theorem[7], hinder any algorithm from determining with certainty every possible flow in a program.

There are two ways to recover a CFG from a program's binary representation. The first approach, henceforth called *static*, tries to recover the program flow via static analysis of the binary program, i.e., from its `.text` section. This is the technique of choice employed by a number of well-known tools, such as DYNINST[8], BAP[9], JAKSTAB[10], SECONDWRITE[11], IDA PRO[12], GNUOBJDUMP[1], and OLLYDBG[2]. The second approach, henceforth called *dynamic*, seeks to construct a CFG out of instruction traces generated during the execution of a program. The dynamic reconstruction of CFGs is not as wide-spread as its static counterpart. We know one industrial-strength tool that provides such capability: Yount's DCFG[13], a software built on top of Intel's PIN, and released in 2015. Dynamic CFG builders can also be found as part of different research artifacts[14,15], a few of which are publicly available[3].

Static and dynamic approaches yield different results. Whereas the static approach gives a conservative approximation of the program's control flow, possibly containing paths that might never be traversed, the dynamic approach gives an under-approximation of the program flow. Every flow discovered by a dynamic tool is a true path within the execution of the program analyzed. However, the dynamic technique might miss paths that are not exercised by the inputs used in the reconstruction. Such differences lead to distinct applications. Static CFG reconstruction is typically used for security analyses[16] and binary optimization[17,18]. Dynamic reconstruction, in turn, is used to build dynamic slices[19,20], and finds services in any situation where such slices are in need[21], such as malware detection, deobfuscation and profiling. Nevertheless, in spite of three decades of progress in dynamic slicing, the dynamic reconstruction of CFGs is still poorly understood, its benefits are often understated, and its engineering still leaves much room for improvement. Motivated by such observations, this work brings forward the following contributions to the recovery of CFGs from binary code:

Completeness: a new definition to quantify the coverage of CFG reconstruction (Section 2) and an empirical evaluation (Section 5.3) that reveals that a standard execution of the SPEC CPU2017 suite yields complete CFGs for 40% of the invoked functions. For CBENCH this number is similar: 37%.

Precision: a suite of techniques that, once combined, yield more precise CFGs than the state-of-the-art approaches available today. Section 3 explains how our techniques support precise profiling information, deal with overlapping instructions and code shared by different functions, handle signals from the operating system, support multi-threaded programs and the incremental construction of CFGs from multiple inputs.

Efficiency: new algorithms (Section 4) that support faster reconstruction of CFGs than state-of-the-art dynamic reconstructors. Our approach is ∼7x faster than DCFG, a tool built over Intel's PINPLAY, and ∼28% faster than BFTRACE, the reconstructor from Gruber *et al.*[14]. Our efficiency is due to extensive use of caching, as Section 5.2 shows.

Complementarity: the demonstration that static and dynamic analyses can be combined to generate more complete CFGs. Section 5.3 shows that the combination of our technique with DYNINST, a state-of-the-art static CFG builder[8], increases coverage in CBENCH from 42% to 57%, and in SPEC CPU2017 from 39% to 46%.

The relevance of the techniques introduced in this paper are demonstrated in CFGGRIND (https://github.com/rimsa/CFGgrind), a dynamic CFG reconstructor. CFGGRIND is mature enough to be used on every program of SPEC CPU2017. It supports the reconstruction of CFGs for programs that run in parallel. It also admits incremental construction of CFGs, meaning that a partial CFG built during one run of the program can be retrofitted into a new execution with different inputs in order to complement it. Thus, if new paths are traversed, more information is added to the CFG. This feature is specially important for programs that require multiple runs to construct a complete CFG. CFGGRIND can be used in tandem with DYNINST, a static binary analyzer, allowing it to discover the target of dynamic jumps, and to handle difficult code sections that would be missed in programs stripped of symbols and debugging information. Additionally, CFGGRIND provides exact profiling information.

[1] GNUOBJDUMP is a disassembler for GNU Linux. To know more, see https://www.gnu.org/software/binutils/
[2] OLLYDBG is a disassembler for Microsoft Windows. To know more, see http://www.ollydbg.de/.
[3] As an example, tools available at https://docs.angr.io/, and https://github.com/toshipiazza/LLVMCFG provide some limited form of CFG reconstruction.

Contrary to sampling based techniques, it tracks how many times every instruction of the target program was executed, respecting the equity of flows: the number of program flows that enter any basic block equals the number of flows that leave it. With this article, we close two years of implementation effort. A first summary of this effort appeared in a paper that we have previously published at the Brazilian Symposium on Programming Languages[22]. That first publication exists in Portuguese only, and describes the algorithms that CFGGRIND uses to reconstruct CFGs. This new work, in turn, now available in a language more widely accessible, describes, in addition to those algorithms, several experiments that demonstrate how CFGGRIND can be effectively used. Moreover, it explains in details how the techniques that we introduce let CFGGRIND outperform state-of-the-art dynamic CFG reconstructors available today.

2 | PRELIMINARY DEFINITIONS

The definition of a CFG is readily available in any compiler textbook; however, given its central role in this paper, this section revisits it. This formalism might differ from standard definitions because it uses a number of terms that are necessary to explain the CFG reconstruction algorithm in Section 4. The building blocks of a CFG are instructions. In the binary representation of a program, each instruction is bound to an address. Each instruction also has an associated textual representation, e.g. *push %rbp*. An instruction can be formally defined as follows:

Definition 1. An **instruction** is a tuple $I = (@addr, size, type, text)$, where $@addr$ is the address of I in memory; *size* is the space that I occupies, measured in bytes; *type* represents a class to which I belongs; and *text* is the assembly textual representation of I. For the purposes of this paper, instructions are classified according to their effect on the flow of control of the program. Therefore an instruction belongs to one of the following types:

standard: flows to the next instruction;

jump(@target, mode: (direct | indirect)): unconditionally jumps to *@target* address, either directly or indirectly;

branch(@target, @fallthrough, mode: (direct | indirect)): conditionally branches to *@target* or *@fallthrough* address, either directly or indirectly.

call(@target, @fallthrough, mode: (direct | indirect)): invokes the function stored at the *@target* address, either directly or indirectly;

return: transfers control back to caller;

The *standard* instructions flow the execution to the instruction immediately after it. The *jump*, *branch* and *call* instructions can transfer control flow directly — the address is embedded in the instruction itself (e.g.: *jmp @addr*); or indirectly — the address is computed either from registers or memory (e.g.: *jmp %rax*). A NIL value is used as the *@target* address in case of indirect control flows. A *return* instruction transfers the execution back to the caller using the address immediately after the corresponding function call; hence, it behaves like an indirect branch. A *return* instruction is usually used to terminate a function, but it can also be used for irregular control flows, either maliciously or not. The tuple (`@0x400580`*, 2, branch(*`@0x40058c`*, @0x400582, direct),* `'jg 0x40058c'`) is an example of an instruction.

Definition 2. A **group** is an ordered sequence of instructions $S = \{I_1, I_2, \ldots, I_n\}$ containing at least one instruction ($|S| > 0$). The instructions in a group are consecutive in the program ($I_{n+1}.@addr = I_n.@addr + I_n.size$). The first instruction of a group is the *leader*. The last instruction is the *tail*. The *leader* is either the first instruction in a program, the target of a *jump*, *branch* or *call*, or the fall-through instruction of a non-taken *branch*. Instructions of type *jump*, *branch*, *call* and *return* cannot be followed by any other instruction.

Instructions are executed in order unless the program flow reaches an operation that diverts execution. Therefore, groups can be formed according to Definition 2 by tracing the sequential execution of instructions from a *leader* to a *tail*. The target of a *tail* instruction will be the *leader* of a next group to be formed. Thus, chains of groups are created during runtime. The sequence of instructions *{(*`@0x400597`*, 1, standard, 'leaveq'), (*`@0x400598`*, 1, return, 'retq')}* is an example of group, assuming that the program flow is diverted to `@0x400597` at some point during execution.

Definition 3. A **Control Flow Graph (CFG)** is a connected, directed graph $G = (V, E)$, where:

- A node $n \in V$ must be in one of the following categories:

  entry: marks the start of the CFG.

  block(group, calls, signals): is a basic block that contains:

  a *group*, according to Def. 2;

  a map of *calls* that associates the addresses of functions with pairs (CFG, count). The first element in the pair is the CFG of a function, and the second is the number of times that function was invoked by a call instruction in the *group*.

  a map of *signals*, similar to the map of calls, except that keys are signal ids, and the CFG, in the pair (CFG, count) is a signal handler with how many times it was invoked.

  phantom(@addr): is an undiscovered node represented by its address.

  exit: marks the return of control to the caller of this CFG.

  halt: marks the stop of the execution of the program — no further instructions can be executed from this point forward.

- An edge $(n_1, n_2, count) \in E$ connects two nodes, n_1 and n_2 ($n_1, n_2 \in V$), with its execution count for profiling information, $count \in \mathbb{N}$, iff:

  One of the following conditions is true:

  1. The *tail* of n_1 is not an unconditional jump and the *leader* of n_2 immediately follows the tail of n_1 in program order.
  2. The *leader* of n_2 is the target of a branch or jump instruction that is the *tail* of n_1.

  And *count* is:

  1. Zero, iff n_2 is a *phantom* node, or when profiling information is not required.
  2. A positive integer with the exact count of how many times this edge was visited during execution.

- *Phantom*, *exit* and *halt* nodes have no successors. The *entry* node has no predecessors. Thus, given an edge $(n_1, n_2, count) \in E$, $n_1 \notin \{phantom, exit, halt\}$ and $n_2 \notin \{entry\}$.

During the reconstruction of a CFG, the algorithm may process branches whose un-taken target has not been visited thus far. These targets are represented by phantom nodes.

Example 1. Figure 1(a) shows the function FMAP written in C, and Figures 1(b,c) show two snapshots of FMAP's CFG. This function receives two parameters: an integer `x`; and a pointer to a function that returns an integer. It performs an indirect function call if `x` is greater than zero. For this example, consider that function `inc` — not shown in Fig. 1(a) — was called. The filled oval in Figures 1(b,c) are *entry* nodes. The double filled ovals represent *exit* nodes. A double filled square denotes the *halt* node. Note that there is only one *exit* node and/or *halt* node per CFG. Functions with multiple exit points, either that terminate the function or terminate the program, must be connected to their respective *exit* or *halt* nodes. Figure 1(b) contains three basic blocks, each one with a group of instructions. The block at address `@0x40058c` holds that a call to function INC was invoked one time. This target is the value stored in the function pointer `*op`. If FMAP is called with other arguments, more target functions will appear in the calls section of this block. Neither block contains invocations of signal handlers, and thus are not shown. The *phantom* nodes are represented with dashed outlines. The question mark represents possible unknown flows when the last instruction in the block is either a *jump*, *branch* or *call* node with *indirect* mode. Note that dashed edges are used to connect question marks, but they serve merely as an indication of an indirect flow for this block.

Previous works have modelled the entire program as a single CFG[6,23,8]. The boundary of functions can still be recorded in such representation, as long as edges in the CFG are marked as *intraprocedural* or *interprocedural*. This formalism departs from that convention: a CFG, according to Definition 3, represents the instructions of a single function. Formalizing a CFG in this way makes it easier to combine the CFG representation with information extracted from compilers such as GCC and LLVM. In particular, representing each function as a separate CFG facilitates the task of tracking the entry and exit points of procedures. In rare occasions, the binary representation of a program can be built in such a way that a set of instructions can be executed through calls to multiple addresses. Meng and Miller solved this problem by allowing a CFG to have multiple entry points[8].

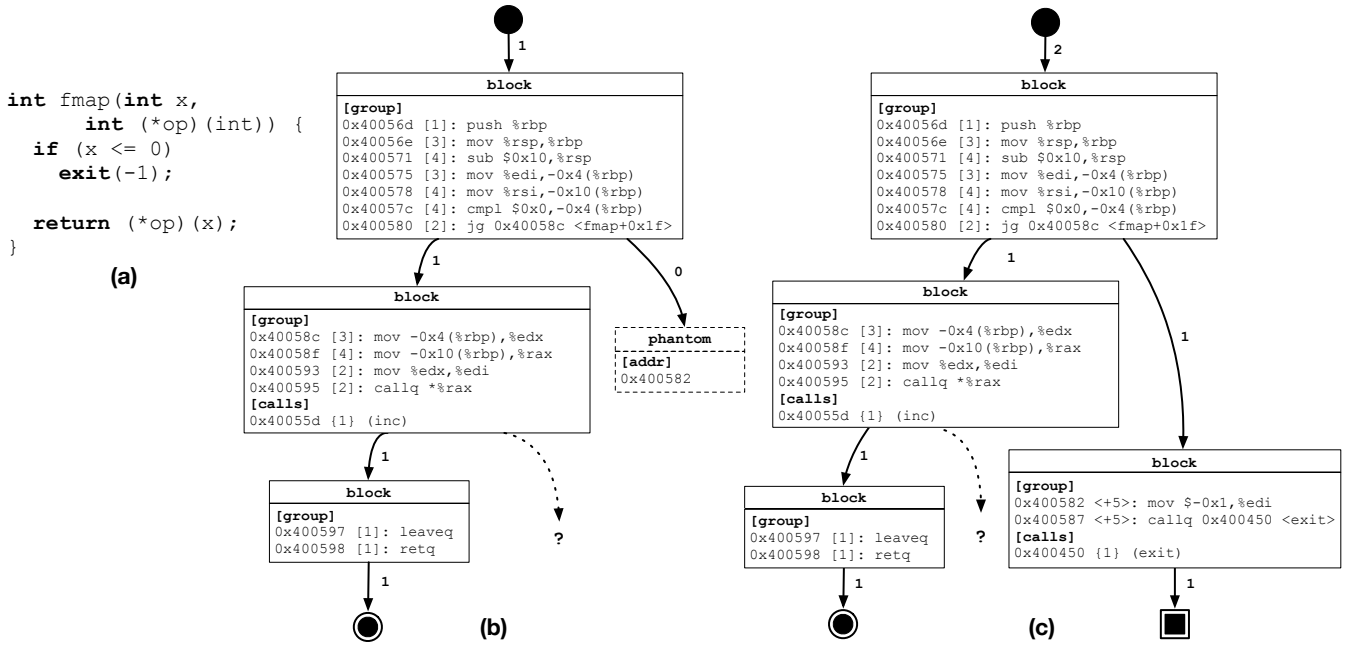

FIGURE 1 (a) Example program. (b) CFG after first call with positive argument for x. (c) Refined CFG after second call with negative argument for x.

We enforce a constraint that a CFG must have a single point of entry. Thus, in CFGGRIND two different CFGs may execute a common subset of instructions. This duplication of information, however, has no penalties in the execution. Example 1 illustrate these concepts.

Completeness

When a CFG contains an indirect jump, an indirect branch, or an indirect call, it is not possible to ensure that all possible execution paths have been discovered. Future executions of the program with different workloads may follow new execution paths. Also, the presence of phantom nodes indicates the existence of paths that have not yet been discovered. The concept of CFG completeness, for the purposed of dynamic reconstruction, can be defined as follows:

Definition 4. Given a control flow graph $G = (V, E)$ (Definition 3), G is said to be **complete** iff, V contains an *entry* and at least an *exit* or a *halt* node; and $\forall n \in V$, the following conditions are true:

1. the successors of n are in V.
2. $n \neq phantom$.
3. if n is a *block* then the *mode* of the *tail* of the *group* of n is *direct*.

A CFG is complete if all its paths are known. Definition 4 uses a more restrictive notion of completeness: even if all indirections are proven to be constrained inside the same CFG, the existence of indirect jumps still classifies this CFG as incomplete.

Example 2. The examples in Figures 1(b-c) present an incomplete CFG because they contains both a *phantom* node and a *block* with an indirect call. Note that edges connecting phantom nodes, although known to exist, are never executed. Thus, they have a count of zero.

This classification is useful in the evaluation of complex code executions. For instance, malware programs may deliberately hide some part of their execution. In order to do so, they rely on constructs such as indirect jumps or calls to avoid exposing the address of the offending code. In such cases, CFGGRIND will mark the CFGs as incomplete. In other words, code that contains indirect control flow will invariably contain either phantom nodes or an indirect marking — the question mark in Figure 1. Further

executions of the same program, with different inputs, might improve coverage; hence, reducing the number of incomplete CFGs. The reconstruction of dynamic CFGs is based on successive refinements. Example 3 shows how re-execution refines CFGs.

