

IntFlow: Improving the Accuracy of Arithmetic Error Detection Using Information Flow Tracking

Anonymous Submission

ABSTRACT

Integer overflow and underflow, signedness conversion, and other types of arithmetic errors in C/C++ programs are among the most common software flaws that result in exploitable vulnerabilities. Despite significant advances in automating the detection of arithmetic errors, existing tools have not seen widespread adoption mainly due to their increased number of false positives. Developers rely on wrap-around counters, bit shifts, and other language constructs for performance optimizations and code compactness, but those same constructs, along with incorrect assumptions and conditions of undefined behavior, are often the main cause of severe vulnerabilities. Accurate differentiation between legitimate and erroneous uses of arithmetic language intricacies thus remains an open problem.

As a step towards addressing this issue, we present *IntFlow*, an accurate compile-time arithmetic error detection tool which combines information flow tracking and static code analysis. By associating sources of untrusted input with the identified arithmetic errors, IntFlow differentiates between undefined but developer-intended arithmetic operations, and potentially exploitable arithmetic bugs. IntFlow examines a broad set of integer errors, covering almost all cases of C/C++ undefined behaviors, and achieves high error detection coverage. We evaluated IntFlow using the SPEC benchmarks and a series of real-world applications, and measured its effectiveness in detecting arithmetic error vulnerabilities and reducing false positives. IntFlow successfully detected *all* real-world vulnerabilities for the tested applications and achieved a reduction of 89% in false positives over standalone static code instrumentation.

Keywords

Static analysis, information flow tracking, integer errors

1. INTRODUCTION

When developing programs in the C and C++ languages, programming practices that involve undefined arithmetic

operations often constitute a well-established status quo. Language specifications and compiler checks are tolerant towards such errors, not only to compensate for developers' imprecise understanding of the C/C++ type systems, but, more importantly, to allow for code optimizations that greatly increase the performance of critical code segments. This freedom comes at a cost: arithmetic operations constitute a major source of errors, often leading to serious security breaches when an erroneous value directly or indirectly affects sensitive system calls or memory operations.

The root of the problem is fundamentally bound to the differences between the physical and machine representations of numbers: although integer and floating point numbers are infinite, their machine representations are restricted by their respective type-specific characteristics (e.g., signedness and bit-length). Furthermore, not all arithmetic operations are well-defined by the language standards, in order to allow for a number of compiler optimizations. For instance, the standard does not specify what the value of a signed integer that has reached its upper/lower limits (and consequently overflows/underflows) should be.

As it is not trivial to determine whether a particular arithmetic operation resulting in undefined behavior is benign or not, bugs due to integer errors are prevalent. Integer errors are listed among the 25 most dangerous software bugs [3], and are the root cause of various vulnerabilities such as buffer overflows [23] and memory disclosures [2]. During the past years, numerous attempts have been made towards their automatic detection and prevention. Such efforts include static [21] and dynamic [7] analysis solutions, tools based on symbolic execution and dynamic test generation [8, 17], as well as compiler extensions [23] that resolve ambiguities at compilation time.

Unfortunately, despite the numerous suggested solutions, there is no generic tool that provides effective and complete detection and sanitization of integer bugs. A crucial reason for this is the inherent difficulty of prevention mechanisms to differentiate between actual—and potentially vulnerable—integer errors, from *intentional* uses of wrap-around behavior, type castings, bit shifts, and other constructs that serve application-specific purposes. Furthermore, existing tools usually focus only on certain integer error classes (mainly examine overflows and underflows). On the other hand, tools with broader scope typically generate a large amount of false positives (in case of static analysis tools) or provide poor source code coverage (in case of dynamic analysis or dynamic test generation tools). These drawbacks prevent their widespread adoption, as developer-intended constructs

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

are typically regarded as errors, and developers are reluctant to make major changes in existing code bases and limit their use of performance enhancing idioms.

As a step towards addressing this issue, in this paper we propose an approach that combines static arithmetic error detection with information flow tracking to improve the accuracy of arithmetic error detection, by focusing on reducing the number of false positives, i.e., developer-intended code constructs that violate language standards. Our tool, IntFlow, uses information flow tracking to reason about the severity of arithmetic errors by analyzing the information flows related to them. The main intuition behind this approach is that arithmetic errors become critical when i) they are triggered by or depend on values originating from untrusted locations, or ii) a value affected by an arithmetic error propagates to sensitive locations such as the arguments of operations like `malloc()` and `strcpy()`.

To demonstrate the effectiveness of our approach, we evaluated IntFlow on real world programs and vulnerabilities and verified that it successfully identifies all the real world vulnerabilities for the applications of our testbed, producing 89% less false positives than state of the art static analysis checkers [10].

Our work makes the following contributions:

- We present an accurate compile-time arithmetic error detection approach that combines information flow tracking and static code analysis.
- We present IntFlow, our prototype implementation for this approach, which operates as an LLVM add-on. IntFlow is freely available as an open source project and, besides static error detection, it can be used as an online defense mechanism for certain classes of applications, as we demonstrate in our evaluation section.
- We evaluate IntFlow using real-world programs and vulnerabilities. Our results demonstrate that IntFlow achieves improved detection accuracy compared to previous solutions, as it suppressed more than 89% of the false positives reported by IOC [10], our reference integer error checking tool.

The rest of this paper is organized as follows: Section 2 provides some background information, and Section 3 present our overall approach. Section 4 describes in detail IntFlow’s design and implementation, off-line analysis methodology, and optimizations, and in Section 5, we present the results of our experimental evaluation using our prototype implementation. After discussing related work in Section 6, we discuss the limitations of our approach in Section 7 and conclude in Section 8.

2. BACKGROUND

To effectively deal with integer errors in real world applications, it is necessary to first define what is considered an error. Doing so is not trivial, as apart from the mere examination of conformance to the language standard, we must also examine whether pieces of seemingly erroneous code—from the perspective of the language specification—are in reality explicitly written in that way due to the developer’s intention, typically for performance, convenience, or other reasons.

