

CSSV: Towards a Realistic Tool for Statically Detecting All Buffer Overflows in C

Nurit Dor*
Tel-Aviv University
nurr@post.tau.ac.il

Michael Rodeh
IBM Research Lab in Haifa
rodeh@il.ibm.com

Mooly Sagiv
Tel-Aviv University
msagiv@post.tau.ac.il

ABSTRACT

Erroneous string manipulations are a major source of software defects in C programs yielding vulnerabilities which are exploited by software viruses. We present **C String Static Verifier (CSSV)**, a tool that statically uncovers *all* string manipulation errors. Being a conservative tool, it reports all such errors at the expense of sometimes generating *false alarms*. Fortunately, only a small number of false alarms are reported, thereby proving that statically reducing software vulnerability is achievable. CSSV handles large programs by analyzing each procedure separately. To this end procedure *contracts* are allowed which are verified by the tool.

We implemented a CSSV prototype and used it to verify the absence of errors in real code from EADS Airbus. When applied to another commonly used string intensive application, CSSV uncovered real bugs with very few false alarms.

Categories and Subject Descriptors

D.2.4 [Software Engineering]: Software/Program Verification—*Assertion checkers, Reliability, Validation*; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs—*Assertions, Pre- and post-conditions*; F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages—*Operational semantics, Program analysis*

General Terms

Algorithms, Reliability, Experimentation, Security, Languages, Verification

Keywords

Error detection, abstract interpretation, static analysis, buffer overflow, contracts

*Partially supported by a grant from the Ministry of Science, Israel and by the RTD project IST-1999-20527 “DAEDALUS” of the european FP5 programme.

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.

PLDI’03, June 9–11, 2003, San Diego, California, USA.
Copyright 2003 ACM 1-58113-662-5/03/0006 ...\$5.00.

1. INTRODUCTION

String manipulation errors are a common source of software defects and lead to many security vulnerabilities. CERT advisories report on many security holes that result from *buffer overflow*, i.e., updates beyond the bounds of a buffer [37]. Furthermore, 60% of the UNIX failures reported in the 1995 FUZZ study [28] are due to runtime string manipulation errors, such as buffer overflow, access beyond the bounds of a string and misuse of the null-termination byte.

Our goal is to perform static analysis that detects *all* string runtime errors with just a few *false alarms*. A false alarm is a reported error that can never occur at runtime. This goal is ambitious. Existing methods either: (i) miss errors (e.g., LCLint [23], Eau Claire [4], and [37]); (ii) yield many false alarms (e.g., [23, 37]); or (iii) cannot handle complex aspects of C, such as multi-level pointers and structures (e.g., [13, 35]). Moreover, the cost of static analysis is considered prohibitive when it comes to large programs.

This paper presents **C String Static Verifier (CSSV)** for short — a tool that demonstrates that uncovering all string problems in C is achievable. CSSV is capable of analyzing realistic procedures and produces rather precise results. Being a conservative static-analysis tool it can never miss a runtime string error. It therefore guarantees the absence of *all* errors at the expense of sometimes generating false alarms.

For every procedure, CSSV allows the programmer to provide a *contract* including (i) a precondition, (ii) a postcondition, and (iii) the potential side-effects of the procedure. Contracts may refer to normal C expressions (including pointers) and can also refer to properties, (such as the number of allocated bytes) that are defined by *instrumented concrete semantics*.

1.1 Analysis of String Errors: CSSV

Fig. 1 shows how CSSV operates. Each procedure is analyzed separately. In the first phase, a source-to-source semantic-preserving transformation is applied to the analyzed procedure P . This transformation exposes the behavior of the procedures invoked by P by essentially inlining their contracts. The generated program yields a runtime error when a contract is violated. In addition, the inliner normalizes the C code to only include statements in a C subset called *CoreC* [38] which simplifies the task of implementing CSSV.

In the second phase, CSSV analyzes pointer interactions. Conducting pointer analysis in a language like C is a non-trivial task. Moreover, it is difficult for programmers to de-




Figure 1: High-level structure of CSSV.

fine contracts in terms of pointer behavior. Fortunately, several flow-insensitive algorithms have been shown to run on whole applications of considerable size, e.g., [8, 18]. Therefore, CSSV does not require contract information regarding pointers. Instead, CSSV applies a whole-program flow-insensitive pointer analysis to detect statically which pointers may point-to the same base address. CSSV then applies an algorithm that extracts *procedural points-to information* for the analyzed procedure P . Our algorithm benefits from the fact that memory locations inaccessible from visible variables of P cannot affect the postcondition of P . In many cases this allows subsequent analyses to perform strong updates [3] when analyzing the procedure’s body. We also compute certain *must-aliases* to improve the precision of the global flow-insensitive pointer analysis.