Example 3. Figure 1(c) shows the CFG that results from a new activation of the same function, but with different arguments. In this case, the branch at `@0x400580` is not taken and leads to the discovery of the block at address `@0x400582`. In this example, the *phantom* node becomes a *block* node that is connected to the *halt* node.

3 | THE NEED FOR CFG RECONSTRUCTION TOOLS

There are at least two tools that perform the dynamic reconstruction of control flow graphs, namely, DCFG [13] and BFTRACE [14]. DCFG[4] is part of PINPLAY[5] — a framework for deterministically replaying a program execution. PINPLAY is publicly available, albeit closed-source. BFTRACE, in turn, is the first part of a four-staged implementation of dependence analysis [14]. It builds intraprocedural control flow graphs and interprocedural call graphs. Revisiting these technologies, the need for further work in this area stems from two simple observations about state-of-the-art tools. On the one hand, the most precise of these tools, DCFG, incurs a heavy performance slowdown that makes its usage prohibitive in programs with long execution traces. On the other hand, BFTRACE, the faster dynamic analyzer, leaves too much information out from the CFGs that it reconstructs — namely, precise profiling data. This paper shows that it is possible to reconstruct CFGs faster than BFTRACE, and still more exactly than DCFG. Section 5 provides empirical evidence to support this efficiency claim. This section explains why CFGGRIND's CFGs are more complete than similar structures produced by the other tools.

TABLE 1 Qualitative comparison of the different tools considered in this paper.

Feature	CFGGRIND	BFTRACE	DCFG	Section
Completeness	Reported	Absent	Absent	3.1
Program exit	Present	Absent	Absent	3.1
OS Signals	Tracked	Absent	Imperfect	3.1
Edge count	Present	Absent	Present	3.2
Flow equity	Present	Absent	Imperfect	3.2
Incremental analysis	Present	Absent	Absent	3.3
Multi-threading	Handled	Not handled	Handled	3.4
Overlapping instructions	Different	Different	Split	3.5
Shared code in functions	Duplicated	Duplicated	Shared	3.5

Table 1 presents a summary comparison of the three tools and indicates the subsection where each feature is discussed. Beware, however, that these tools are not strictly equivalent: being conceived with different goals, each of them has a distinct representation for CFGs. For instance, BFTRACE is part of a larger system whose purpose is to track dependences between memory regions in order to advise for or against program parallelization. Nonetheless, BFTRACE is a standalone application whose sole purpose is to reconstruct a program's CFGs and call graph. DCFG is also part of a larger system, PINPLAY, which logs program state to allow re-execution, e.g., to support debugging. The code of DCFG is not open; hence, we cannot affirm that its only purpose is to reconstruct CFGs for PINPLAY. Nevertheless, from what we could infer from DCFG's documentation, such seems to be the case.

[4]DCFG: https://software.intel.com/en-us/articles/pintool-dcfg
[5]https://software.intel.com/en-us/articles/program-recordreplay-toolkit

3.1 | On the Precise Representation of CFGs

BFTRACE, DCFG and CFGGRIND adopt different representations for the program's control flow graph. CFGGRIND and BFTRACE associate a CFG for each identified program function, while DCFG provides a single, flattened, CFG for the entire program. However, the CFGs produced by CFGGRIND have a few features that are absent from the CFGs produced by at least one, and sometimes both, of the other tools.

First, CFGGRIND reports the **completeness**, a notion formalized in Definition 4, of a CFG. Neither DCFG nor BFTRACE let users know if a CFG had all its basic blocks visited during the execution of the program. CFGGRIND provides this functionality by augmenting the concept of a CFG with *phantom* nodes and annotations for indirect flows.

Second, the precise recognition of **exit points** is another feature missing in DCFG and BFTRACE. These tools, like CFGGRIND, track paths between different functions along the program's call graph. However, in both DCFG and BFTRACE it is not possible to know if a basic block ends only a function, or terminates the entire program. Our experience using CFGGRIND as a debugger tells us that such differentiation is important to correctly identify the points where no other instructions can be executed.

Third, CFGGRIND tracks signal events that may occur during the program execution. Signals are particularly difficult to handle because they do not originate from specific instructions, e.g., `call` or `jmp`. Some instructions, such as `div`, `mod`, `store` and `load` can produce signals (`SIGSEGV`, `SIGILL`, `SIGFPE`, etc). Signals can come from outside the program, e.g., due to interruptions (`SIGINT`), or can be scheduled to happen, e.g., due to alarms (`SIGALRM`). Example 4 compares the support for this feature in CFGGRIND in contrast with the other tools.

Example 4. Figure 2 shows a sample program that has a signal handler (a) with its respective assembly code (b), and shows how signals are processed by CFGGRIND (c) and DCFG (d). When a signal handler is activated, BFTRACE crashes and is unable to produce the CFGs for such a program. CFGGRIND records the address of the function handler called with its associated signal id at the basic block where the event originated. DCFG creates a special edge marking the function handler as a context switch, but without an associated signal id. Also, DCFG fails to track reliably the correct execution flow after the return of the signal handler. In Figure 2(d) the edge at address `@0x400610` in `BB 28`, is misidentified as a fall-through edge, whereas it should have been marked as a return edge. Furthermore, the target of this edge points to a special `Unknown` node due to an invalid target address calculated at this point. Note that if a signal handler is never activated, none of the three tools are able to find its CFG.

3.2 | On Exact Profiling Information

A profiler provides users with either exact or approximate information. In the latter category we have all the *sampling-based profilers*. In the former, we have *instrumentation* and *emulation* based profilers. CFG reconstructors can be used as a supporting infrastructure to build exact profilers. To fulfill this goal, three features are desirable: **edge count**, **call count** and **signal count**. Edge count gives the number of times each edge in the CFG was traversed by the program flow. Call count provides the number of times each function has been called during the execution of the program. Signal count holds similar information, but for signal handlers instead of functions calls. Both, CFGGRIND and DCFG provide these three features. They are absent in BFTRACE.

Edge counts, when available, should be subject of the **Law of Flows**, which Tarjan[24], among other graph theoreticians, have postulated as: "the sum of incoming flows must equal the sum of outgoing flows on each vertex of a directed graph, except on its start and end nodes." In the context of this work, the count in the incoming edges must add up to the sum of the counts of the outgoing edges for any basic block traversed during program execution. The two exceptions are the program entry point, whose in-degree is zero, and the program exit point, whose out-degree is zero. This principle is true for CFGGRIND; however, it is not entirely true for DCFG.

Example 5. Figure 3(a) shows an example program where the compiler can optimize the invocation of function `add` depending on its calling context. As can be observed by the assembly code produced in Fig 3(b), the call in function `normx` was optimized to use a jump instruction. However, the compiler was unable to use the same strategy for the call in function `twice`. Thus, function `add` is used in two distinct contexts. Similar situation is commonly observed in code in general. For example, LIBGFORTRAN (version 3.0.0) has some data transfer functions, e.g. `transfer_integer` or `transfer_real`, to copy data between different container types. These functions are used externally — using call instructions — , but are also used internally — using jump instructions after tail call optimization. Therefore, CFG reconstruction tools must be able to handle properly such cases. The three CFGs in Fig. 3(c) were produced both by CFGGRIND and by BFTRACE. Each CFG has its own distinctive copy of a shared block — the block with address `@0x400507` is duplicated in the CFGs for `normx` and `add`. DCFG on the other hand uses a

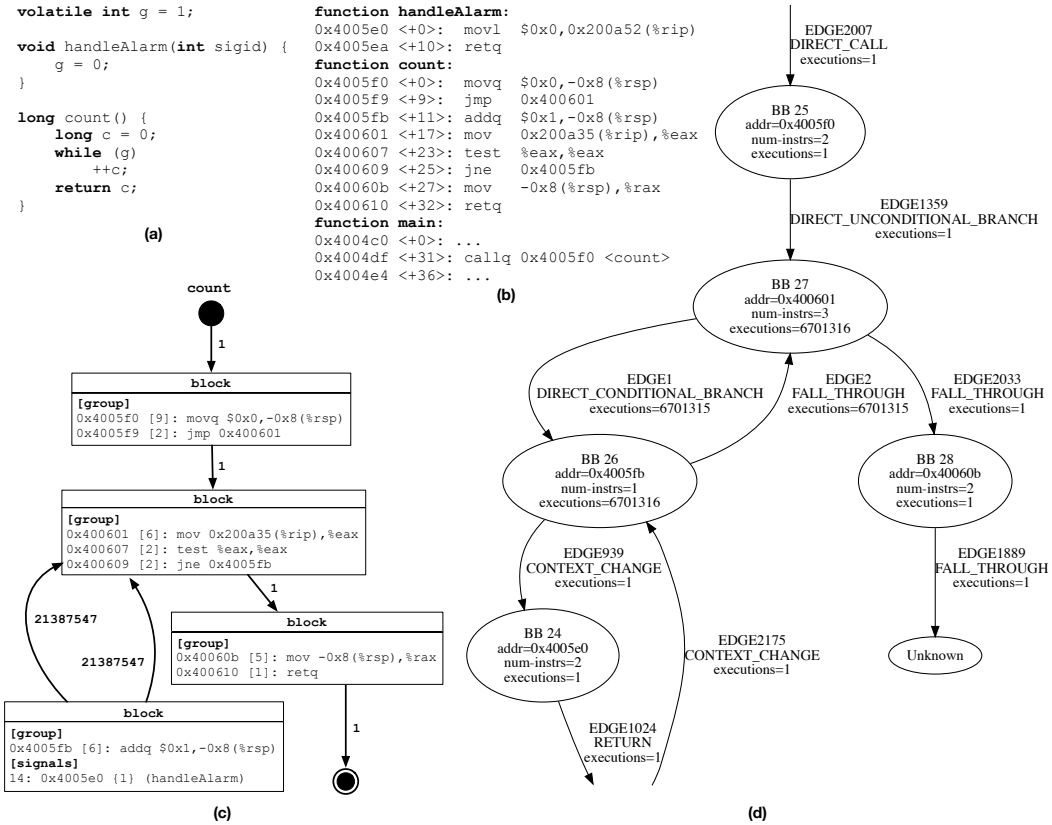

FIGURE 2 (a) Example program with alarm handler. (b) Assembly code for this example. (c) CFG obtained with CFGGRIND. (d) Simplified CFG obtained with DCFG, showing “unknown" node that emerges after signal handling.

single block (`BB 17`) in both contexts. Consequently, it is not possible to determine if the transfer of control happened due to a function call or due to an unconditional branch. In Figure 3(d), there are two return edges going out of `BB 17`, but only one incoming call — the other return edge is due to the jump from `BB 18`.

3.3 | On the Incremental Construction of CFGs

Dynamic analyses require good datasets: the more inputs are available for a program, the more information can be inferred from the program’s behavior. This principle applies to the dynamic reconstruction of CFGs. However, neither DCFG nor BFTRACE support the **incremental construction** of CFGs. In other words, it is not possible to combine events observed in two different executions of a program to build a refined version of its CFGs. CFGGRIND provides this functionality, as Example 1 illustrates. Thus, additional program inputs lead to successive refinements of this program’s CFG; hence, increasing code coverage. Section 5.4 quantifies the benefits of incremental construction in the CBENCH suite. More details on how CFGGRIND supports incremental constructions of CFGs can be found in Section 4.2.

3.4 | On the Execution of Multi-Threaded Programs

A parallel program can span multiple threads during its execution. Both CFGGRIND and DCFG supports tracking the execution of such threads; however, BFTRACE crashes in this scenario. DCFG provides detailed profiling information, where each edge in the control flow graph contains the execution count for each thread separately; CFGGRIND compounds the result of all threads as a total for each edge.

CFGGRIND leverages the serialization performed natively by VALGRIND, where the execution of multi-threaded programs is converted into a single-threaded application by using VALGRIND’s own scheduling policy. CFGGRIND tracks each thread’s

context switch to account for the correct execution flow of programs. More details about the implementation of this feature can be found in Section 4.3. It is unclear how DCFG works internally to support this feature.

3.5 | On Other Assembly Idiosyncrasies

Although low-level assembly code is usually derived from high-level languages via a compilation chain, some aggressive optimizations can dramatically change the structure of the target code. For instance, optimizations might force code sharing between multiple functions. Also, some sections of assembly code can have overlapping instructions. Overlapping happens mostly in hand-crafted code, which either implements some optimization or encodes malware. In this last category, we have examples of return oriented programming attacks[25]. In all these cases, binary code presents idiosyncrasies that a reconstructor must handle.

Example 6. Figure 4 exemplifies the first situation: instructions shared between functions. A snippet of an object dump of mapped symbols available in GLIBC (version 2.17) for two function: *__read* and *__read_nocancel*. The former function is mapped between addresses `@0xeb86a-@0xeb88c`; while the latter is mapped between `@0xeb860-@0xeb8d7`. Since there is an overlap of these two ranges, some instructions are shared by these functions. DCFG approaches this situation using an unique node that is shared across multiple parts of the entire control flow graph. This node can be interpreted as if a section of code has multiple access points. On the other hand, CFGGRIND and BFTRACE build a CFG for each function, which means that each CFG has its own copy of a block that contains these shared instructions. The same approach is employed by CFGGRIND

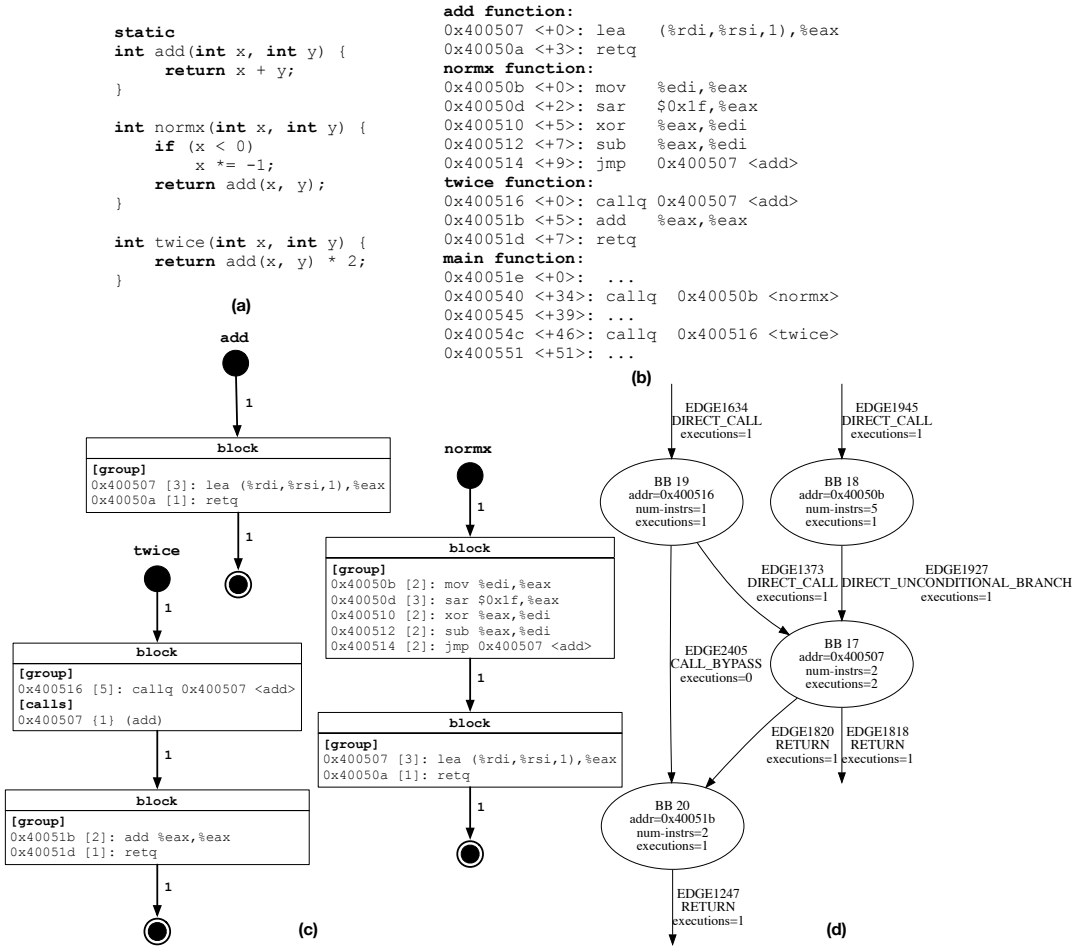

FIGURE 3 (a) Example program with two distinct calls to function `add`. (b) Assembly code with a tail call optimization for this example. (c) CFGs obtained with CFGGRIND for each function. (d) Simplified CFG obtained with DCFG for this program.

to support functions with multiple entry points. Each entry point spawns a different CFG with its own copies of the shared instructions. Thus, every CFG in CFGGRIND has only a single entry point.