In this section, we discuss how the C/C++ language standards define correctness for arithmetic operations, and examine why developers often write code that deviates from the language specification. We also present examples of exploitable vulnerabilities caused by integer errors, and demonstrate the importance of good programming practices.

2.1 Integer Errors and Undefined Behavior

Although the C/C++ standards explicitly define the outcome of most integer operations, they still leave a number of corner cases undefined. As an example, the C11 standard considers an `unsigned` integer overflow a well-defined operation, whose result is the minimum value obtained after the wrap-around. In contrast, it leaves `signed` integer overflows undefined. This choice leaves room for compiler implementations to handle them in a way that produces optimized binaries [22]. For instance, signed integer overflows (or underflows) enable compiler developers to implement constant propagation [5], an optimization that infers invariants from expressions such as `i+1 < i` and replaces them with a constant Boolean value.

Table 1 lists special cases of integer operations and their definedness. It should be noted that although more instances of undefined behavior (not necessarily restricted to integer operations) are declared in the language specification, we only consider those relevant to this work.

Arithmetic Operation	Definedness
Unsigned overflow (underflow)	defined
Singed overflow (underflow)	undefined
Signedness conversion	undefined*
Implicit type conversion	undefined*
Oversized shift	undefined
Division by zero	undefined

*if value cannot be represented by the new type

Table 1: Summary of defined and undefined arithmetic operations according to the C/C++ language specification.

As in practice not all cases of undefined behavior necessarily result in actual errors, the difficulty of dealing with these types of bugs lies in distinguishing *critical* integer errors from *developer-intended* errors. The intention of a developer, however, cannot be formally defined or automatically derived, as the code patterns that are present on a piece of software are deeply related to each developer’s knowledge, ability, and programming style.

Although writing code that intentionally relies on undefined operations is generally considered a bad programming practice (as the outcome of those operations can be arbitrary, depending on the architecture or the compiler), there are several cases in which the community has reached consensus on what is the expected behavior of the compiler in terms of the generated code, mainly due to empirical evidence. This explains why we still see idioms that actively take advantage of undefined behavior so frequently: although, *according to the standard*, the result of the operation is undefined, developers have an empirically derived expectation that compilers will always handle such cases in a consistent manner.

This expectation creates serious complications whenever developers check the validity of their code with state-of-the-

```

1  UINT_MAX = (unsigned) -1;
2  INT_MAX = 1 << (INT_WIDTH - 1) + 1;

```

Listing 1: Widely used idioms that according to the standard correspond to undefined behavior.

```

1  /* struct containing image data, 10KB each */
2  img_t *table_ptr;
3  unsigned int num_imgs;
4  num_imgs = get_num_imgs();
5  ...
6  unsigned int alloc_size = sizeof(img_t) * num_imgs;
7  ...
8  table_ptr = (img_t*) malloc(alloc_size);
9  ...
10 for (i = 0; i < num_imgs; i++)
11     { table_ptr[i] = read_img(i); } /* heap overflow */

```

Listing 2: An unsigned integer overflow as a result of a multiplication operation (line 8), which in turn results in an invalid memory allocation (line 8) and unintended access to the heap (line 13).

art static analysis tools. These tools evaluate code based on strict conformance to the language specification, and consequently generate a large amount of false positives. Thus, the generated reports are often overlooked by developers who struggle to distinguish which of the reported bugs are actual errors and which are not. Unfortunately, tools based on dynamic code analysis also do not provide strong guarantees in these cases, as they suffer from low code coverage.

To further shed light on the complexity of this issue, in the following we present two characteristic integer error examples and discuss the complications introduced by the use of undefined operations.

2.2 Integer Error Examples

While the task of automatically detecting undefined integer operations is relatively easy, the true difficulty lies in identifying the developer’s intention behind their legitimate or illegitimate use of arithmetic operations that do not conform to the language standard.

Listing 1 presents two C statements in which developers intentionally rely on undefined behavior, mainly for performance reasons. Both are based on assumptions on the numerical representation used by the underlying system (two’s complement). Line 1 shows a case of signedness casting in which the original value cannot be represented by the new type. In Line 2, a shift operation of `INT_WIDTH - 1` is also undefined¹ but it conventionally returns the minimum value of the type, while the addition operation incurs a signed overflow whose behavior is again undefined. Although these cases are violations of the language standard, the desirable operation of an integer overflow checker would be to *not* report them, as they correspond to developer-intended behavior—otherwise, such cases are considered *false positives* [10].

In contrast, in the example of Listing 2, the unsigned integer variable (`num_imgs`) might overflow as a result of

¹It is undefined according to the C99 and C11 standards. The ANSI C standard defines this behavior.

the multiplication operation at line 8. This behavior is well-defined by the standard, but the overflow may result in the allocation of a memory chunk of invalid (smaller) size, and consequently, to a heap overflow. An effective arithmetic error checker should be able to identify such dangerous cases, which can potentially lead to exploitable vulnerabilities, as it is clear that the developer did not intend for this behavior.

3. APPROACH

The security community is still failing to completely eliminate the problem of integer errors even after years of effort [10, 17, 23, 21]. One of the main reasons is the problem of differentiating between *actual* errors, which can potentially lead to reliability or security flaws, from intended uses of undefined constructs, as a result of concise developer choices—the latter are still considered errors by rigorous static checkers that strictly follow language standards. We differentiate between the two by considering *developer-intended* and *critical* arithmetic errors. Before describing the design of IntFlow, we first provide a more concrete definition of what we consider as critical arithmetic errors.