In the third phase, the procedure code and points-to information are fed into the *C2IP* transformer. *C2IP* generates a procedure that manipulates integers. *C2IP* guarantees that if there is a runtime string manipulation error in a procedure invocation then either (i) the procedure’s precondition did not hold on this invocation, or (ii) an `assert` statement in the resultant integer program is violated on a corresponding input. In addition, *C2IP* checks pointer assertions if specified in the contracts.

In the fourth phase, the resultant integer program is analyzed using a conservative integer-analysis algorithm to determine all potential violations of `assert` statements. Because the integer and pointer analyses are sound and because contracts are verified both at call sites and at the procedural level, all string errors are reported. In particular, the integer analysis reports an error when the specified postcondition is not guaranteed to hold. For minimizing the number of false alarms, CSSV uses a rather precise integer analysis that represents linear relationships on integer variables. The final result is a list of potential errors. For every error, a counter-example is generated that can assist the programmer in determining if a message is a real error or a false alarm. False alarms may occur due to (i) erroneous or overly weak contracts, (ii) abstractions conducted by *C2IP*, or (iii) imprecision of the pointer or integer analyses.

As opposed to alternative interprocedural program analysis techniques, CSSV’s approach has important advantages: (i) Each potentially recursive procedure can be analyzed separately, exactly once. (ii) The tool is applicable even if the source code is not available in its entirety. (ii) Contracts offer user control in a way similar to “design by contract” [30]. In particular, it enables CSSV to more effectively locate the actual source location in which the error occurs. (iii) Contracts can improve the precision of the analysis by providing information which can be hard to statically infer via an interprocedural analysis. (iv) By using the contracts to analyze procedure calls, CSSV applies a rather precise intraprocedural algorithm to reduce false alarms.

1.2 The Burden of Contracts

Contracts exert additional burden on the programmer. In the case of CSSV, this deficiency is minimized because the pre- and post-conditions need not describe the procedure’s complete behavior. Moreover, unlike tools such as Eau Claire and LCLint, CSSV does not require annotations within the code itself such as loop invariants. Also, unlike these unsound approaches, since CSSV is sound, with any given contract, runtime errors cannot go undetected. Depending on the contracts, errors will be identified when analyzing the body of the procedure or at the procedure invocations. Clearly, when a procedure code is omitted as in the case of library functions, CSSV assumes its contract is correct and cannot verify it.

Pointer information is automatically collected by CSSV, and therefore contracts usually omit information about how pointers are used. In addition, interprocedural modification side-effect analysis algorithms already exist (e.g., [34]). They can generate automatically the modification clause. Therefore, it is always possible to run CSSV with *vacuous contracts* including only the side-effect information and a `true` pre- and post-condition.

This paper presents preliminary algorithms for automatically strengthening the pre- and post-conditions. The effectiveness of these algorithms is measured by comparing the number of false alarms obtained: (i) with the vacuous contracts, (ii) when using automatically derived contracts, and (iii) when using manually provided contracts.

The derivation procedure uses a forward sound integer-analysis algorithm called *ASPost* to compute an Approximation to the Strongest Postcondition of the integer program. Similarly, a backward sound integer-analysis algorithm called *AWPre* is used to compute an Approximation to the Weakest liberal Precondition [11]. The generated postcondition (precondition) is not necessarily the strongest (weakest) because information is lost during the static integer analysis. Both *ASPost* and *AWPre* yield integer conditions. Therefore, the process can be repeated iteratively by running the derivation process given the generated integer con and further restricting the existing postcondition (precondition).

We also present a conservative method that uses the procedural points-to information to convert an integer expression for postcondition (precondition) into a C expression that can be used to strengthen the initial contract.

1.3 Main Results

The contributions of this paper can be summarized in the following way:

- A conservative static-analysis algorithm for detecting string runtime errors is presented. The algorithm reduces the problem of checking string manipulation to that of checking integer manipulations—a problem for

which well-known solutions exist. In comparison to our previous algorithm, presented in [13], it handles the *full* spectrum of C language constructs, including dynamically allocated structures, multi-level arrays, multi-level pointers, function pointers, and casting. In addition, this algorithm is an order of magnitude better in its asymptotic time and space requirements.