FIGURE 4 Objdump snippet for functions *__read* and *__read_nocancel*, from glibc 2.17, that share code. To see that overlapping happens, notice that $[eb86a_{16}, eb86a_{16} + 22_{16}] \cap [eb860_{16}, eb860_{16} + 77_{16}] \neq \emptyset$.

```
000eb86a l     F .text   00000022     __read_nocancel
000eb860  w    F .text   00000077     __read
```

Example 7. Figure 5 exemplifies the second situation: a block of contiguous bytes can be interpreted as different sequences of assembly instructions. Such situation occurs when there is a jump or call to an unaligned target address. Fig. 5(b) shows that CFGGRIND and BFTRACE obtained the same CFG, while Fig. 5(b) shows that DCFG splits the nodes incorrectly, leading to an unrealistic execution flow. A call to address `@0x4004b7` activates `Sequence 1` with two instructions. Its last instruction is a relative jump to the unaligned address `@0x4004b8`. Thus, `Sequence 2` is activated. Of the three instructions of this sequence, the last one is never executed due to the return instruction. CFGGRIND and BFTRACE capture the correct behavior by treating the instructions individually in the blocks. However, DCFG treats the block as a range of addresses, disregarding how the instructions are read inside this range. This modus-operandi leads to the flawed split at node `BB 16`. Although seemly artificial, the unaligned access that this example illustrate is a key component in several real-world ROP-based program exploits, some of which are catalogued in the CVE database[26,27,28].

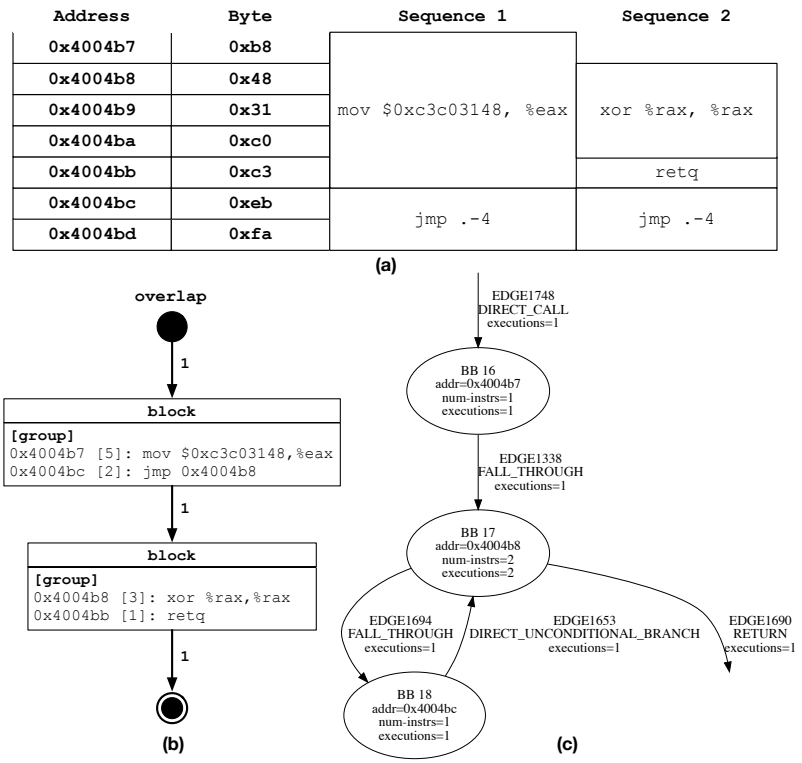

FIGURE 5 (a) Hand-crafter example of two overlapping sequences of assembly instructions. (b) CFG obtained with CFGGRIND and BFTRACE. (d) Simplified CFG obtained with DCFG.

4 | DYNAMIC RECONSTRUCTION OF CFGS

This section uses pseudo-code to explain the dynamic reconstruction of CFGs. CFGGRIND, the tool that prototypes the ideas presented in this paper, is implemented in C, on top of VALGRIND. However, for ease of understanding, the algorithms in this section are presented in a Python-like format. Executable versions of these algorithms can be downloaded from CFGGRIND's repository.

4.1 | The Machine

In the context of this work, a *machine* is any technology, be it based on interpretation, emulation or instrumentation, that produces traces representing the execution of programs. Typical machines include tools such as QEMU, PIN, GDB and VALGRIND. The instructions that appear in a trace are partitioned into groups according to Definition 2. Traces can be processed online, as soon as they are produced by the machine; or offline, as a *post-morten* analysis. The algorithm described in Section 4.2 is agnostic to this processing mode. CFGGRIND, implemented in VALGRIND, uses the online approach. The following example illustrates the notion of a trace.

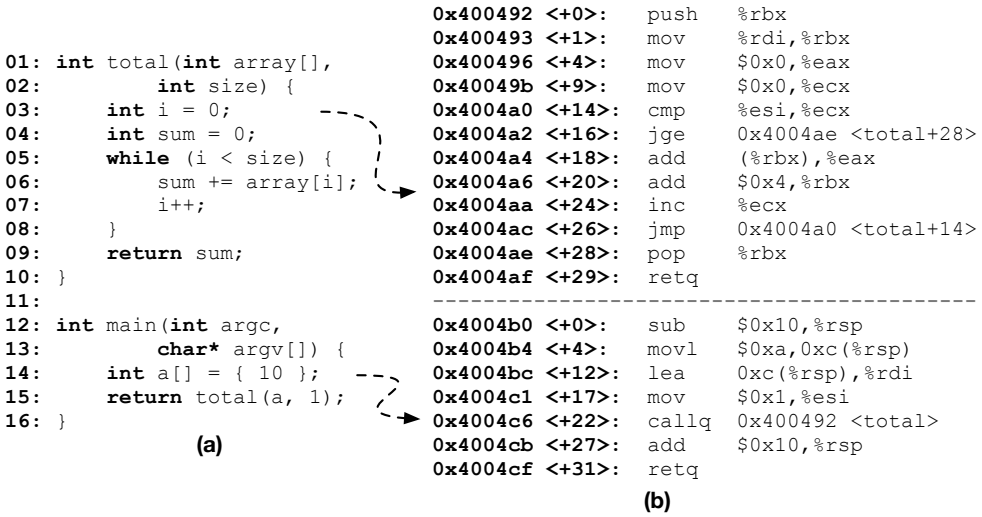

FIGURE 6 (a) Program written in C. (b) static assembly representation of the program.

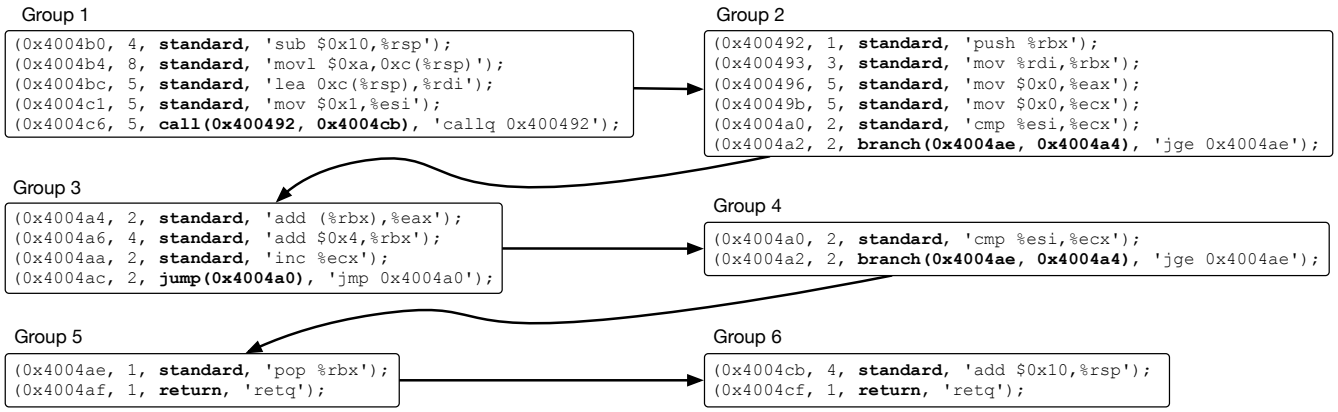

FIGURE 7 Execution trace of the program in Figure 6. Instructions are grouped according to Def. 2. Arrows show order in which groups are processed.

Example 8. Figure 6 shows a program (a) with two functions and its assembly representation (b). The execution of this program in a machine produces a trace formed by those assembly instructions. Such trace represents the paths traversed by the execution of the program. Figure 7 shows the different groups formed by the analysis of this execution trace. The *jump*, *branch* and *call* instructions in this trace are all direct, and thus the *mode* of each instruction is omitted. In this example, the body of the while loop in function `total` (Lines 5-8) executes only once.

4.2 | The Algorithm

Central to the understanding of Algorithms 1-3, is the notion of a *state*, defined as follows:

Definition 5. A **state** is a tuple $S = (current, callstack)$. *Current* is a pair (*cfg*, *working*), where *cfg* is the CFG ($G = (V, E)$, Def. 3) currently being reconstructed, and *working* is one of this CFG's nodes ($working \in V$). The *callstack* is a stack of (*current*, @*ret_addr*) pairs, where @*ret_addr* is a return address. The *callstack*'s *current* pair is similar to the one in the *state*, except *working* must perform a function call (*working.tail.type* is *call*), and the @*ret_addr* is the address of the instruction immediately after this call (*working.tail.type.fallthrough*).

During the reconstruction of CFGs, the algorithms discussed in this section operate on a *state*. The processing of groups, such as those shown in Fig. 7, leads to changes in this state. Thus, Algorithms 1-3 are state-transition functions that map a state-group pair into another state ($state \times group \mapsto state$). When the algorithm processes a *working* node in the *current* CFG, another node becomes the *working* node. When the algorithm processes a function call, the *current* pair is pushed onto the *callstack* and its return address is set. A function return to an address matching a @*ret_addr* in the *callstack* causes the stack to pop elements until this point is reached. The *current* pair associated with this return address is then restored as the *current* pair of the *state*. At initialization, the *current* is set to NIL and the *callstack* is empty ($S = (\text{NIL}, [])$).

Example 9. Figure 8 shows the *state* after each one of the six groups in Figure 7 is processed. In this multi-layer representation, the front layer presents the *current* state, e.g., (*cfg*, *working*). Underneath layers represent the state's *callstack*. The front layer in Figures 8(a) and 8(f) represents the `main` function. The front layer in Figure 8(b-e) corresponds to the `total` function.

The algorithms discussed in this section use a core data structure, the *cfg*, with the following operations:

- *add_node(node)*: adds a new node to the *cfg* if this new node is not already there.
- *add_edge(src, dst, count)*: adds a new edge to the *cfg* from node *src* to node *dst* with *count* as the number of executions. If the edge already exists, increment the previous execution count by the value of *count*.
- *find_node_with_addr(@addr)*: searches for a block node with instruction at @*addr*, or a phantom node at @*addr*; returns NIL if not found.
- *phantom2block(phantom_node, block_node)*: replaces the phantom node with the block node, including moving its predecessors edges to the new node.
- *split(block_node, @addr)*: finds instruction i_j with address @*addr* in the group of the block node such that $i_1 < i_j \leq i_n$, moves the instructions $\{i_1, \ldots, i_{j-1}\}$, and its predecessor to a new block node, and finally connects them with a new edge.

Processing Programs.

Algorithm 1 is the entry point for the process of CFG reconstruction. The algorithm assumes the existence of a global *state*, initialized as (NIL, []), that is readily available during processing. This global *state* can be externally manipulated to support features such as multi-thread programs and signal handlers (Sec. 4.3). The algorithm receives a *machine* and a *mapping* of CFGs indexed by their addresses. The *mapping* can be either empty or pre-populated with CFGs loaded from a previous run. This is they key to support incremental construction of CFGs as described in Section 3.3. By loading previously computed CFGs, the algorithms described in this section can further improve them, as they continue to refine the CFGs as new paths are explored during the execution.

Algorithm 1 follows a program execution to reconstruct the CFGs dynamically by processing each group obtained from the machine individually (Lines 2-12). Once the machine halts, i.e. no more groups are generated, the algorithm finalizes the remaining CFGs by connecting the *working* nodes, of the *state*'s *current* pair or of the *callstack* if present, to the *halt* node (Lines 13-17). Finally, Algorithm 1 returns the updated mapping with all reconstructed CFGs at line 17.

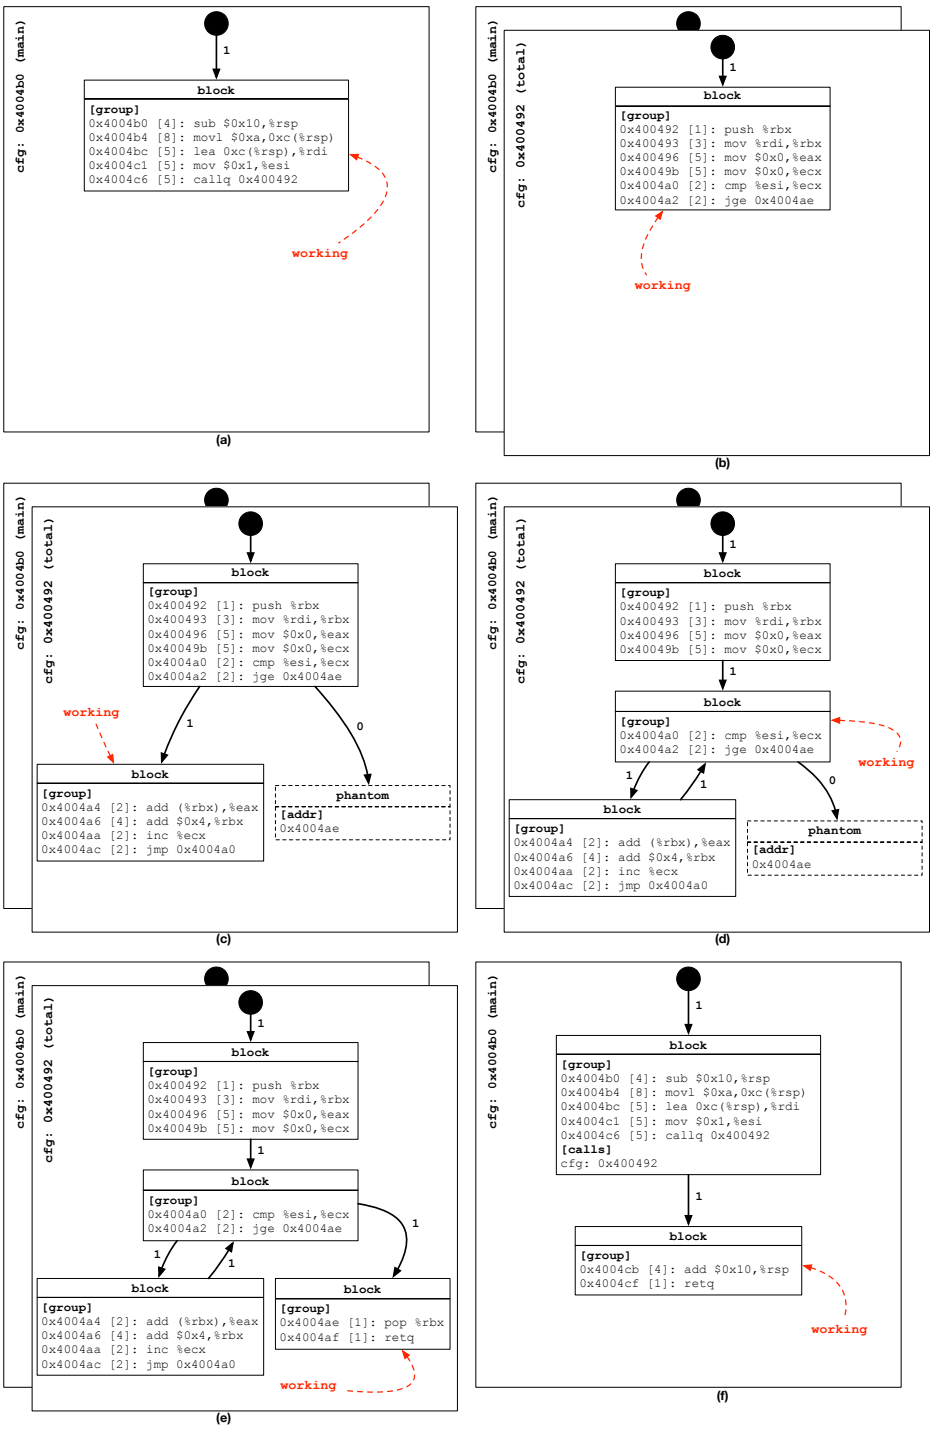

FIGURE 8 State after processing each of the six groups listed by Figure 6.