Definition 1. *An arithmetic error is potentially critical if it satisfies one of the following conditions:*

1. *At least one of the operands in an erroneous statement originates from an untrusted source. Examples of untrusted sources include the network, user input, files, and so on.*
2. *The result of an erroneous operation may propagate to a sensitive program location. Examples of sensitive locations include memory allocation functions (`mmap()`, `malloc()`, ...), string handling functions (`strcpy()`, `strcat()`, ...), and so on.*

As capturing the intention of developers is a hard problem, IntFlow focuses on the detection of arithmetic errors that might constitute exploitable vulnerabilities or cause reliability issues. This is achieved not only by focusing on the identification of arithmetic errors (according to the language standard), which in itself is a tractable problem, but also by considering the information flows that affect the erroneous code. The rationale behind this definition is that (a) arithmetic errors influenced by external inputs are probably not developer-intended, and may potentially be exploited through carefully crafted inputs, and (b) arithmetic errors typically result in severe vulnerabilities only when they affect inputs to sensitive library and system operations.

Figure 1 illustrates different types of information flows that may involve erroneous arithmetic operations. Critical errors are related to information flows that originate from untrusted inputs, or that eventually reach a sensitive operation, such as a system call, through value propagation. In cases where the input of an arithmetic operation is untrusted (condition 1.1) or a sensitive sink is reached (condition 1.2), the error is flagged as critical. On the other hand, arithmetic errors influenced only by benign inputs are considered less likely to be used in exploitation attempts.

The information flow based approach enables us to handle cases similar to the examples presented in Section 2.2. IntFlow can also silence error reports caused by statements similar to those presented in Listing 1, since it can trace that the origin of the value that lead to the undefined behavior is a




Figure 1: Information flows to and from the location of an arithmetic error.

constant initialization which is *de facto* developer-intended. The code snippet in Listing 2, contains two different types of flows. The first one connects `get_num_imgs()` (line 6) with the multiplication operation (line 8) while the second one connects the result of the multiplication with a memory allocation function (`malloc()`) which is considered a sensitive system location. The former is a Type 1 flow because the value of the multiplication operand originated from an untrusted input location, while the latter is a Type 2 flow since the result of the operation affects a sensitive function call. IntFlow would retain the checking mechanism in both cases, ensuring that a maliciously crafted input leading to an arithmetic error would be detected and reported.

4. DESIGN AND IMPLEMENTATION

In this section we present the design and implementation of IntFlow, a tool that combines Information Flow Tracking [18] with a custom integer error checking tool [10] to improve the accuracy of arithmetic error detection. The main goal of IntFlow is to reduce the number of false positives produced by previous static arithmetic error checkers. In this context, the term “false positive” refers to reporting developer-intended errors (according to our definition in Section 2.1) as critical errors. Although from the perspective of the language standard these correspond to erroneous code, the prevalence of such constructs makes reports of such issues a burden for security analysts, who are interested only in critical errors that may form exploitable vulnerabilities.

IntFlow integrates a set of components to its custom analysis engine to perform data flow tracking along with integer error checking. An integer error checker [10] instruments potentially erroneous operations and a static information tracking tool [18] performs inter-procedural analysis. In the following, we discuss the rationale of our decision to use these components along with a detailed description of each technology. Moreover we describe their integration into a system that reduces the number of reported developer-intended errors by leveraging information flow tracking.

4.1 Main Components

IOC operates at the Abstract Syntax Tree (AST) produced by Clang [1], a C/C++ front-end of LLVM [14]. It instruments all arithmetic operations such as additions, multiplications and shifts, as well as most of unary, casting and type conversion operations. In contrast to previous tools that focus on a subset of integer errors (typically overflows and underflows), IOC provides protection against a broader set of integer errors. Even though it focuses mainly on errors

with undefined behavior based on the language standards, it can also protect against errors that do not fall into this category, covering most of the integer error classes presented in Table 1.

IOC instruments *all* arithmetic operations that may lead to an erroneous result, and inserts checks accordingly. Essentially, for each integer operation inside a basic block, additional basic blocks that implement the error-checking logic are inserted into the control flow graph, and users are allowed to register callback functions for error handling. The fact that IOC instruments blindly all arithmetic operations is a major source of false positives, common in many integer error detection systems thus IntFlow actively provisions for reducing false positives by eliminating checks for developer-intended errors. IOC is a major component of our architecture, as it provides assurance that all potentially serious arithmetic errors can be checked. It is then up to the information flow analysis to differentiate between developer-intended and critical errors, and report only the latter.

For IntFlow’s information tracking mechanism, we employ `llvm-deps` [18], an LLVM compiler pass implementing static information flow tracking in a manner similar to the classic data flow analysis [5]. It is designed as a context sensitive (inter-procedural) analysis tool that allows forward and backward slicing on source and sink pairs of our choice, implementing the DSA [15] algorithm. DSA performs context-sensitive, unification-based points-to analysis that allows us to track data flows among variables referred by pointer aliases. It is important to note that the analysis scope of `llvm-deps` is limited to a single object file, as it is implemented as a compile-time optimization pass and not as a link-time optimization pass.

4.2 Putting It All Together

Figure 2 illustrates the overall architecture of IntFlow: First, IOC adds checks to the integer operations that are exposed by Clang in the AST and then `llvm-deps` performs static IFT analysis on the LLVM intermediate representation (IR).

To reduce unnecessary checks that lead to false positives, IntFlow uses `llvm-deps` to examine only certain flows of interest. Specifically, IntFlow examines only flows stemming from untrusted sources, or ending to sensitive sinks, with respect to Definition 1 in Section 2.1. Initially, IntFlow performs *forward* slicing: starting from a particular source—essentially a program variable—used in a potentially erroneous arithmetic operation, it examines whether the result of the operation flows into sinks of interest. Once such a variable is found, IntFlow performs *backward* slicing, to verify that the sink is actually affected by that particular source. Since the flow tracking mechanism does not offer full code coverage, we employ this two-step process to gain confidence on the accuracy of the flow and verify its validity. Once the source is reached when using backward slicing starting from the sink, the flow is considered established.