For each group (Lines 2-12), Algorithm 1 manipulates the state in two phases:

Phase 1 (Lines 4-10): takes an action based on the previous *working* node. In the absence of the *working* node, initializes the first CFG (Lines 4-6). The initial CFG is either fetched from the mapping based on the address of the *group*'s leader instruction if existent, or it is created and set in the mapping (Line 5). Then, the state's *current* pair is configured with this CFG and its entry node (Line 6). Otherwise, ensures that *working* node is a basic block (Line 8) and activates Algorithm 2 (Line 9) passing the type of the *tail* instruction of the *working* node and the address of the next instruction of the group.

Algorithm 1 Process program by handling each group of instructions generated by a machine during execution.

```
    global: state
    input: machine, mapping
    output: mapping
 1: function PROCESS_PROGRAM(machine, mapping)
 2:     for group in machine.RUN() do
 3:         @addr = group.leader.addr
 4:         if not state.current then
 5:             initial = mapping.GET(@addr) if mapping.HAS(@addr) else mapping.PUT(@addr, CFG())
 6:             state.current = (initial, initial.entry)
 7:         else
 8:             assert state.current.working instanceof Block
 9:             mapping = PROCESS_TYPE(mapping, state.current.working.group.tail.type, @addr)
10:         end if
11:         state.current.working = PROCESS_GROUP(state.current.cfg, state.current.working, group)
12:     end for
13:     while state.current do
14:         state.current.cfg.ADD_EDGE(state.current.working, state.current.cfg.halt, 1)
15:         state.current = state.callstack.POP() if not state.callstack.EMPTY() else nil
16:     end while
17:     return mapping
18: end function
```

Phase 2 (Line 11): activates Algorithm 3. This algorithm is responsible for building a new path or following an existing one in the CFG. It may create or split nodes in this process, but it will never transition between CFGs. At the end, Algorithm 3, sets the *working* node to the node which its *tail* is last instruction of the processed group.

Example 10. Each one of the six frames in Figure 8 is a snapshot of the state after each iteration of Algorithm 1. Snapshots are taken immediately after the processing of the group by Algorithm 3 (Line 11).

Group 1: (Figure 8(a)) In phase 1, the CFG for function MAIN is created with its *entry* node set as the *working* node. In phase 2 this group is processed leading to the creation of the block with address `@0x4004b0` with all the instructions of the group. A new edge was created with execution count of 1 from the previous *working* node, i.e. *entry* node, to the current *working* node, i.e. the newly created block.

Group 2: (Figure 8(b)) In phase 1, the pending call of the previous block is processed. The CFG for function TOTAL at address `@0x400492` is created and inserted into *mapping*. This CFG is added to the call map of the *working* node (block `@0x4004b0`). Then, the *state*'s *current* pair is pushed onto the *state*'s *callstack* with the return address `@0x4004cb` — the fall-through of the instruction call. Finally, there is a switch to the new CFG by setting the *state*'s *current* pair with this CFG and its entry node. In phase 2, the second group is processed in a similar fashion as the previous. A block with address `@0x400492` is created, connected from the *entry* with execution count of one, and set as the new *working* node.

Group 3: (Figure 8(c)) In phase 1, the pending branch of the previous block is processed. The algorithm creates a phantom node with address `@0x4004ae` for the target address of this branch. Note that no *phantom* node is created for this branch's fall-through address, since this path will be covered in phase 2 for this group. Thus, in phase 2 the block `@0x4004a4` is created, connected, and set as the *working* node.

Group 4: (Figure 8(d)) In phase 1, there is no action for the jump instruction of the previous block, since the jump target will be handled by this group. In Phase 2, there is a jump to the instruction `@0x4004a0` that is inside block `@0x4004a0`. Therefore, this block must be split in two blocks: block `@0x400492` with four instructions and block `@0x4004a0` with two instructions. Then, a new edge with one execution is created between blocks `@0x4004a4` and `@0x4004a0`. All the instructions of this group are matched against the ones in block `@0x4004a0`; thus no new information is added at this point. Afterwards, this block becomes the *working* node.

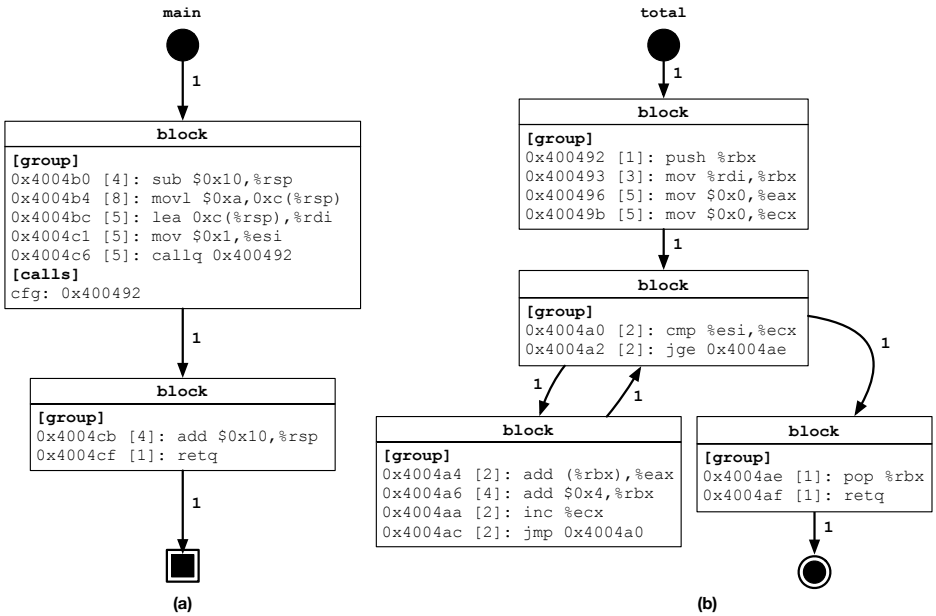

FIGURE 9 The resulting CFGs for functions MAIN (a) and TOTAL (b) in the *mapping* produced by Algorithm 1.

.

Group 5: (Figure 8(e)) In phase 1, the branch of instruction `@0x4004a2` is processed again, but both paths it can follow have already been covered; thus nothing is changed for this CFG. The execution followed the target of the branch, which lead to this group. In phase 2, the leader of this group matches the address of the phantom node at `@0x4004ae`. Thus, the *phantom* node is converted to a *block* node and it is populated with the instructions of this group and the update count of the edge increased by one. Finally, this new block becomes the *working* node.

Group 6: (Figure 8(f)) In phase 1, a function return occurs, because the tail instruction of the *working* node is a return. First, the *working* node is connected to the *exit* node with an edge count of one. Then, the algorithm checks if there is a return address in the *callstack* that matches the address of the leader of this group. In this case, the leader address of this group `@0x4004cb` matches the top of the stack. Thus, the *state*'s *current* pair is restored by popping the top of the stack. At this point, the current *working* node is block `@0x4004b0`, and the *cfg* the CFG of the MAIN function. In phase 2, the block `@0x4004cb` is created, connected, and set as the *working* node.

After processing all the groups, Algorithm 1 connects the state's *working* node to the *halt* node to conclude the execution (Lines 13-16). Then, Algorithm 1 returns the *mapping* containing the functions MAIN and TOTAL that were invoked during the execution of this program (Line 17). The final CFG for both functions can be seen in Figure 9.

Processing the Type of a Group's Tail Instruction.

Algorithm 2, invoked at Line 9 of Algorithm 1, performs an action based on the *type* of the *tail* instruction of the previously processed group. This *tail* instruction is obtained from the last instruction of the *working* node, which is always a basic block. The function PROCESS_TYPE of Algorithm 2 receives a *mapping* of all CFGs discovered so far, the *type* of the *tail* instruction of the previous group, and the target address (*target_addr*) of the *leader* instruction of the next group obtained from the machine. This function returns the updated *mapping*, in case new control flow graphs are discovered. Note that this function may also affect the global *state*.

According to Definition 2 the *type* of the *tail* instruction of a group must be either *jump*, *branch*, *call*, or *return*. If the *type* is an unconditional *jump*, then no special action is required (Lines 2-3). In this case, only one program flow is possible in the CFG and it will be handled when processing the next group. If the *type* is a conditional *branch* then Algorithm 2 models the possible execution flows for this instruction (Lines 4-21). First, it builds a list of the possible target addresses: the branch's target if it is known — in case of a direct branch — , and the branch's fall-through address (Lines 5-8). Then, for each target *@addr* (Line 9) that is not the *target_addr* of the next block, Algorithm 2 either: (1) splits its block, if *@addr* is not in the first instruction (Lines

Algorithm 2 Process the type of the tail instruction of a group.

```
    global: state
    input: mapping, type, @target_addr
    output: mapping
 1: function PROCESS_TYPE(mapping, type, @target_addr)
 2:     if type instanceof Jump then
 3:         # do nothing
 4:     else if type instanceof Branch then
 5:         addrs = [type.fallthrough]
 6:         if type.direct then
 7:             addrs.APPEND(type.target)
 8:         end if
 9:         for @addr in addrs do
10:             if @addr ≠ @target_addr then
11:                 node = state.current.cfg.FIND_NODE_WITH_ADDR(@addr)
12:                 if node then
13:                     if node instanceof Block and node.group.leader.addr ≠ @addr then
14:                         node = state.current.cfg.SPLIT(node, @addr)
15:                     end if
16:                 else
17:                     node = state.current.cfg.ADD_NODE(Phantom(@addr))
18:                 end if
19:                 state.current.cfg.ADD_EDGE(state.current.working, node, 0)
20:             end if
21:         end for
22:     else if type instanceof Call then
23:         called = mapping.GET(@target_addr) if mapping.HAS(@target_addr) else mapping.PUT(@target_addr, CFG())
24:         state.current.working.ADD_CALL(called, 1)
25:         state.callstack.PUSH(state.current, type.fallthrough)
26:         state.current = (called, called.entry)
27:     else if type instanceof Return then
28:         pops = state.callstack.POPS_COUNT(@target_addr)
29:         while pops > 0 do
30:             state.current.cfg.ADD_EDGE(state.current.working, state.current.cfg.exit, 1)
31:             state.current = state.callstack.POP()
32:             pops−−
33:         end while
34:     else
35:         error "Unreachable code"
36:     end if
37:     return mapping
38: end function
```

12-15); or (2) creates a new phantom node, if *@addr* does not belong to a known block (Line 16-18). Regardless of the case, the *working* node is connected to this new node without updating its execution count, since this path has not been traversed yet (Line 19).

Example 11. Figure 8(c) shows that Algorithm 2 created the phantom node `@0x4004ae`. Said node corresponds to the target of the branch `jge` at the end of group `@0x400492` that was not taken.

If the *type* is a *call*, then a different CFG will be visited (Line 22-26). First, the CFG is obtained either from the mapping if it already exists, or a new instance is created otherwise (Line 23). Then, this CFG is added to the call list of the *working* node (Line 24) with the execution count incremented by one. Later, Algorithm 2 pushes the *current* pair with the *cfg* and *working* node onto the state's *callstack* with the expected return address, at the fall-through of this call after it is completed (Line 25). Finally, the *state*'s *current* pair is updated with the called cfg and its entry node (Line 26).

Example 12. Figure 8 shows the transition in the *state* for the CFGs that happens when the function `main` (Fig. 8(a)) makes a call to another function `total` (Fig. 8(b)). At this point, the *working* node points to the entry node of function `total` — situation prior to the Figure 8(b). Also, the called CFG is added to the call list of block `@0x4004b0` of function `main`, as seen in Figure 9(a).

If the *type* is a *return*, then Algorithm 2 restores the *state*'s *current* pair if the target address matches the return address of an entry in the call stack (Lines 27-33). First, the Algorithm 2 calls the CFG auxiliary function POPS_COUNT to scan the state's *callstack*, from top to bottom, searching for an entry whose *@ret_addr* is the same as the *@target_addr*. It returns how many pops, or hops, are necessary to find the matching entry. Then, while the pop count is positive (Line 29), Algorithm 2 adds an edge from *working* to *exit*, or increments that edge's counter by one (Line 30). The *state*'s *current* pair is restored with a pop in the call stack (Line 31). Also, the pop count is decremented by one (Line 32). If there is no entry matching the target address with the return address in the call stack, the return is treated as an unconditional jump. In this case, no further action is performed.

Example 13. Figure 8(e) shows the moment before the return at block `@0x4004ae` is processed by Algorithm 2. The target address `@0x4004cb` matches the top of the call stack, hence a pop is required. The *working* node of Figure 8(e) is connected to the exit node as can be seen in Figure 9(b). Then, the *current* pair is restored to the CFG of MAIN function as in Figure 8(f), but with working node at `@0x4004b0` and before the creation of block `@0x4004eb`.

Algorithm 2 ensures that variable *type* can only be one of: *jump*, *branch*, *call*, or *return* according to Definition 2. Any other type results in an error (Lines 34-35). In the end, Algorithm 2 returns the *mapping*, which might have been updated.

Processing Groups of Instructions.

Algorithm 3 processes each instruction in a group in the order defined by their addresses. When processing instructions, Algorithm 3 either builds a new path in the current CFG, follows an existing path, or does a combination of both. While the same group might pertain to different CFGs[8], a CFG cannot contain only part of a group. Therefore, the only part of the state that matters to Algorithm 3 is the *current* pair — the *cfg* and the *working* node. Algorithm 3 is parameterized by these two arguments, plus the group to be processed. The algorithm follows the instructions of the group, updating the *working* node in the process. It returns a block that has the tail instruction of the group, which Algorithm 1 uses to update the *working* node (Alg. 1-Line 15).

Algorithm 3 processes instructions individually (Line 3). There is an auxiliary variable *curr_instr* to ensure that the first instruction of the group belongs to the successor of the *working* node (Line 2). When *curr_instr* is defined, Algorithm 3 takes no action (Lines 4-5). In this case, the *instr* already exists in the basic block of the *working* node. When *curr_instr* is NIL, there is a switch from one basic block to another. Such event happens in several scenarios, each one implying different actions:

1. The program flow is moving onto a phantom node. Algorithm 3 "resurrects" it, that is to say, turns this phantom node into a basic block (Lines 9-11). A new edge is created between the *working* node and the revived block with the execution count increased by one (Line 16). The new block becomes the current *working* node (Line 17).

2. The program flow is moving onto the middle of a sequence of instructions previously thought to be a single basic block. Algorithm 3 splits this block (Lines 12-13). Then, the same steps of the previous case happens to connect the *working* node to this block and make it the new *working* node (Lines 16-17).

3. The program flow is visiting an instruction that should belong to the *working* node. Algorithm 3 appends the new instruction to the *working* node (Lines 20-23), if: (1) the instruction is not be the leader of the group — group leaders must always be the first instruction of a block; (2) the *working* node has no calls nor signal handlers to other functions, and neither does it have successors nodes.

4. The program flow is visiting a block leader for the first time. Algorithm 3 creates a new node to represent this block of instructions, and sets the *working* pointer to it (Lines 25-28).

Algorithm 3 Process group by handling each instruction individually.

```
    input: cfg, working, group
    output: working
 1: function PROCESS_GROUP(cfg, working, group)
 2:     curr_instr = nil
 3:     for instr in group.INSTRS() do
 4:         if curr_instr then
 5:             assert curr_instr == instr
 6:         else
 7:             node = cfg.FIND_NODE_WITH_ADDR(instr.addr)
 8:             if node then
 9:                 if node instanceof Phantom then
10:                     node = cfg.FIND_NODE_WITH_ADDR(instr.addr)
11:                     node = cfg.PHANTOM2BLOCK(node, Block(Group(instr)))
12:                 else if node.group.leader ≠ instr then
13:                     node = cfg.SPLIT(node, instr.addr)
14:                 end if
15:                 assert node.group.leader == instr
16:                 cfg.ADD_EDGE(working, node, 1)
17:                 working = node
18:             else
19:                 if (instr ≠ group.leader) and (working instanceof Block) and (not working.calls.EMPTY()) and
20:                     (not working.signals.EMPTY()) and (not cfg.SUCCS(working).EMPTY()) then
21:                     assert (working.group.tail.addr + working.group.tail.size) == instr.addr
22:                     working.group.ADD_INSTR(instr)
23:                 else
24:                     node = cfg.ADD_NODE(Block(Group(instr)))
25:                     cfg.ADD_EDGE(working, node, 1)
26:                     working = node
27:                 end if
28:             end if
29:         end if
30:         curr_instr = working.group.NEXT(instr)
31:     end for
32:     return working
33: end function
```

When there is a mismatch between an instruction of the group (*instr*) with the tracker pointer (*curr_instr*), Algorithm 3 takes one of two possible actions. (1) a block must be created (event 4) or modified (events 1 and 2) which guarantees *instr* is the leader. Edges may be added to connect the previous *working* block with this new block. (2) *instr* must be the tail of the *working* block (event 3).

Example 14. Figure 8(e) shows that the phantom node `@0x400492` in Figure 8(d) was replaced with an actual basic block. This happens during the processing of Group 5 in Figure 8, when the branch `jge` is visited. Figure 8(d) shows the splitting of block `@0x400492` into two new blocks: `@0x400492` now with four instructions and `@0x4004a0` with the remaining two instructions. This happens during the processing of Group 4 in Figure 8, because of the jump to `@0x4004a0`.

4.3 | Extensions to the Basic Algorithm

The core algorithms described in Section 4.2 support extensions for performance and precision. Regarding efficiency, Algorithm 1 admits a caching strategy to avoid unnecessary recomputations. Regarding precision, the algorithm supports multi-threaded programs and signal handlers. These extensions are key for CFGGRIND to provide the functionalities described in Section 3.

Caching Strategy.

By Definition 2, once the program flow reaches the leading instruction of a group g, every instruction within g will be executed. As a consequence of this observation, it is not necessary to invoke Algorithm 3 on groups that have already been visited in the same context. In other words, the outcome of function PROCESS_GROUP (Alg. 3), invoked at Line 11 of Algorithm 1, is always the same for a given triple (*cfg*, *working*, *group*). Therefore, as an optimization, the algorithm associates a cache in each node of the *cfg*. When a pair formed by a *working* node and a group is processed for the first time, the algorithm caches the next *working* node. This cache is a table $working_{src} \times group \mapsto (working_{dst}, count)$. The execution *count* is updated in case of cache hits, and flushed in case of cache misses. This optimization is implemented by Algorithm 4, which augments Algorithm 1 with a cache.

Algorithm 4 Algorithm 1 update to use a caching strategy to avoid recomputation of function PROCESS_GROUP.

```
    global: state
    input: machine, mapping
    output: mapping
 1: function PROCESS_PROGRAM(machine, mapping)
 2:     for group in machine.RUN() do
 3:         addr = group.leader.addr
            ...
11:         idx = addr mod CACHE_SIZE
12:         (cached_group, cached_working, cached_count) = state.current.working.cache[idx]
13:         if cached_group == group then
14:             state.current.working.cache[idx] = (cached_group, cached_working, cached_count + 1)
15:             state.current.working = cached_working
16:         else
17:             if cached_count > 0 then
18:                 state.current.cfg.FLUSH_COUNTS(state.current.working, cached_group, cached_working, cached_count)
19:             end if
20:             prev_working = state.current.working
21:             state.current.working = PROCESS_GROUP(state.current.cfg, state.current.working, group)
22:             prev_working.cache[idx] = (group, state.current.working, 0)
23:         end if
24:     end for
        ...
29:     for (addr, cfg) in mapping do
30:         for src in cfg.NODES() do
31:             for (group, dst, count) in src.cache do
32:                 if count > 0 then
33:                     cfg.FLUSH_COUNTS(src, group, dst, count)
34:                 end if
35:             end for
36:         end for
37:     end for
38:     return mapping
39: end function
```

Example 15. The cache avoids work due to repeated loop iterations. The loop in Figure 6(a) iterates once, because its input is a single-element array. However, running this program with a longer array, only the first loop iteration would change the structure of the CFG. The other iterations would just update the execution counters. In this case, the *working* pointers would be moving between the loop condition (block `@0x4004a0`) and loop body (block `@0x4004a4`) without generating new information, except updating its execution count.

The cache avoids the $O(i)$ cost of Algorithm 3, where i is the number of instruction in the group, upon cache hits. The performance evaluation in Section 5 indicates that such situations abound. To support this optimization, the CFG is augmented with the following operation:

- *flush_counts(src, group, dst, count)*: flush execution counts of *group* for *count* times by following the edges starting from *src* node until it reaches the *dst* node.

Every *entry* and *block* node has a cache with n triples like (*group*, *working*, *count*). Entries are indexed by the *leader* address of the group (Alg. 4, Lines 11). The cache size n is configurable at compile-time. The *working* node is updated without invoking Algorithm 3 if the current *group* plus its *working* node gives us a cache hit (Lines 13-15). The cache is updated; hence, increasing by one its execution count.

If we have a cache miss, then a number of actions must be taken, before a new entry is added to the cache. In particular, Algorithm 4 needs to update counters associated with paths stored in the cache. This "flush" is necessary because the cache avoids processing instructions, including updating edge counters. Flushing happens at Lines 17-19 of Algorithm 4. After flushing the cache, Algorithm 4 processes the new group and updates the cache (Lines 20-22). The first entry of a group in the cache is associated with a counter of zero (Line 22), because this first information is recorded directly in the current CFG's edges when PROCESS_GROUP is invoked at Line 21 of Algorithm 4.

Support to Multi-Threaded Programs

The prototype of CFGGRIND, implemented in VALGRIND, has support for multi-threaded programs. VALGRIND natively serializes the execution of such programs into a single-threaded application by implementing its own scheduling policy. CFGGRIND leverages this feature by maintaining a *state* per thread of execution. In Algorithm 5, the *state* is the same global variable used by Algorithms 1-3. The global variable *current_thread* keeps all the information regarding the execution of the active thread, including its ID. The global map *thread_states* holds the states for all the threads indexed by the thread IDs. Initially, each thread state is initialized with an empty state *(*NIL, *[])*, similarly to how the global *state* is configured prior to program execution (Sec. 4.1). A context switch occurs as an event external to the process that runs Algorithms 1-3. Thus, this operation is of no consequence to the inner workings of these algorithms.

Algorithm 5 Context switch from the *current_thread* to another *next_thread*.

```
   global: state, current_thread, thread_states
   input: next_thread
1: procedure SWITCH_CONTEXT(next_thread)
2:     assert current_thread ≠ next_thread
3:     thread_states[current_thread.id] = state
4:     state = thread_states[next_thread.id]
5:     current_thread = next_thread
6: end procedure
```

In a context switch from the active thread (*current_thread*) to a different thread (*next_thread*), Alg. 5 saves the *state* of *current_thread* in the map *thread_states*, indexing it by *current_thread*'s ID (Line 3). Then, the previously saved state of *next_thread* is restored to the global *state* (Line 4). Finally, Alg. 5 updates variable *current_thread* to refer to *next_thread* (Line 5).

Handling Signal Events.

The prototype of CFGGRIND has the ability to precisely track signal events. CFGGRIND relies on VALGRIND's capabilities to identify when an event is raised by the machine and when this event is properly handled by the program. To support signal

events, CFGGRIND employs a strategy similar to the one employed for multi-threaded programs: it manipulates the global *state* externally. Algorithm 6 is responsible to prepare the state when a signal event is raised, whereas Algorithm 7 is responsible to recover the state once the signal was handled by the program. Both Algorithms 6 and 7 hold a global variable for the active thread (*current_thread*), and a global map of state's queue indexed by thread IDs (*signal_states*). Signals occur in the scope of a thread; thus each thread must have its own signal handlers. A queue is required to hold the thread state in case multiple signals are raised simultaneously — they must be treated separately and in sequence.

Algorithm 6 Process a raised signal event.

```
   global: state, current_thread, signal_states
   input: mapping, sigid, @target_addr
1: procedure ENTER_SIGNAL(mapping, sigid, @target_addr)
2:     called = mapping.GET(@target_addr) if mapping.HAS(@target_addr) else mapping.PUT(@target_addr, CFG())
3:     assert state.current.working instanceof Block
4:     state.current.working.ADD_SIGNAL(sigid, called, 1)
5:     signal_states[current_thread.id].ENQUEUE(state)
6:     state = (nil, [])
7: end procedure
```

In Algorithm 6, ENTER_SIGNAL receives the map of CFGs (*mapping*), the ID of the signal raised (*sigid*), and the address of the first instruction to be executed afterwards (*@target_addr*). This address is the entry point of a function that will be called to handle the signal. Algorithm 6 obtains the called CFG for the signal handler based on *@target_addr* (Line 2). Then, it adds this CFG to the list of signal handlers associated with the *working* node with an execution count of one (Lines 3-4). Notice that this *working* node must be of the *block* type. Later, the current *state* is pushed onto the *signal_states* queue for the active thread (Line 5). Finally, the state is initialized as empty to proceed with the execution of Algorithms 1-3.

Algorithm 7 Process a handled signal event.

```
   global: state, current_thread, signal_states
1: procedure LEAVE_SIGNAL( )
2:     while state.current do
3:         state.current.cfg.ADD_EDGE(state.current.working, state.current.cfg.exit, 1)
4:         state.current = state.callstack.POP() if state.callstack else nil
5:     end while
6:     assert not signal_states[current_thread.id].EMPTY()
7:     state = signal_states[current_thread.id].DEQUEUE()
8: end procedure
```

When leaving a signal handler, Algorithm 7 connects the *working* node of the current *state* and all the *working* nodes in its *callstack* if present, to the *exit* node with execution count of one (Line 2-6). This step is similar to the one present in Algorithm 1, Lines 13-16: it is used to ensure consistency of CFGs. Then, the *state* is restored by popping the *signal_states* queue for the *current_thread*.

5 | EVALUATION

The techniques introduced in this paper are integrated into a tool, CFGGRIND, which is publicly available at https://github.com/rimsa/CFGgrind. Although a research artifact, CFGGRIND's implementation is sufficiently solid to enable the exploration of several research questions:

RQ1 How efficient is CFGGRIND, when compared with tools with similar purpose?

RQ2 What is the impact of the cache (Algorithm 4) on the performance of CFGGRIND?

RQ3 What is the ratio between complete and incomplete CFGs in large programs?

RQ4 What is the impact of different input sets in the incremental refinement of CFGs?

RQ5 How much information does CFGGRIND add onto a static CFG reconstructor?

RQ6 What is the time complexity of CFGGRIND in practice?

Benchmarks. This evaluation of CFGGRIND uses two benchmarks, CBENCH (http://ctuning.org/) and SPEC CPU2017 (https://www.spec.org/). CBENCH contains 32 C programs. Each program has 20 available input sets, except `bzip2d` and `bzip2e` with 32 inputs each. The CBENCH programs were modified to compile and run in a 64-bit architecture. SPEC CPU2017 contains 43 programs, written in C, C++ and Fortran. They are organized in 4 categories: integer and floating point, separated into single- and multi-thread versions. Multi-thread programs were configured to execute as single-thread, since VALGRIND serializes the execution. Thus, executing the programs with a single thread in one core for the experiments is sufficient because the performance of CFGGRIND is not affected by the number of threads in an execution (Sec. 3.4). All programs in CBENCH and SPEC CPU2017 are compiled with GCC at the -O2 optimization level. In the comparison with DYNINST (**RQ5** in Section 5.5), the code was stripped from debugging symbols.
Runtime Setup. The current version of CFGGRIND has been implemented in VALGRIND version 3.15.0. Results reported for CBENCH were produced on a 16-core Intel(R) Xeon(R) E5-2630 at 2.40GHz with 16GB of RAM running CentOS 7.5. For SPEC CPU2017, the results were obtained on an 8-core Intel(R) Core(TM) i7-4790 at 3.60GHz with 16GB of RAM running CentOS 7.6. We use two machines to run the experiments in parallel.
Measurement Methodology. Performance numbers for CBENCH are the average of three executions for each program. On average, a run on the entire CBENCH is completed in ~22.5h. The difference between the fastest and slowest of the three runs is less than one minute. Due to this small difference — one minute in 22.5 hours, we shall not report standard deviations in our results. Performance numbers for SPEC CPU2017 were measured only once because of the long run times. Executing a single set of experiments for `intrate` takes ~30.1h, `fprate` ~35.5h, `intspeed` ~40.6h, and `fpspeed` ~295.6h. The experimental evaluation used all the inputs available in both benchmarks for the simulations. To answer RQ1 5.1 and RQ2 5.2, the average is computed using the geometric mean. The variance between each program run times is high: the fastest program in CBENCH executes in ~2s, while the slowest in ~20s; in SPEC CPU2017 the fastest runs in ~4m, while the slowest in ~57m. To answer RQ3 5.3, the average is computed using the arithmetic mean. The total number of complete, incomplete, and unreached CFGs is divided by the total number of CFGs in the benchmark suite.

5.1 | RQ1: Efficiency

Dynamically reconstructing CFGs with CFGGRIND during the execution of a program results in significant overhead. For instance, the execution of the 32 programs of CBENCH with CFGGRIND is ~19 times slower than an equivalent non-instrumented baseline execution. For the 43 programs in SPEC CPU2017, CFGGRIND has a slowdown of ~29 times. This runtime cost is on par with other tools built on top of VALGRIND, whose manual we quote below[29]:

> 'The amount of instrumentation code added varies widely between tools. At one end of the scale, `Memcheck` adds code to check every memory access and every value computed, making it run 10-50 times slower than natively. At the other end of the spectrum, the minimal tool, called `Nulgrind`, adds no instrumentation at all and causes in total "only" about a 4 times slowdown.'

The empirical results in this Section evidence that the instrumentation overhead is also high for other tools that reconstruct CFGs. Figure 10 presents, in logarithmic scale, a comparison of the slowdown for three different tools that reconstruct CFGs: CFGGRIND, BFTRACE[14] and DCFG[13]. The baseline for these comparisons is the original program. Figure 10 also shows the slowdown caused by CALLGRIND, a VALGRIND tool that builds the call graph of a program. Results for the 32 programs available in CBENCH are reported; however, we omit for SPEC CPU2017 because DCFG takes a prohibitively long time to process the larger SPEC CPU2017 suite.

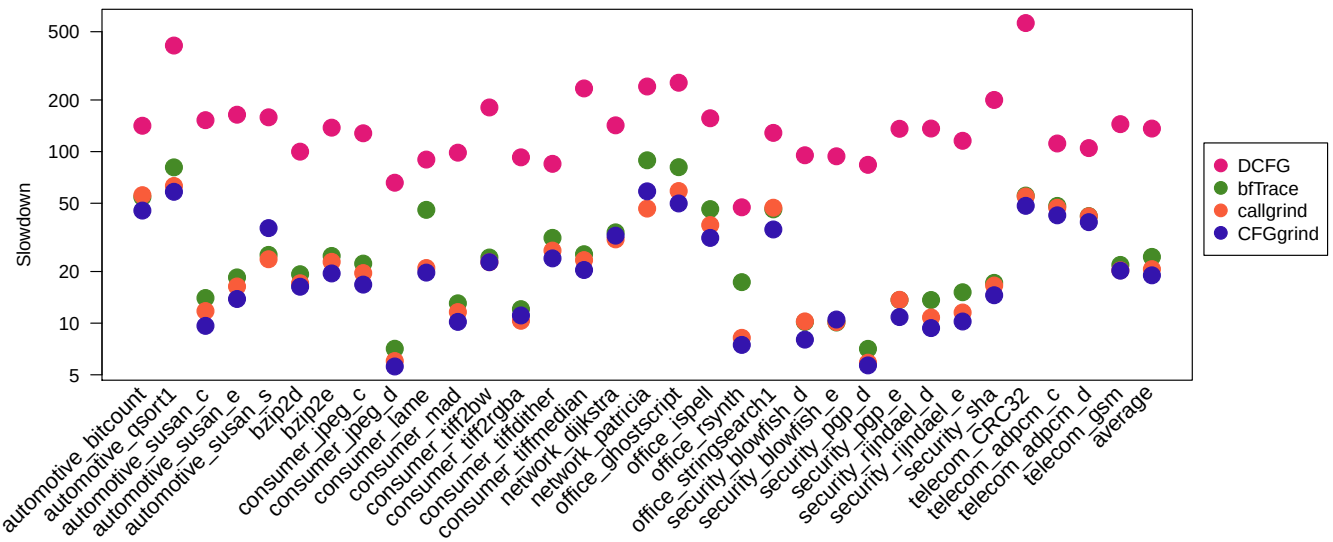

FIGURE 10 Slowdown of different tools that reconstruct CFGs relative to the original program execution without instrumentation for CBENCH.