4.3 Modes of Operation

As discussed in Section 3, IntFlow uses two different types of information flow to pinpoint errors. The first associates untrusted inputs with integer operations whilst the second associates the outcome of integer operations with its use in sensitive system operations. Once IOC inserts checks in all arithmetic operations that may lead to an error, IntFlow




Figure 2: Overall architecture of IntFlow.

eliminates unnecessary checks operating in one of the following modes:

- In *whitelisting* mode, all checks for operations whose arguments come from trusted sources are removed.
- In *blacklisting* mode, IntFlow only maintains checks for operations whose operands originate from untrusted sources and *removes all other checks*.
- In *sensitive* mode, IntFlow only maintains checks for operations whose results may propagate into a sensitive sink.

In the following we elaborate more on the specifics of each of the previous modes:

4.3.1 Trusted and Untrusted Inputs

For each operation that may result in an arithmetic error, IntFlow’s IFT analysis determines where the values of the involved operands originate from. If their origin corresponds to some input source, IntFlow classifies it into either trusted or untrusted, and handles it accordingly, using one of the following two modes of operation.

White-listing: Erroneous arithmetic operations for which all operands originate from trusted sources are unlikely to be exploitable. Thus, for those cases, IntFlow safely removes the error checks inserted by IOC at the instrumentation phase. Before an operation is verified as safe, IntFlow needs to examine the origin of *all* data flowing to that operation. Values derived from constant assignments or from safe system calls and library function calls, e.g., `gettimeofday()` or `uname()`, are typical examples of sources that can be trusted, and thus white-listed.

Black-listing: Input sources that can be affected by users are considered untrusted, since carefully crafted inputs might lead to integer errors, and potentially, successful exploitation. If any of the operands has a value affected by such a source, IntFlow retains the error checking instrumentation. System and library calls that read from untrusted sources, such as `read()` and `recv()`, are examples of this type of sources.

By combining the above two approaches, IntFlow selects the unsafe integer operations that will be instrumented with protection checks. These modes of operation can be complemented by IntFlow’s third mode, which refines the analysis results for the surviving checks.

4.3.2 Sensitive Operations

In this mode of operation, IntFlow reports flows that originate from integer error locations and propagate to sensitive sinks, such as memory-related functions, system calls

etcetera. Moreover, in contrast to the previous mode, whenever an integer error occurs, the error is not reported at the time of its occurrence, but only once it propagates as input into one of the sensitive sinks. This is very effective in suppressing false positives, since errors that do not lead to a sensitive operation from some path in the control flow graph are not generally exploitable. In the following, we explain how IntFlow achieves runtime detection of errors whenever execution reaches a sensitive sink.

To report errors at the sensitive sink locations, IntFlow performs the following operations:

- Initially, the tool identifies all integer operations whose results may propagate into a sensitive sink at runtime. This step further eliminates checks that do not lead into a sensitive sink and therefore are probably not exploitable. A global array is created for each sensitive sink, holding one entry per arithmetic operation affecting the sink.
- Whenever an integer operation produces an erroneous result, the respective entry for this operation in the affected global tables is set to `true` (denoting that the flow originated from an operation that produced an erroneous result). If the execution continues and does not reach the sensitive location, the respective entries are set to `false`, denoting that the result of the sensitive operation will not be affected by the integer operation.
- If the execution reaches a sensitive operation, the respective global array is examined. If one or more entries are set to `true`, execution is interrupted since an erroneous value from the integer operation can have affect on the sensitive system operation.

While it is better to combine the two modes to establish end-to-end monitoring and detection of suspicious flows, we can also use each mode of IntFlow independently: the first mode to generally reduce the number of false positives and the second mode to detect exploitable vulnerabilities. In Section 5, we evaluate the effectiveness of IntFlow using the white-listing mode and in Section 7 examine further the characteristics of each of its mode of operations suggesting their ability to not only detect critical errors but also to evaluate the severity of each reported error.

4.4 Implementation Details

IntFlow is implemented as an LLVM [14] pass written in ~3,000 lines of C++. Briefly, it glues together its two

main components (IOC and `llvm-deps`) and supports fine-tuning of its core engine through custom configuration files. In the following, we outline the main characteristics of IntFlow regarding its core engine and its support for developer customizations.

4.4.1 Pass Characteristics

With respect to IntFlow’s integration with Clang, IntFlow can be invoked by simply passing the appropriate flags to the compiler, without any further action needed from the side of the developer. Although IOC has been integrated into the LLVM main branch since version 3.3, for the current implementation of IntFlow we used an older branch of IOC that supports a broader set of error classes than the latest. IntFlow’s LLVM pass is placed at the earliest stage of the LLVM pass dependency tree, in order to prevent a subsequent optimization pass from optimizing away the semantics of integer operations. During the compilation process, arithmetic error checks are inserted by IOC, and subsequently selectively filtered by IntFlow. Any basic blocks corresponding to IOC checks that have been disabled by IntFlow are then optimized away by the compiler.

4.4.2 Configuration Options

IntFlow offers developers the option to explicitly specify arithmetic operations or sources that need to be whitelisted or blacklisted. In addition, it can be configured to overall exclude any specific file from its analysis or ignore specific lines in the source code. Developers can also specify the operation mode (whitelisting, blacklisting or sensitive-sink analysis) that IntFlow will use, as well as override or extend the default set of sources and sinks that will be considered during information flow analysis. Finally, they can specify particular callback actions that will be triggered upon an error, such as activating runtime logging or exiting with a suitable return value whenever an error occurs. These features offer great flexibility to developers, enabling them to fine-tune the granularity of the generated reports and adjust the built-in options of IntFlow to the exact characteristics of their source-code.

5. EVALUATION

In this section, we present the results of our experimental evaluation using our prototype implementation of IntFlow. To assess the effectiveness and performance of IntFlow, we look into the following aspects:

- What is the accuracy of IntFlow in detecting and preventing critical arithmetic errors?
- How effective is IntFlow in reducing false positives? That is, how good is it in omitting developer-intended errors from the reported results?
- When used as a protection mechanism, what is the runtime overhead of IntFlow compared to native execution?