The CFG-reconstruction tools compared in Fig. 10 are not equivalent (see Section 3). CALLGRIND is not a CFG reconstructor, but it is included in the comparison because it runs on VALGRIND, like CFGGRIND does. Hence, CALLGRIND serves as a performance baseline for readers that are familiar with the VALGRIND's ecosystem. The other three tools in Fig. 10 reconstruct CFGs. However, their outputs, although similar, are not directly comparable because each tool uses its own program representation. For instance, DCFG produces a single CFG for the entire program; CFGGRIND and BFTRACE, in turn, splits it per function.

Figure 10 indicates that CFGGRIND and BFTRACE are much faster than DCFG. DCFG, built onto PINPLAY, saves program state for posterior re-execution — an overhead absent on the other tools. On average, DCFG is ∼7x slower than CFGGRIND. CFGGRIND runs faster than BFTRACE, although by a lower margin: ∼28%. CFGGRIND is also faster than CALLGRIND: ∼9%. Figure 10 also shows the runtime for the original programs. Binaries analyzed by CFGGRIND experiment a slowdown of ∼19x when compared to the original programs — viz., without any emulation. To put these numbers in perspective, DCFG causes a slowdown of ∼136x and BFTRACE, ∼24x. VALGRIND, without any tool, slows CBENCH down by ∼3.6x on average; however, for the sake of readability, we omit this result from Figure 10.

Figure 11 compares the runtime of CFGGRIND for the SPEC CPU2017 suite against the non-instrumented baseline program and other VALGRIND builtin tools. These results indicate that CFGGRIND has a performance on par with CALLGRIND — CFGGRIND is actually ∼7% faster than CALLGRIND. Even though CFGGRIND has a slowdown of ∼29x in relation to the original program, it is only ∼4.5x slower than running VALGRIND without any tool (NULGRIND). Notice that CFGGRIND delivers substantial more information than CALLGRIND. CFGGRIND's CFGs encapsulates the call graph of programs, in addition to all the instructions and paths traversed during the program flow. CFGGRIND is as suitable as other tools in the VALGRIND ecosystem for practical use.

5.2 | RQ2: The Impact of the Cache

The cache implemented by Algorithm 4 is key to boost CFGGRIND's performance. As explained in Section 4, the caching strategy avoids the re-execution of Algorithm 3 for a previously visited pair (*working*, *group*). A cache hit enables the algorithm to move directly to the next *working* node without processing all the instructions of the group. In CFGGRIND, the size of the cache is configurable at compilation time: for each *working* node there are n entries indexed by different group addresses. However, increasing n past a certain value results in diminishing returns. Figure 12 illustrates this trend on the training set for the `intrate` part of the SPEC CPU2017 suite.

Figure 12 makes it clear that the cache is important: the average performance improvement from the introduction of a cache with $n = 2$ is ∼1.6 times. This benefit is substantial in loop-intensive programs, such as `xz`. Setting $n > 2$ produces mixed results

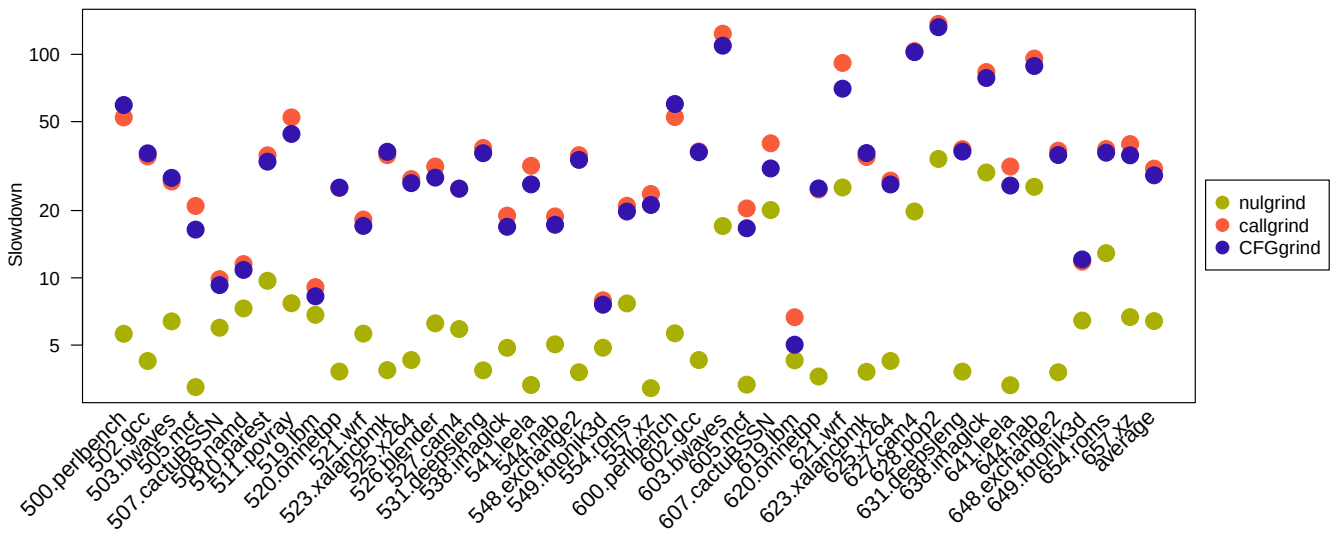

FIGURE 11 Slowdown of builtin tools in VALGRINDand CFGGRIND relative to the original program execution without instrumentation for SPEC CPU2017.

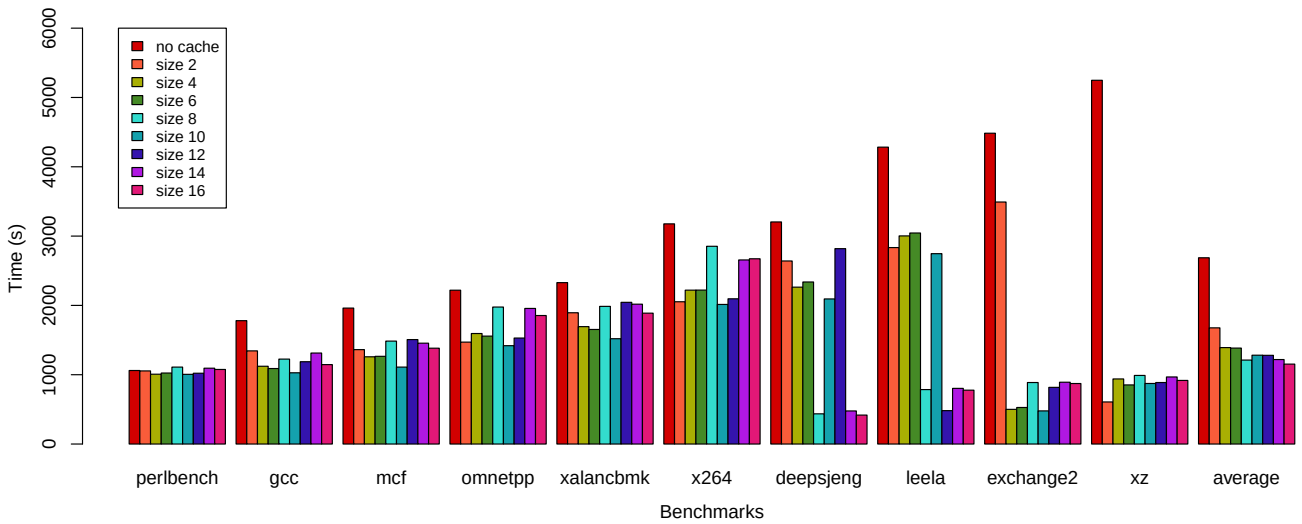

FIGURE 12 The impact of cache size on the runtime of CFGGRIND.

because most of the basic blocks in a program have only one or two successors. Exceptions to this rule are due to indirect jumps, such as those used to implement switch statements. Thus, although the last column of Figure 12 tends to report increasingly better results, improvements past $n = 8$ are too small to be of practical consequence. Larger cache sizes might even provoke slowdowns due to heavier memory usage. Based on these results, the experiment performed to answer **RQ1** (Sec. 5.1) used a fixed cache size of $n = 8$, which provides a good balance between efficiency and memory requirements.

5.3 | RQ3: CFG Completeness

When used to support software testing, CFGGRIND accurately recovers the portions of code traversed by test inputs, with exact profiling information (as explained in Section 3). And, contrary to classic approaches[30,31], it also recovers the CFG of library code. Program coverage through dynamic reconstruction of CFGs can be estimated by answering the following question: “how

many functions had all their instructions executed at least once by a particular input?". These functions are called *complete*, as stated in Definition 4. This section provides an answer to this research question.

Determining Functions of Interest.

Even though CFGGRIND can track the execution of dynamically shared libraries, this study of completeness considers only functions available in the source code of each benchmark. This restriction enables the computation of a ratio of completeness because the total number of functions that can be invoked is available when the source code is accessible. The `.text` section of binary files, compiled with debugging symbols, is used to identify source-code functions. SPEC CPU2017 has 172,268 functions scattered across 43 programs; CBENCH has 7,250 functions in 32 programs.

Invocation Ratio.

The *invocation ratio* of a set of inputs for a benchmark suite is the number of functions that are invoked over the total number of functions in the programs in the benchmark suite. For the SPEC CPU2017, with all the reference inputs, the invocation rate is ∼25%, while for CBENCH, with 20 inputs, the invocation rate is ∼38%.

Completeness Ratio.

Figures 13-14 show, in logarithmic scale, the number of complete, incomplete, and unreached CFGs for SPEC CPU2017 and CBENCH, respectively. Both were executed with the benchmark's reference inputs. The data collection for both figures, from a single run of the entire suite, required 402 hours (almost ∼17 days) for SPEC CPU2017 and 22.5 hours for CBENCH. The *completeness ratio* for a benchmark suite with a given workload is the number of functions for which the entire CFG was discovered divided by the number of functions that were invoked with that workload. For the SPEC CPU2017 suite with the reference inputs, the completeness ratio is ∼40%, and for the CBENCH suite the completeness ratio is ∼37%.

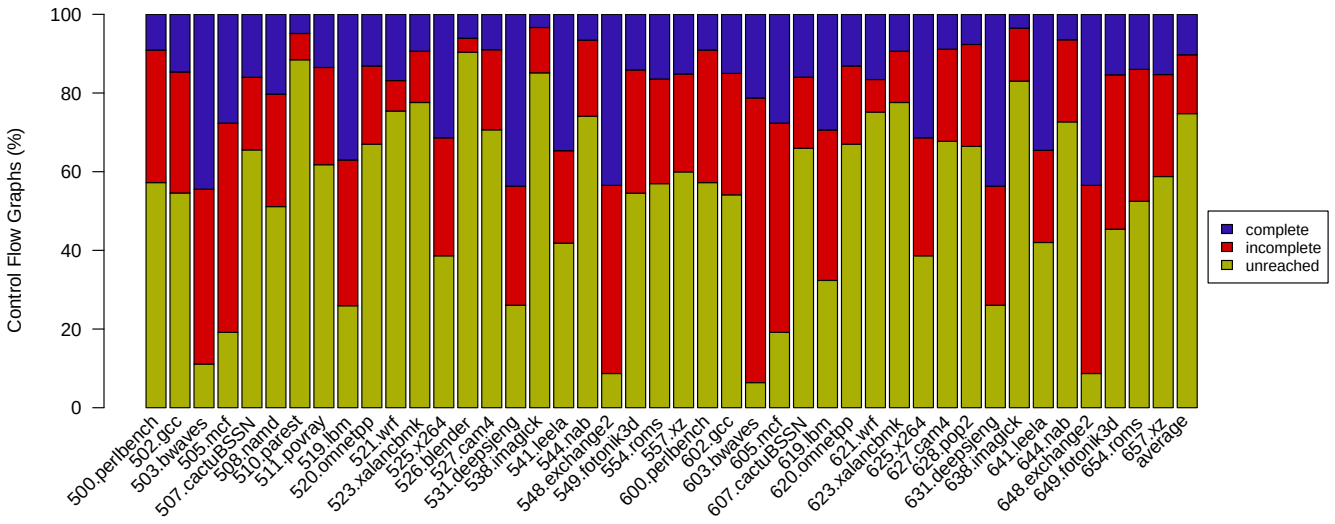

FIGURE 13 Number of complete/incomplete/unreached flow graphs for each of the 43 programs in SPEC CPU2017.

Figure 15 shows the correlation between completeness ratio and (a) number of blocks, or (b) number of instructions per CFG. To improve readability, the graph shows results only for CFGs with up to 100 blocks (a) and up to 1,000 instructions (b). For example, SPEC CPU2017 has 392 CFGs that contain exactly 20 block nodes. Out of those, 50 are complete and 342 are incomplete. The same is true for instructions: out of the 15 CFGs of SPEC CPU2017 with exactly 200 instructions, 4 are complete and 11 are not. Most of the CFGs in programs in the SPEC CPU2017 suite are small; hence, the negative slopes in Figure 15(a-b). A similar behaviour is observed in CBENCH, although not shown in this manuscript. This decreasing rate is much more accentuated for complete CFGs. This trend indicates that, as expected, the probability of finding complete CFGs decreases as the size of the CFG increases.

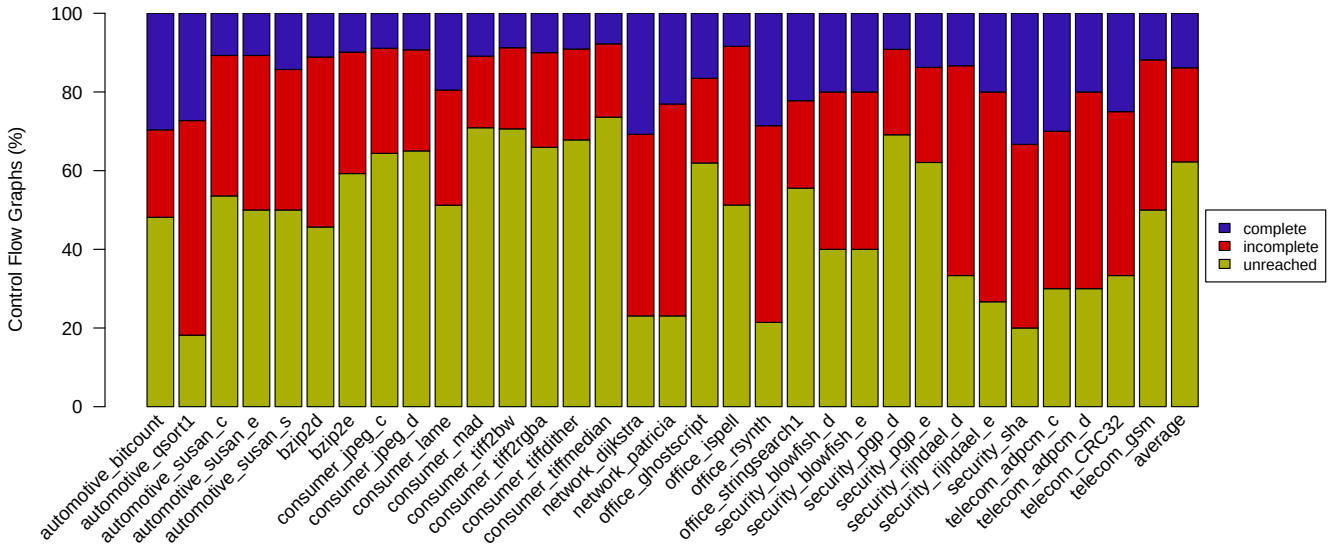

FIGURE 14 Number of complete/incomplete/unreached flow graphs for each of the 32 programs in CBENCH.