Our first set of experiments aims to evaluate the tool’s ability to identify and mitigate critical errors. For this purpose, we use two datasets, consisting of artificial and real-world vulnerabilities, respectively. Artificial vulnerabilities were inserted to a set of real-world applications, corresponding to various types of MITRE’s Common Weakness Enumeration (CWE) [4]. This dataset provides a broad test

suite that contains instances of many different types of arithmetic errors, which enables us to evaluate IntFlow in a well-controlled environment, knowing exactly how many bugs have been inserted, and the nature of each bug. Likewise, our real-world vulnerability dataset consists of applications such as image and document processing tools, instant messaging clients, and web browsers, with known CVEs, allowing us to get some insight on how well IntFlow performs against real-world, exploitable bugs.

In our second round of experiments, we evaluate the effectiveness of IntFlow’s information flow tracking analysis in reducing false positives, by running IntFlow on the SPEC CPU2000 benchmark suite and comparing its reported errors with those of IOC. IOC instruments all arithmetic operations, providing the finest possible granularity for checks. Thus, by comparing the reports produced by IntFlow and IOC, we obtain a base case for how many developer-intended errors are correctly ignored by the IFT engine.

Finally, to obtain an estimate of the tool’s runtime overhead, we run IntFlow over a diverse set of applications of varying complexity, and establish a set of performance bounds for different types of binaries. All experiments were performed on a system with the following characteristics: 2× Intel(R) Xeon(R) X5550 CPU @ 2.67GHz, 2GB RAM, i386 Linux.

5.1 Accuracy Evaluation

In all experiments for evaluating accuracy, we configured IntFlow to operate in *whitelisting* mode, since this mode produces the greatest number of false positives, as it preserves most of the IOC checks among the three modes. Thus, whitelisting provides us with an estimation of the worst-case performance of IntFlow, since the other two modes will always be implementing a more fine-tuned instrumentation.

5.1.1 Evaluation Using Artificial Vulnerabilities

To evaluate the effectiveness of IntFlow in detecting critical errors of different types, we used seven popular open-source applications with planted vulnerabilities from nine distinct CWE categories.² Table 2 provides a summary of the applications used and the respective CWEs.

Each application is replicated to create a set of test-case binaries. In every test-case binary—essentially an instance of the real-world application—a vulnerability of a given type is planted and then the application is compiled with IntFlow. Subsequently, each test-case binary is executed over a set of benign and malicious inputs (inputs that exploit the vulnerability and result in abnormal behavior). A correct execution is observed when the binary executes normally on benign inputs or terminates before it can be exploited on malicious inputs.

Overall, IntFlow was able to correctly identify 79.30% (429 out of 541) of the planted artificial vulnerabilities. The 20.7% missed are due to the accuracy limitations of the IFT mechanism, which impacts the ability of IntFlow to correctly identify flows, and also due to vulnerabilities triggered by implicit information flows (i.e., non-explicit data flows realized by alternate control flows), which IFT is inherently unable to detect. We discuss ways in which accuracy can be further improved in Section 7.

²All modified applications for this experiment were provided to us by MITRE [6].

Applications	CWEs
Cherokee 1.2.101	CWE-190
Grep 2.14	CWE-191
Nginx 1.2.3	CWE-194
Tcpdump 4.3.0	CWE-195
W3C 5.4.0	CWE-196
Wget 1.14	CWE-197
Zshell 5.0.0	CWE-369
	CWE-682
	CWE-839

Table 2: Summary of the applications and CWEs used in the artificial vulnerabilities evaluation.

Program	CVE Number	Type	Detected?
Dillo	CVE-2009-3481	Integer Overflow	Yes
GIMP	CVE-2012-3481	Integer Overflow	Yes
Swftools	CVE-2010-1516	Integer Overflow	Yes
Pidgin	CVE-2013-6489	Signedness Error	Yes

Table 3: CVEs examined by IntFlow.

To examine IntFlow’s ability to ignore developer-intended errors, we exercised IntFlow over sets of benign inputs on vanilla versions of each application. Since those inputs produce the expected output, we assume that all the errors they trigger are developer-intended either explicitly or implicitly. With this in mind, we compared the error checks of IntFlow with those of IOC to quantify the ability of our tool in removing unnecessary checks. Overall, IntFlow eliminated 90% of the false checks (583 out of 647) when tested with the default set of safe inputs. This reduction was achieved due to the successful identification of constant assignments and the whitelisting of secure system calls, as discussed in Section 3.

5.1.2 Mitigation of Real-world Vulnerabilities

In our next experiment, we examined the effectiveness of IntFlow in detecting and reporting real-world vulnerabilities. For this purpose, we used four widely-used applications and analyzed whether IntFlow detects known integer-related CVEs included in these programs. Table 3 summarizes our evaluation results. IntFlow successfully detected *all* the exploitable vulnerabilities under examination. From this small-scale experiment, we gain confidence that IntFlow’s characteristics are maintained when the tool is applied to real world programs, and therefore is suitable as a detection tool for real-world applications.

5.1.3 False Positives Reduction

Reducing the number of false positives is a major goal of IntFlow, and this section focuses on quantifying how effective this reduction is. Our first measurement is made using SPEC CPU2000, a suite that contains C and C++ programs representative of real-world applications. Since IOC is a core component of IntFlow, we chose to use CPU2000 to examine what are the improvements of IntFlow’s IFT in comparison to the results of IOC as those were previously reported [10]. We ran the SPEC benchmarks using the “test” data sets for both IOC and IntFlow, so that we could manually analyze all the reports produced by IOC and classify them as true




Figure 3: Number of critical and developer-intended arithmetic errors reported by IOC and IntFlow for the SPEC CPU2000 suite. IntFlow identifies the same number of critical errors, while it reduces significantly the number of reported developer-intended errors.

	Overall	Dillo	Gimp	Pidgin	SWFTools
IOC	330	31	231	0	68
IntFlow	82	26	13	0	43

Table 4: Reported errors by IOC and IntFlow for some of the vulnerable real-world programs.

or false positives (i.e., critical or developer-intended errors, respectively). Once all reports were categorized based on Definition 1, we examined the respective results of IntFlow. We report our findings in Figure 3.

IntFlow was able to correctly identify all the critical errors (64 out of 64) that were triggered during execution, and reduced the reports of developer-intended errors by $\sim 89\%$ (419 out of 471).

Real-world Applications.

In Section 5.1.2 we demonstrated how IntFlow effectively detected known CVEs for a set of real-world applications. Here, we examine the reduction in false positives achieved when using IntFlow’s core engine instead of static instrumentation with IOC alone. To collect error reports, we ran each application with benign inputs as follows: for Gimp, we scaled the ACSAC logo and exported it as gif; for SWFTools, we used the pdf2swf utility with a popular e-book as input; for Dillo, we visited the ACSAC webpage and downloaded a PDF with the list of notable items published in 2013; and for Pidgin, we performed various common tasks (registering a new account, logging-in and out, and so on). Table 4 shows the reports produced by IOC and IntFlow respectively.

Overall, IntFlow was able to suppress 75% of the errors reported by IOC during the execution of the applications on benign inputs. Although this evaluation does not provide full coverage on the number of generated reports (for instance, we did not observe any false positives for pidgin with the tests we performed), it allows us to obtain an estimate of how well IntFlow performs in real world scenarios. It should be noted that the effectiveness in the reduction of false pos-

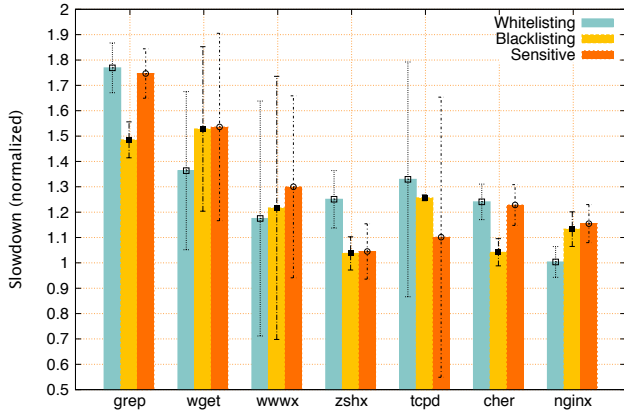


Figure 4: Runtime overhead for the applications of Section 5.1.1 (normalized over native execution).

itives is highly dependant on the nature of each application, as well as on the level of the execution’s source coverage. That is, the more integer operations occur throughout the execution, the greater the expected number of false positives. For instance, Gimp’s functionality is tightly bound to performing arithmetic operations for a number of image processing actions, and thus IOC reports many errors, most of which are developer-intended, while Dillo does not share the same characteristics and as a result exhibits a smaller reduction in false positives.

5.2 Runtime Overhead

Although IntFlow was not designed as a runtime detection tool but rather as an offline integer error detection mechanism, one may wonder whether it could be customized to offer runtime detection capabilities. In this section, we seek to examine the performance of IntFlow for various applications, when running them with all the automatically inserted arithmetic error checks. For this purpose, we perform a set of timing measurements on the applications used in Section 5.1.1. For each run, we measured the time that was required to complete a series of tasks for each of IntFlow’s modes of operation, and then normalized the running time with respect to the runtime of the native binary. Reported results are mean values over ten repetitions of each experiment, while the reported confidence intervals correspond to 95%. We ran all binaries natively, and measured user time with the `time` utility.

For *grep*, we search for strings matching a complex regular expression in a 1.2GB file. For *wget*, we download a 1GB file over a 1Gbps link from a remote server. To measure *wwwx*, we perform verbose GET requests on a series of websites. For *zshx*, we execute a series of commands that are internal to the shell. For the web servers, Cherokee was configured for IPv4 only, while for Nginx all configurations options were left to their default setting. We measured both servers’ performance using Apache’s *ab* benchmarking utility and static HTML files.

Figure 4 shows the results of our evaluation. IntFlow incurs a high overhead in applications that involve a large number of integer operations, such as *grep*.³ Although in

³Our experience with the SPEC CPU2000 benchmarks

such applications the overhead is rather prohibitive, and cancels out the benefits of using a static mechanism, in other cases, such as the *cher* and *nginx* servers, the overhead is within an acceptable 20%. Thus, it could be the case that IntFlow might be used as a runtime defense for certain types of applications, i.e., I/O-bound. As each of IntFlow’s modes of operation targets different flows and can be fine-tuned by developers, customization can result in different overheads, as different flows dominate the execution of different applications. This is the reason we observe different slowdowns per mode: depending on whether the dominating flows involve sensitive calls (as is the case with web servers), the sensitive mode will be slower or faster than the other two modes, and so on.

6. RELATED WORK

During the past years, as the protection mechanisms against buffer overflows became more mature, great focus was placed upon efficiently dealing with integer overflows. This section summarizes the main characteristics of the several approaches that have been followed so far for addressing integer overflows and outlines the connection between the current work and existing research on the field.

6.1 Static Analysis

Static analysis tools provide good coverage but generally suffer from a high rate of false positives. IntPatch [23] is built on top of LLVM [14] and detects vulnerabilities utilizing the type inference of LLVM IR. Similarly to our tool, IntPatch uses forward & backward analysis to classify sources and sinks as sensitive or benign. Each sensitive variable is located through slicing. If a variable involved in an arithmetic operation has an untrusted source and the respective sink may overflow, IntPatch will insert a check statement after that vulnerable arithmetic operation. If an overflow result is used for sensitive actions such as memory allocations, IntPatch considers it a real vulnerability. Contrary to the current work though, IntPatch does not deal with all types of integer overflows and also does not address programmer-inserted sanitization routines.