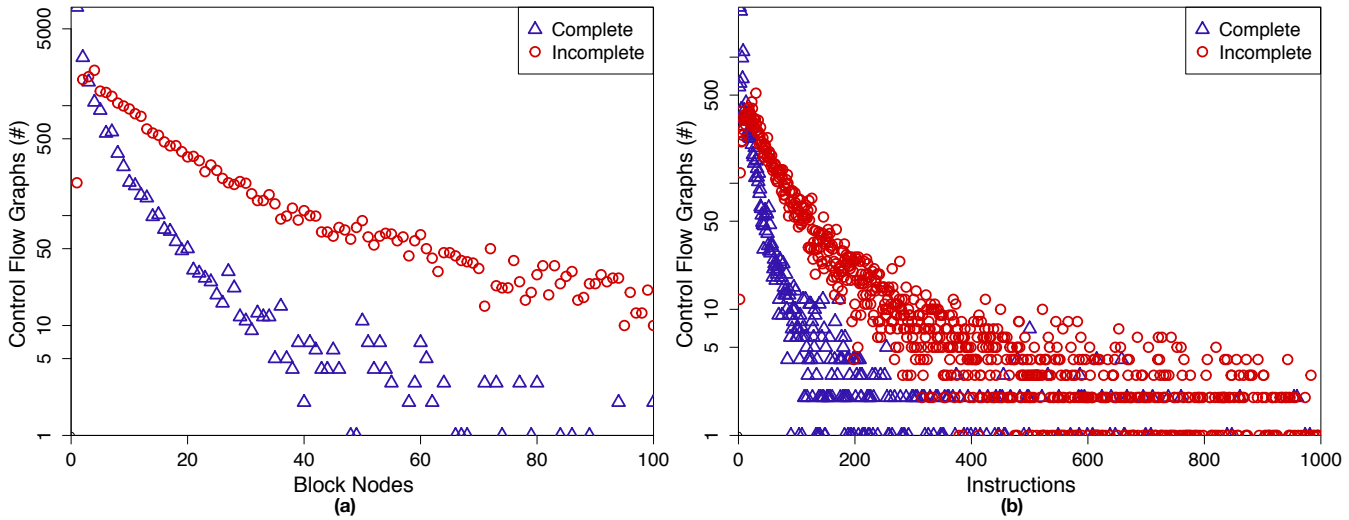

FIGURE 15 Relation between the number of complete/incomplete CFGs (y-axis, ln scale) per number of blocks (a) and instructions (b) for the SPEC CPU2017 benchmark. X-axis shows number of blocks (a) and instructions (b).

5.4 | RQ4: Incremental Construction of CFGs

Multiple invocations of the same function during a single run of a program might lead to more complete CFGs when new paths are explored. To capitalize on this observation, the results produced by a run of CFGGRIND can be forwarded as input to another run. If the new execution flows into unexplored program areas, this information will be added to the CFGs produced. Entire CFGs can be included when new functions are called, and existing CFGs can be expanded when phantom nodes or unmapped areas are visited. The more inputs are given to CFGGRIND, the more complete the reconstruction of the program's control flow. Note that neither BFTRACE nor DCFG support incremental construction of CFGs, as discussed in Section 3.3.

Figure 16 shows how extra inputs contribute to augment the number of visited instructions in CBENCH. This benchmark is well suited for this experiment because each program comes with 20 data sets, except for `bzip2d` and `bzip2e` that comes with 32 inputs each. In this case, the 32 inputs were evenly distributed as 20 inputs in Figure 16. Each tick in the X-axis of Figure 16 shows the number of instructions observed up to the n^{th} execution of a program ($1 \leq n \leq 20$). Following the methodology

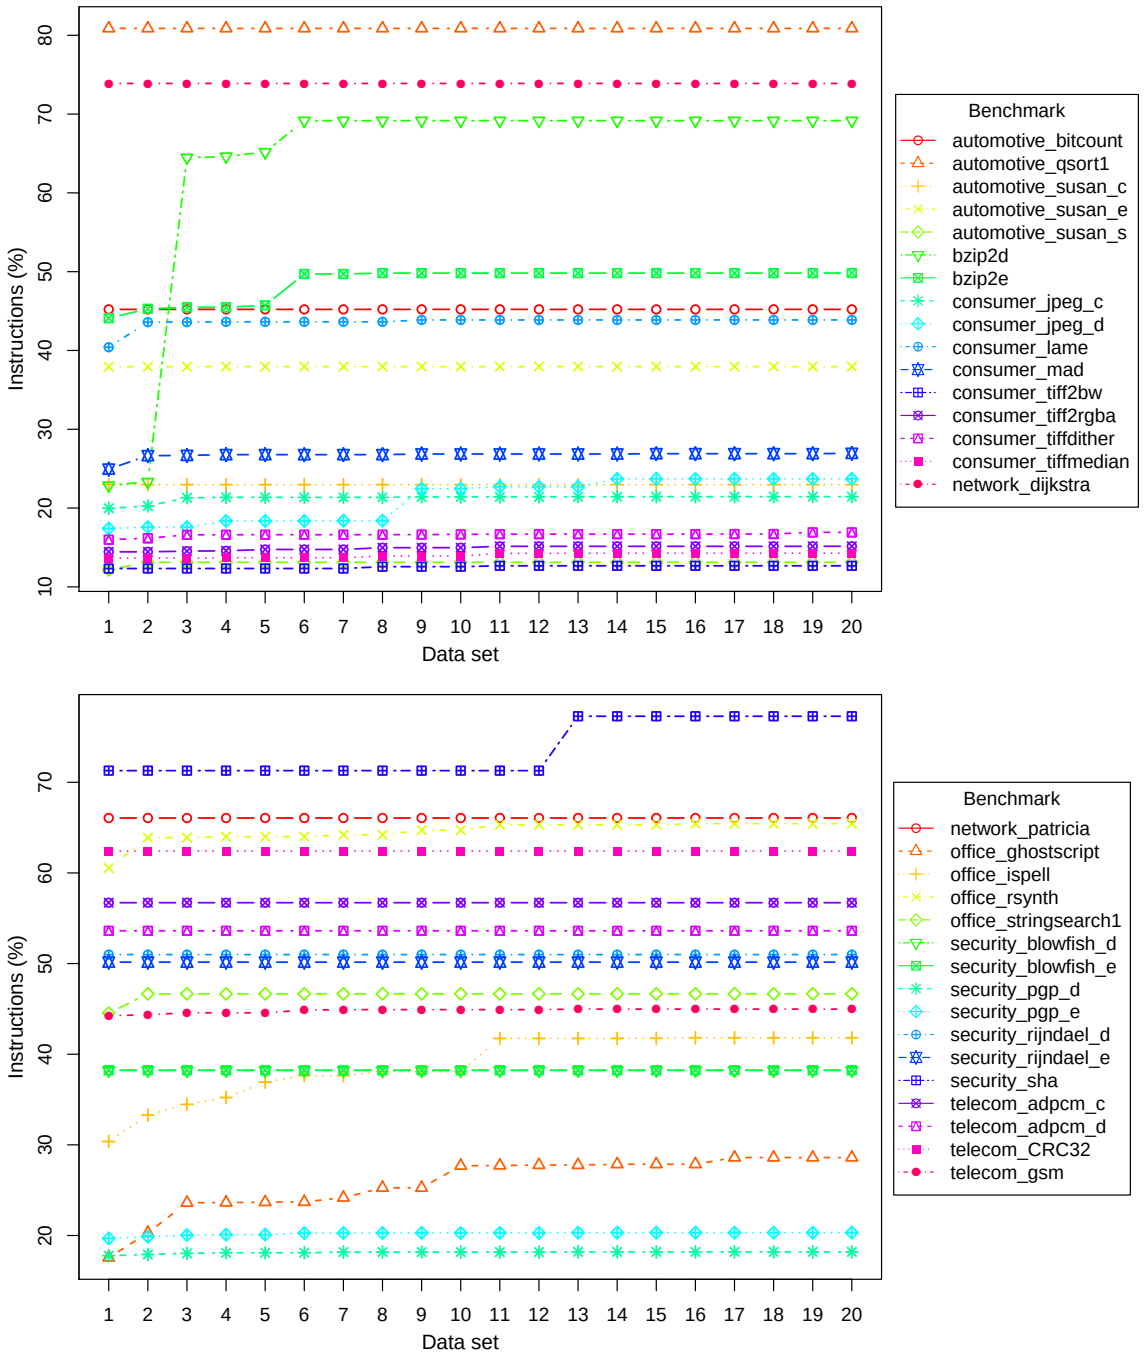

FIGURE 16 Evolution of instruction coverage (y-axis) due to incremental execution of inputs for cBENCH.

used in Section 5.3, library functions are excluded from this analysis. Considering all 32 cBENCH programs, the 19 extra inputs augment the number of instructions visited from 127,016 to 163,750 — an increase of ∼29%. The largest growths were observed in `bzip2d`, 19 new CFGs were added of the 81 available CFGs for the entire program (∼23%), and `office_ghostscript`, 312 new CFGs added onto 3,488 (∼9%). Comparing the first and last executions of all the programs reveals that CFGGRIND was able to identify 378 new CFGs — a growth of 5.21% over a universe of 7,250 CFGs available in the text section of cBENCH. Applying the same principles for instructions, 36,736 new unique instructions were executed — a growth of 6.01% upon 601,345 instructions in cBENCH. This experiment indicates that, at least for cBENCH, extra inputs have a mild effect on code coverage:

they provide new information about the program execution. Although a great extent of each program was already observed in the first execution.

5.5 | RQ5: Combining Static and Dynamic CFG Reconstruction

DYNINST, a state-of-the-art static CFG reconstructor[8], can be used to extend CFGGRIND's coverage, and vice-versa. This section uses these two tools in tandem to analyze CBENCH and SPEC CPU2017. For this experiment, we have compiled the benchmarks without debugging information — the typical way in which production code is distributed. Table 2 shows the result of this comparison for CBENCH. The invocation ratio of CFGGRIND is ∼38%; hence, it identifies 2,738 out of 7,250 possible CFGs. The invocation ratio of DYNINST is 42%; hence, it finds 3,049 CFGs. The two techniques found 1,633 common CFGs, i.e., ∼23% of the total. Similarly, Table 3 shows results for SPEC CPU2017. The invocation ratio of CFGGRIND is ∼25%. This percentage means that it identifies 43,485 out of 172,268 CFGs. The invocation ratio of DYNINST is ∼39%; hence, it finds 66,552 of the CFGs. The two techniques found 30,825 common CFGs in SPEC CPU2017, i.e., ∼18% of the total.

	CFGGRIND (A)	DYNINST (B)	$A \cap B$	$A \setminus B$	$B \setminus A$
CFGs	2,738	3,049	1,633 (59.6%/53.6%)	1,105 (40.4%)	1,416 (46.4%)
Basic blocks	33,316	76,456	23,608 (70.9%/30.9%)	9,708 (29.1%)	52,848 (69.1%)
Edges	52,980	111,732	37,345 (70.5%/33.4%)	15,635 (29.5%)	74,387 (66.6%)
Instructions	163,752	332,189	124,338 (75.9%/37.4%)	39,414 (24.1%)	207,851 (62.6%)
Calls	7,596	18,728	4,120 (54.2%/22.0%)	3,476 (45.8%)	14,608 (78.0%)

TABLE 2 Comparison between CFGGRIND and DYNINST for CBENCH. In the column between the intersection of CFGGRIND and DYNINST, the percentage is given in relation to CFGGRIND and DYNINST, respectively.

	CFGGRIND (A)	DYNINST (B)	$A \cap B$	$A \setminus B$	$B \setminus A$
CFGs	43,485	66,552	30,825 (70.9%/46.3%)	12,660 (29.1%)	35,727 (53.7%)
Basic blocks	939,568	3,466,454	714,309 (76.0%/20.6%)	225,259 (24.0%)	2,752,145 (79.4%)
Edges	1,429,277	5,098,624	1,096,006 (76.7%/21.5%)	333,271 (23.3%)	4,002,618 (78.5%)
Instructions	4,968,718	17,712,186	4,161,470 (83.8%/23.5%)	807,248 (16.2%)	13,550,716 (76.5%)
Calls	302,929	3,160,257	198,561 (65.5%/06.3%)	104,368 (34.5%)	2,961,696 (93.7%)

TABLE 3 Similar to Table 2, but for SPEC CPU2017.

Binaries without debugging information hurt static analyses, whereas dynamic analyses require good program inputs to be effective. Combining these two techniques can improve code coverage.

The combined analyses find 4,154 CFGs for CBENCH — an invocation rate of ∼57%. CFGGRIND finds 1,105 new CFGs that DYNINST was unable to recover statically. In other words, CFGGRIND adds ∼15% more CFGs onto the collection observed by DYNINST. Similarly, the combined analysis for SPEC CPU2017 yields 79,212 CFGs — an invocation rate of ∼46%. Of those, 12,660 (∼7%) were previously unknown to DYNINST. The remaining metrics in Tables 2-3, e.g., blocks, edges, instructions and calls, collide in ways that are hard to quantify. For instance, DYNINST identifies instructions that are never executed, such as those used for padding. Some CFG edges mark impossible paths — they arise due to conservative estimates of indirect branches, for instance. Also, a basic block in one analysis can intersect partially with one or more basic blocks in other analysis. Thus, because there is no one-to-one correspondence between these four metrics in both analyses, the numbers presented must be understood as approximate results.

5.6 | RQ6: Empirical Estimate of Asymptotic Complexity

The reconstruction of CFGs increases the complexity of program execution because of accesses to the cache discussed in Section 4.3. The cache is implemented as a hash-table. In the absence of collisions, the next working node is retrieved in constant time, i.e., with an overhead for this access of $O(1)$. However, collisions might happen. The current implementation of CFGGRIND minimizes collisions via a simple expedient. If occupation of the hash-table reaches 80% of its size, then a new table, twice as large is allocated, and data is copied from the old cache to the new one. If collisions happen, then CFGGRIND uses a list to store multiple entries. We have opted for a list, instead of a balanced tree, for two reasons. First, the list has lower startup cost; hence, it outperforms the balanced trees for small number of elements. Second, the resize-and-copy procedure tends to reduce the number of collisions; thus, in practice, it is unlikely that CFGGRIND's cache will contain a large number of entries with the same hash code. The other components of the algorithms discussed in Section 4.2 contribute only a constant factor to the processing of each instruction. The algorithms follow instructions in the order in which they are executed. The loop in Lines 2-12 of Algorithm 1 processes each group in order, while the loop in Lines 3-31 of Algorithm 3 processes each instructions of the groups sequentially. Algorithm 2 processes the tail instruction of the group, and run in $O(1)$ for jumps, calls and returns. For branches, the algorithm needs to find a successor in a list, but since most blocks have a small number of successors, the search cost is low. Furthermore, operations that add a node or an edge, or replace a phantom node with a block node run in $O(1)$. The operation to find the node with a specific address, or the exact program point where to split a node requires a search in a hash-table; but this operation tends to run in constant time due to lower collisions. Therefore, in practice it is still possible to reconstruct the CFG of programs with a constant time per instruction; or, in other words, with a linear cost in terms of number of executed instructions. Figure 17 supports this observation with empirical data.

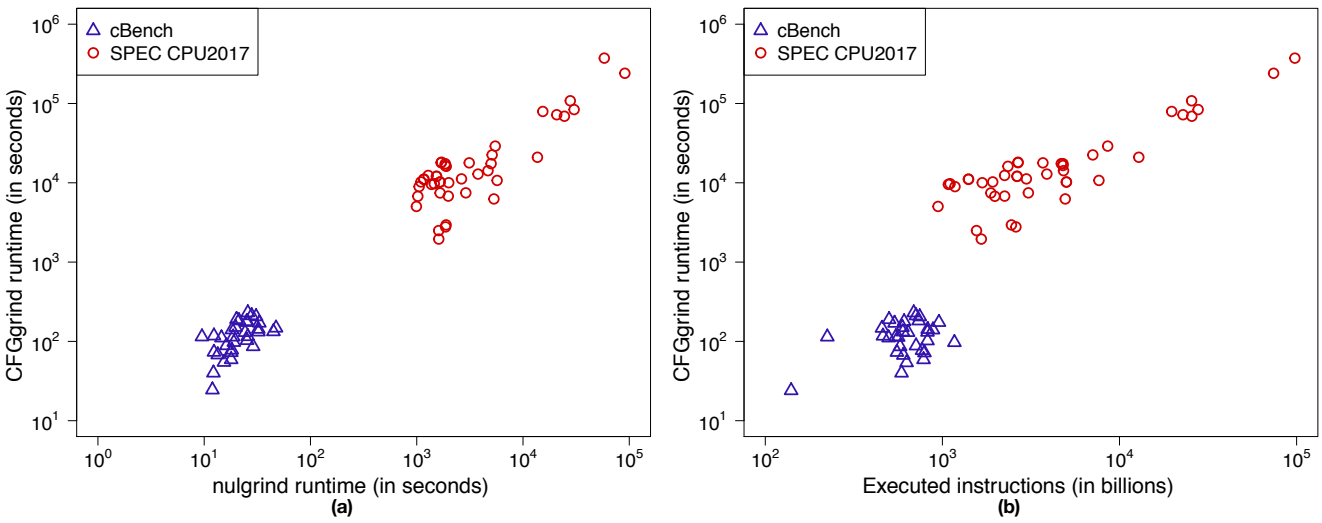

FIGURE 17 Relation between execution time of non-instrumented programs with NULGRIND (emulation only) and programs instrumented with CFGGRIND, for CBENCH and SPEC CPU2017 (a); Relation between the number instructions observed during the execution of CBENCH and SPEC CPU2017 programs, and the running time of these programs, when instrumented with CFGGRIND (b).

Figure 17(a) correlates the running time of programs executed with NULGRIND and the running time of programs instrumented with CFGGRIND. NULGRIND runs VALGRIND on the target program without instrumentation, as a emulation only tool. Visual inspection of the figure indicates strong linear correlation. Indeed, the coefficient of determination between these two running times is 0.905: very strong evidence of linear behavior. The linear relation between the number of instructions that are fetched during program execution, and the running time of CFGGRIND is even stronger. Figure 17(b) supports this statement by presenting such relation for CBENCH and SPEC CPU2017 programs. In this case, the coefficient of determination is 0.990: very close to 1.0, which would be a perfect linear relation.

6 | RELATED WORK

Related research includes the reconstruction of CFGs for the analysis of binary programs; two alternative approaches for dynamic reconstruction of CFGs; the reconstruction of CFGs in dynamic program slicing, which is typically done through program instrumentation; and several approaches for the static reconstruction of CFGs. This section reviews these related topics.

Dynamic Reconstruction of CFGs.

Dynamic analyses of binary programs have been used to detect malware[32], to improve test coverage[33], to de-obfuscate programs[34], to locate out-of-bounds memory accesses[35], and to detect memory dependences[14]. All these uses of dynamic analysis of binaries had to reconstruct the CFG in order to perform the analysis. However, the description of these analyses do not detail the method used to reconstruct the CFG. Therefore, it is difficult to discern the advantages and shortcomings of the CFG reconstruction in each of them. Moreover, none of them provide a publicly available artifact that would allow for an evaluation or comparison with the approach described in this paper. To the best of our knowledge, only three tools focus on dynamic CFG reconstruction: FXE by Xu *et al.*[15], BFTRACE by Gruber *et al.*[14] and PINPLAY by Yount *et al.*[13]. In this paper, we have experimented with the last two of them.

PINPLAY and BFTRACE are the only implementations of dynamic CFG reconstruction that are available for public scrutiny. The experimental results presented in Section 5 indicate that the techniques described in this paper improve on both tools, in terms of efficiency and completeness. Indeed, many of the design decisions in the development of CFGGRIND were motivated by the possibility to use it to augment the precision of DYNINST, a static CFG reconstructor. Integration with static analyzers is not a driving force behind neither PINPLAY's implementation nor BFTRACE's; hence, such possibility is not discussed in the papers that introduce those tools.

FXE combines static and dynamic analysis[15]. Like BFTRACE, FXE interprets a program using QEMU, whereas PINPLAY uses PIN, and CFGGRIND uses VALGRIND. However, instead of simply interpreting each instruction with the state produced by the normal execution of the program, FXE tries to *force* the execution of each branch that it finds while building the program's CFG; hence, a CFG produced by FXE does not correspond to a dynamic slice of a program's execution. In other words, upon finding a phantom node, FXE saves the current state at that program point, and marks it as active. While there are active branches, FXE backtracks, and re-evaluates the branch condition, forcing the visit of the phantom block. Although elegant, this approach has a much higher runtime complexity. Therefore, to keep reconstruction practical, FXE foregoes the analysis of library code, which is a serious limitation for its practical use. According to Xu *et al.*[15]:

> "When FXE detects a function call pointing to external code, it forces the execution to immediately return to the call site and continue along the fall-through."[15]

Dynamic Program Slicing.

Much of the literature on the dynamic reconstruction of CFGs was influenced by the notion of *Dynamic Program Slicing*. This concept was introduced by Korel and Laski[20]. Yet, the formulation of Agrawal and Horgan[36], introduced five years later, seems to be the most standard today. If P is a program, $I \in P$ is an instruction of P and ι is an input of P, then the dynamic slice S is a subset of P's instructions that, when executed, always causes the interpretation of I as in P. Dynamic program slicing has been the focus of much research, and remains a trendy topic even today[37,38].

A survey of the literature on Dynamic Slicing reveals that most work on the area relies on code instrumentation. In contrast, CFGGRIND, PINPLAYand BFTRACE rely on program emulation. Code instrumentation has a key advantage: it simplifies the task of linking runtime events with source code. On the other hand, it has a major disadvantage: it requires the availability of the source code; hence, it is unable to handle library code.

Static Reconstruction of CFGs.

Most papers about the analysis of binary code deal with the static reconstruction of control flow graphs. Seminal work on binary code analysis, such as Cifuentes'[3], Gao's[39] and Balakrishnan's[40], used static reconstruction of CFG. More recent techniques to reconstruct CFGs also use static reconstruction. For example, the binary optimizers that appeared in 2019, such as BOLT[17] and Janus[18]. As discussed in Section 1, the static methodology has advantages and disadvantages over its dynamic counterpart. This paper presents a dynamic CFG reconstruction technique that improved the precision of DYNINST, a static CFG reconstructor, created by Meng *et al.*[8]. To the best of our knowledge, DYNINST is the most precise static CFG reconstructor to date.

7 | CONCLUSION

This paper provided evidence that the dynamic reconstruction of CFGs from the execution of a program can result in more precise CFGs than the ones obtained solely via static analyses. However, to correctly reconstruct CFGs in this fashion, it was necessary to revisit the definition of CFGs to account for phantom nodes and signals. New algorithms had to be engineered into an efficient and robust tool. This tool, CFGGRIND (https://github.com/rimsa/CFGgrind), was used to analyze several large programs. The experimental evaluation determined that CFGGRIND outperforms, in terms of precision and efficiency, two other tools that support dynamic CFG reconstruction: DCFG and BFTRACE. CFGGRIND also improves the precision of DYNINST, a state-of-the-art static binary analyzer by augmenting it with the ability to handle binaries stripped of debugging information. The experimental results also evidenced that typical data sets distributed with benchmarks already let a dynamic reconstructor completely recover a substantial part of all the active functions in large programs. Although intrinsically dependent on program inputs, this complete recovery has been observed in a large number of programs, including SPEC CPU2017 and CBENCH.

Future work that stems from this paper includes the use of CFGGRIND in different scenarios. First, CFGGRIND's ability to track non-aligned and overlapping instructions in the binary representation of a program gives can be useful to reconstruct return-oriented programming attacks[25], even when they are built via Checkoway *el al*[41]'s approach based on indirect jumps. Second, CFGGRIND's exact profiler is likely to give binary optimizers, such as BOLT[17], more information to improve the instruction layout of programs. The performance improvements that can be derived from this extra information remains to be evaluated. Finally, CFGGRIND's dynamic approach can also be useful for the recovery of the control flow of programs obfuscated with control flow flattening[42], the nemesis of static deobfuscation. How much information can be recovered via CFGGRIND when it is used to analyze obfuscated programs is an open question.

References

1. Aho AV, Lam MS, Sethi R, Ullman JD. *Compilers: Principles, Techniques, and Tools (2nd Edition)*. Boston, Massachusetts, USA: Addison Wesley . 2006.
2. Allen FE. Control flow analysis. *SIGPLAN Not.* 1970; 5: 1–19.
3. Cifuentes C, Gough KJ. Decompilation of Binary Programs. *Softw. Pract. Exper.* 1995; 25(7): 811–829.
4. Schwarz B, Debray S, Andrews G. Disassembly of Executable Code Revisited. In: IEEE Computer Society; 2002; Washington, DC, USA: 45–.
5. Sites RL, Chernoff A, Kirk MB, Marks MP, Robinson SG. Binary Translation. *Commun. ACM* 1993; 36(2): 69–81.
6. Theiling H. Extracting safe and precise control flow from binaries. In: ACM; 2000; New York, NY, USA: 23-30.
7. Rice HG. Classes of recursively enumerable sets and their decision problems. *Transactions of the American Mathematical Society* 1953; 74(1): 358–366.
8. Meng X, Miller BP. Binary Code is Not Easy. In: ACM; 2016; New York, NY, USA: 24–35.
9. Brumley D, Jager I, Avgerinos T, Schwartz EJ. BAP: A Binary Analysis Platform. In: Springer-Verlag; 2011; Berlin, Heidelberg: 463–469.
10. Kinder J, Veith H. Jakstab: A Static Analysis Platform for Binaries. In: Springer-Verlag; 2008; Berlin, Heidelberg: 423–427.
11. Smithson M, Elwazeer K, Anand K, Kotha A, Barua R. Static binary rewriting without supplemental information: Overcoming the tradeoff between coverage and correctness. In: IEEE; 2013; Koblenz, Germany: 52–61.
12. Eagle C. *The IDA Pro Book: The Unofficial Guide to the World's Most Popular Disassembler*. San Francisco, CA, USA: No Starch Press . 2011.
13. Yount C, Patil H, Islam MS, Srikanth A. Graph-matching-based simulation-region selection for multiple binaries. In: IEEE Computer Society; 2015; Washington, DC, USA: 52–61.

14. Gruber F, Selva M, Sampaio D, et al. Data-flow/Dependence Profiling for Structured Transformations. In: ACM; 2019; New York, NY, USA: 173–185

15. Xu L, Sun F, Su Z. Constructing Precise Control Flow Graphs from Binaries. tech. rep., University of California, Davis; Davis, CA, USA: 2009.

16. Song D, Brumley D, Yin H, et al. BitBlaze: A New Approach to Computer Security via Binary Analysis. In: Springer-Verlag; 2008; Berlin, Heidelberg: 1–25.

17. Panchenko M, Auler R, Nell B, Ottoni G. BOLT: A Practical Binary Optimizer for Data Centers and Beyond. In: IEEE Press; 2019; Piscataway, NJ, USA: 2–14.

18. Zhou R, Jones TM. Janus: Statically-driven and Profile-guided Automatic Dynamic Binary Parallelisation. In: IEEE Press; 2019; Piscataway, NJ, USA: 15–25.

19. Agrawal H, Horgan JR. Dynamic Program Slicing. In: ACM; 1990; New York, NY, USA: 246–256.

20. Korel B, Laski J. Dynamic Program Slicing. *Inf. Process. Lett.* 1988; 29(3): 155–163.

21. Tip F. A Survey of Program Slicing Techniques. tech. rep., IBM T. J. Watson Research Center; Amsterdam, The Netherlands, The Netherlands: 1994.

22. Rimsa A, Amaral JN, Quintão FM. Efficient and Precise Dynamic Construction of Control Flow Graphs. In: Sociedade Brasileira de Computação. Association for Computing Machinery; 2019; New York, NY, USA: 19–26

23. Bernat AR, Miller BP. Structured Binary Editing with a CFG Transformation Algebra. In: ACM; 2012; New York, NY, USA: 9-18.

24. Tarjan RE. Testing Graph Connectivity. In: ACM; 1974; New York, NY, USA: 185–193

25. Shacham H. The geometry of innocent flesh on the bone: return-into-libc without function calls (on the x86). In: ACM; 2007: 552–561.

26. Gorelik M. [CRITICAL ALERT] CVE-2018-4990 ACROBAT READER DC DOUBLE-FREE VULNERABILITY.; 2018. http://blog.morphisec.com/critical-alert-cve-2018-4990-acrobat-reader-dc-double-free-vulnerability.

27. Seebug M. Tenda AC15 Router - Unauthenticated Remote Code Execution(CVE-2018-5767).; 2018. https://vulners.com/seebug/SSV:97161.

28. Alvarez-Perez D. In depth analysis of malware exploiting CVE-2017-11826.; 2017. https://www.gradiant.org/noticia/analysis-malware-cve-2017/.

29. Seward J. The Valgrind Manual.; 2019. valgrind.org/docs/manual.

30. Fraser G, Arcuri A. EvoSuite: Automatic Test Suite Generation for Object-oriented Software. In: ACM; 2011; New York, NY, USA: 416–419

31. Lemos OAL, Ferrari FC, Masiero PC, Lopes CV. Testing Aspect-oriented Programming Pointcut Descriptors. In: ACM; 2006; New York, NY, USA: 33–38

32. Moser A, Kruegel C, Kirda E. Exploring Multiple Execution Paths for Malware Analysis. In: IEEE Computer Society; 2007; Washington, DC, USA: 231–245

33. Godefroid P. Micro Execution. In: ACM; 2014; New York, NY, USA: 539–549

34. Zhen L. Control Flow Graph Based Attacks: In the Context of Flattened Programs. Master's thesis. KTH. Stockholm, Sweden: 2014.

35. Kimball WB. *A Formal Approach to Vulnerability Discovery in Binary Programs*. PhD thesis. Air Force Institute of Technology, Wright-Patterson AFB, OH, USA; 2013.

36. Agrawal R, Imielinski T, Swami AN. Mining Association Rules between Sets of Items in Large Databases. In: ACM; 1993; New York, NY, USA: 207-216.

37. Hu Y, Zhang Y, Li J, Wang H, Li B, Gu D. BinMatch: A Semantics-based Hybrid Approach on Binary Code Clone Analysis. In: IEEE; 2018; Madrid, Spain: 104-114.

38. Lin Y, Sun J, Tran L, Bai G, Wang H, Dong J. Break the Dead End of Dynamic Slicing: Localizing Data and Control Omission Bug. In: ACM; 2018; New York, NY, USA: 509–519.

39. Gao D, Reiter MK, Song D. BinHunt: Automatically Finding Semantic Differences in Binary Programs. In: Springer-Verlag; 2008; Berlin, Heidelberg: 238–255

40. Balakrishnan G, Reps TW. Analyzing Memory Accesses in x86 Executables. In: Springer; 2004; Berlin, Germany: 5–23

41. Checkoway S, Davi L, Dmitrienko A, Sadeghi AR, Shacham H, Winandy M. Return-Oriented Programming without Returns. In: ACM; 2010: 1-14.

42. Blazy S, Trieu A. Formal Verification of Control-Flow Graph Flattening. In: ACM; 2016; New York, NY, USA: 176–187

AUTHOR BIOGRAPHY

Andrei Rimsa, a Ph.D. student at the Federal University of Minas Gerais under supervision of professor Fernando Pereira. He worked as a web developer and as an embedded software engineer. His research interests include programming languages, compiler technologies, static and dynamic code analysis, binary analysis and computer security. He worked on the PHC compiler—a php compiler, on the LLVM compiler and on the binary instrumentation framework VALGRIND. He is an assistant professor at the Computer Department of Centro Federal de Educação Tecnológica de Minas Gerais (CEFET-MG) since 2014. At CEFET-MG he teaches programming languages and automata theory.

José Nelson Amaral, a Computing Science professor at the University of Alberta, Ph.D. from The University of Texas at Austin in 2004, has published in optimizing compilers and high-performance computing. Scientific community service includes general chair for the 23rd International Conference on Parallel Architectures and Compilation Techniques in 2014, for the International Conference on Performance Engineering in 2020, and for the International Conference on Parallel Processing in 2020. Accolades include ACM Distinguished Engineer, IBM Faculty Fellow, IBM Faculty Awards, IBM CAS "Team of the Year", awards for excellence in teaching, and the GSA Award for Excellence in Graduate Student Supervision.

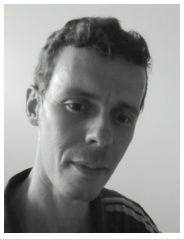

Fernando M. Q. Pereira, Ph.D. at the University of California, Los Angeles, in 2008. Since November of 2009 he is an associate professor at the Department of Computer Science of the Federal University of Minas Gerais. He does research in compilers, and is interested in the design and implementation of static analyses and code optimizations. At UFMG he teaches programming languages and compilation technology. Some of his projects are today part of open-source software, such as LLVM and Mozilla Firefox. Fernando's research is supported by public research agencies, such as INRIA, FAPEMIG, CAPES and CNPq, and by private enterprises, such as Intel, LGE, Google and Maxtrack.

How to cite this article: A. Rimsa, J. Amaral, and F. Pereira (2020), Practical Dynamic Reconstruction of Control Flow Graphs, *SPE.*, *2020;00:0–0.*