KINT [21] is a static tool that generates constraints representing the conditions under which an integer overflow may occur. It operates on LLVM IR and defines untrusted sources and sensitive sinks via user annotations. KINT avoids path explosion by performing constraint solving at the function level and by statically feeding the generated constraints into a solver. After this stage, a single path constraint for all integer operations is generated. Unfortunately, despite the optimization induced by the aforesaid technique, the tool’s false positives remain high and there is a need for flagging false positives with manual annotations in order to suppress them. Moreover, contrary to this work, KINT attempts to denote all integer errors in a program and does not make a clean distinction between classic errors and errors that constitute vulnerabilities.

SIFT [16] uses static analysis to generate input filters against integer overflows. If an input passes through such filter, it is guaranteed not to generate an overflow. Initially,

showed that the overhead on benchmarks with very frequent integer operations, such as *gzip*, is in the range of $\sim \times 10$, prohibiting IntFlow from being used as a generic runtime detection mechanism for such applications.

the tool creates a set of critical expressions from each memory allocation and block copy site. These expressions contain information on the size of blocks being copied or allocated, and are propagated backwards against the control flow, generating a symbolic condition that captures all the points involved with the evaluation of each expression. The free variables in the generated symbolic conditions represent the values of the input fields and are compared against the tool’s input filters. A significant difference of this paper in comparison to SIFT is that the latter nullifies overflow errors but does not detect them nor examines whether they could be exploitable.

IntScope [19] decompiles binary programs into IR and then checks lazily for harmful integer overflow points. To deal with false positives, IntScope relies on a dynamic vulnerability test case generation tool to generate test cases which are likely to cause integer overflows. If no test case generates such error, the respective code fragment is flagged appropriately. This approach varies significantly from the one used in our tool as it relies on the produced test cases to reveal a true positive: if a test case does not generate an overflow, that does not guarantee that no overflow occurs. In addition, IntScope regards all errors as generic, without focusing particularly on errors leading to vulnerabilities.

Finally, RICH [7] is a compiler extension which enables programs to monitor their own execution and detect potential attacks exploiting some integer vulnerability. Although RICH is very lightweight, it does not handle cases of pointer aliasing and also produces false positives in cases where developers intentionally abuse undefinedness features of the C/C++ standard, both of which are basic components of IntFlow’s design.

6.2 Dynamic & Symbolic Execution

Dynamic Execution tools use runtime checks to prevent unwanted program behavior. IOC [10] uses dynamic analysis to detect the occurrence of overflows in C/C++ programs. The tool performs a compiler-time transformation to add inline numerical error checks and then relies on a runtime handler to prevent any unwanted behavior. The instrumentation transformations operate on the Abstract Syntax Tree (AST) in the Clang front-end, after the parsing, type-checking and implicit type conversion stages. IOC checks for overflows both in shifting and arithmetic operations and makes a clear distinction between well-defined and undefined behaviors, that is, it does not consider all overflows malicious. Our tool adopts this perspective and complements IOC in the sense that it addresses the issue of high false positives by integrating static instruction flow tracking to the performed analysis.

Symbolic Execution tools provide low false positives but can’t easily achieve full coverage. They usually use dynamic test generation [11] to detect violations. SmartFuzz [17] generates a set of constraints defining unwanted behavior and determines whether some input could trigger an unwanted execution path. This tool does not require source code and makes use of the Valgrind framework for its symbolic execution and scoring. Coverage and bug-seeking queries are explored in a generational search, whilst queries from the symbolic traces are solved generating a set of new test cases. Thus, a single symbolic execution feeds the constraint solver with many queries, which themselves generate new test cases etc. KLEE [8] uses symbolic execution to automatically gen-

erate tests that achieve high coverage for large-scale programs but it is not focused on integer errors, thus it does not achieve as good results against integer overflows as other tools that targeted towards integer operations.

7. DISCUSSION

7.1 Static Information Flow Tracking

A core component of IntFlow is `llvm-deps` [18], which, as an implementation of static information flow tracking, is expected to provide good source code coverage with low runtime overhead. However, we should note that `llvm-deps` suffers from inherent inaccuracy issues, largely due to the limitations of its points-to analysis [15] and due to its data flow analysis mechanism. These limitations are amplified when one wishes to extend the scope of the technique by performing inter-procedural analysis. Fortunately, as our experience revealed, sources and sinks typically reside within a single function. This can be viewed as an instance of the classic trade-off between accuracy and performance: for cases where accuracy has the maximum priority, we may choose to incorporate dynamic IFT [13] and attempt to reduce any increased runtime overhead using techniques that combine static and dynamic analysis [12, 9].

7.2 IntFlow for Runtime Detection

While the primary use case of IntFlow is to help users analyze existing code bases at during the development phase by reducing the amount of false positives, another use case is to deploy the tool as a runtime defense against zero-day vulnerabilities. For this purpose, the two main issues that must be addressed are i) the increased runtime overhead due to the inserted checks, and ii) any remaining false positives after IntFlow’s analysis.

7.2.1 Fast Integer Error Checking

The main source of runtime overhead, as shown in Section 5.2, can be attributed to IOC’s checks, as it replaces each arithmetic operation with at least three basic blocks to perform the checking operation. Given the significance of the problem, there have been many previous proposals for implementing fast and efficient checking operations [20], so IntFlow could adopt them to improve performance.

7.2.2 False Positives

As shown in Section 5, IntFlow was able to identify a large portion of the developer-intended errors in the programs under examination, but still missed some cases. In order to provide as a broad coverage of false positives as possible, IntFlow supports manual labeling of false positives. Developers can dedicate a separate off-line phase to apply IntFlow using a trusted input set over their application, and pinpoint the locations in which IntFlow falsely flags a benign operation as malicious.

A more suitable solution for this use case would be the incorporation of dynamic IFT, which again would impose high runtime overhead. For this, we can consider an efficient dynamic IFT approach [9] that again leverages DataFlow analysis.

7.3 Quality of the Produced Reports

Another advantage of IntFlow’s design is the fact that its three different modes of operation offer an estimation of how

critical a particular bug is. Errors reported by the sensitive mode have the highest risk, as they involve sensitive operations and are more likely to be exploitable. Likewise, in black-listing mode, IntFlow examines flows originating from untrusted locations and thus the produced reports are of moderate priority. Finally, the whitelisting mode is likely to generate the largest amount of warnings. Thus, if developers wish to examine as few locations as possible, e.g., due to limited available time for performing code auditing, they can first examine the reports generated by IntFlow in the sensitive mode, and if time permits, the follow with the blacklisting mode, and so on.

Throughout our evaluation, we noticed that many of the reports generated by IntFlow follow a particular pattern, mainly due to code reuse from the side of developers. We believe that using simple pattern matching and lexical analysis of the source code, in combination with the reports of IntFlow, could further increase the accuracy with which IntFlow classifies errors as malicious or not—the more the occurrences of a particular error, the less likely for this error to not be developer-intended. We will explore this approach as part of our future work.

8. CONCLUSION

We have presented IntFlow, a tool that identifies a broad range of arithmetic errors, and differentiates between actual, critical errors and developer-intended constructs that rely on undefined behavior but do not constitute potential vulnerabilities. IntFlow uses static information flow tracking to associate flows of interest with erroneous statements, and greatly reduces false positives without removing checks that would prevent the detection of critical errors. The results of our evaluation demonstrate the effectiveness of IntFlow in distinguishing between the two types of errors, allowing developers and security analysts to detect and fix critical errors in an efficient manner, without the need to sift through numerous non-critical developer-intended errors. The significant reduction in false positives that IntFlow achieves over IOC, which has been integrated into Clang since version 3.3, demonstrates the need for arithmetic error protection and the importance of our contribution.

9. REFERENCES

- [1] Clang C language family frontend for LLVM. <http://clang.llvm.org/>.
- [2] CVE - CVE-2006-3824. <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2006-3824>.
- [3] CWE - 2011 CWE/SANS top 25 most dangerous software errors. <http://cwe.mitre.org/top25/>.
- [4] CWE - Common Weakness Enumeration. <http://cwe.mitre.org/>.
- [5] A. V. Aho, R. Sethi, and J. D. Ullman. *Compilers: Principles, Techniques, and Tools*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1986.
- [6] A. Benameur, N. S. Evans, and M. C. Elder. Minestrone: Testing the soup. In *Proceedings of CSET*, Washington, D.C., 2013. USENIX.
- [7] D. Brumley, T. Chiueh, R. Johnson, H. Lin, and D. Song. RICH: Automatically Protecting Against Integer-Based Vulnerabilities. *Proceedings of NDSS*, page 28, 2007.
- [8] C. Cadar, D. Dunbar, and D. R. Engler. KLEE: Unassisted and Automatic Generation of High-Coverage Tests for Complex Systems Programs. *OSDI*, 2008.
- [9] W. Chang, B. Streiff, and C. Lin. Efficient and extensible security enforcement using dynamic data flow analysis. *Proceedings of the 15th ACM conference on Computer and communications security*, pages 39–50, 2008.
- [10] W. Dietz, P. Li, J. Regehr, and V. Adve. Understanding integer overflow in C/C++. In *Proceedings of ICSE*. IEEE Press, June 2012.
- [11] P. Godefroid, N. Klarlund, and K. Sen. DART. *ACM Sigplan Notices*, 2005.
- [12] K. Jee, G. Portokalidis, V. P. Kemerlis, S. Ghosh, D. I. August, and A. D. Keromytis. A general approach for efficiently accelerating software-based dynamic data flow tracking on commodity hardware. In *In Proc. of the 19th NDSS*, 2012.
- [13] V. P. Kemerlis, G. Portokalidis, K. Jee, and A. D. Keromytis. libdft: practical dynamic data flow tracking for commodity systems. In *Proc. of VEE*, 2012.
- [14] C. Lattner and V. Adve. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *Proceedings of CGO*. IEEE Computer Society, Mar. 2004.
- [15] C. Lattner, A. Lenharth, and V. Adve. Making context-sensitive points-to analysis with heap cloning practical for the real world. In *Proceedings of PLDI*. ACM Request Permissions, June 2007.
- [16] F. Long, S. Sidiropoulos-Douskos, D. Kim, and M. Rinard. Sound input filter generation for integer overflow errors. In *the 41st ACM SIGPLAN-SIGACT Symposium*, New York, New York, USA, 2014.
- [17] D. Molnar, X. C. Li, and D. A. Wagner. Dynamic test generation to find integer bugs in x86 binary linux programs. pages 67–82, 2009.
- [18] S. Moore. thinkmoore/llvm-deps. <https://github.com/thinkmoore/llvm-deps>. (Visited on 06/07/2014).
- [19] T. Wang, T. Wei, Z. Lin, and W. Zou. Intscope: Automatically detecting integer overflow vulnerability in x86 binary using symbolic execution. In *NDSS*, 2009.
- [20] X. Wang. Fast integer overflow detection. <http://kqueue.org/blog/2012/03/16/fast-integer-overflow-detection/>.
- [21] X. Wang, H. Chen, Z. Jia, N. Zeldovich, and M. F. Kaashoek. Improving integer security for systems with KINT. In *Proceedings of OSDI*. USENIX Association, Oct. 2012.
- [22] X. Wang, N. Zeldovich, M. F. Kaashoek, and A. Solar-Lezama. Towards optimization-safe systems. In *the Twenty-Fourth ACM Symposium*, pages 260–275, New York, New York, USA, 2013. ACM Press.
- [23] C. Zhang, T. Wang, T. Wei, Y. Chen, and W. Zou. IntPatch: Automatically fix integer-overflow-to-buffer-overflow vulnerability at compile-time. pages 71–86, 2010.