- An algorithm that computes procedural pointer information from a given whole-program flow-insensitive pointer information is presented. The goal is to reduce the number of false alarms when analyzing well-behaved programs. Specifically, the algorithm can infer that a formal parameter points to a single location throughout the procedure and that a local variable must be aliased to a formal parameter. Obtaining procedural pointer information is a fundamental problem that clients of whole-program flow-insensitive pointer information face. The problem stems from the fact that a formal parameter may point to different abstract locations. Thus, a naive implementation will perform weak updates which may lead to many false alarms. Hence, we believe that our algorithm can be used to improve the precision of other clients of whole-program flow-insensitive algorithms, such as slicing tools and program optimizers (e.g., [15, 25]).
- Preliminary program-analysis algorithms for strengthening pre- and post-conditions are presented. The algorithms reduce the burden on the programmer. They analyze the input procedure using existing (potentially vacuous) contracts and yield a new, more restrictive, contract for this procedure.
- We have implemented CSSV using the AST-Toolkit [32], CoreC, the Golf pointer analysis [8, 9], and the polyhedra integer analysis of [6] from [19]. We have applied the implementation to real-life programs. CSSV verified an intricate string library from EADS Airbus yielding only 6 false alarms. In the application `fixwrites`, part of `web2c`, CSSV uncovered 8 errors with 2 false alarms. Finally, we implemented the derivation algorithms and applied them to automatically generate pre- and post-conditions. The results show that in some cases this brought about contracts equivalent to the manually specified ones.

1.4 Outline of the Rest of this Paper

The rest of the paper is organized as follows: Section 2 introduces CoreC, a contract language, a running example, and our instrumented concrete sem Section 3 describes CSSV. Section 4 describes the contract derivation algorithms. Section 5 describes the prototype implementation and the experimental results. Section 6 discusses related work.

2. BACKGROUND

2.1 CoreC

CoreC is a subset of C with the following restrictions: (i) Control-flow statements are either `if`, `goto`, `break`, or `continue`; (ii) expressions are side-effect free and cannot be nested; (iii) all assignments are statements; (iv) declarations do not have initializations; (v) address-of formal variables is not allowed. An algorithm for transforming C programs to

Attribute	Intended Meaning
<code>exp.base</code>	The base address of <code>exp</code>
<code>exp.offset</code>	The offset of <code>exp</code> , i.e., <code>exp - exp.base</code>
<code>exp.is_nullt</code>	Is <code>exp</code> pointing to a null-terminated string?
<code>exp.strlen</code>	The length of the string pointed-to by <code>exp</code>
<code>exp.alloc</code>	The number of bytes allocated from <code>exp</code>

Table 1: Attributes in the contract language.




Figure 2: Graphical representation of the contract-language attributes.

CoreC is presented in [38]. Given a C program, it generates an equivalent CoreC program by adding new temporaries. CSSV is defined and implemented for CoreC. In the rest of this paper, CoreC is used instead of C.

2.2 Contracts

Contracts are used to describe expected inputs, side-effects, and expected output of functions. In this paper, we write contracts in the style of Larch [24]. Our implementation actually supports a more general executable language similar to [29], which can include loops. Contracts are specified in the `.h` file. Every prototype declaration of a function f has the form:

```

<type> f(⋯)  requires  ⟨e⟩
              modifies ⟨e⟩, ⟨e⟩, ⋯, ⟨e⟩
              ensures  ⟨e⟩;

```

defining the precondition required to hold whenever f is invoked, the side-effects of the function f , i.e., the objects that may be modified during invocations of f , and the post-condition that is guaranteed to hold on the modified objects. Here, $\langle e \rangle$ is a C expression, without function calls, over global variables and the formal parameters of f . We allow *attributes* of the form defined in Table 1 and displayed in Fig. 2. A designated variable `return_value` denotes the return value of f . The special syntax $\llbracket \langle e \rangle \rrbracket_{pre}$ denotes the value of $\langle e \rangle$ when f is invoked. Although not required, the contract mechanism enables specifying pointer values. In addition a shorthand expression `is_within_bounds(arg)` is allowed to indicate that `arg` points within the bounds of a buffer.

2.3 Running Example

The CoreC version of the function `RTC_Si_SkipLine` from EADS Airbus (`SkipLine` for short) is shown in Fig. 3. `SkipLine` inserts `NbLine` newline characters starting at the location pointed-to by `*PtrEndText`, appends a null-termination character and sets `*PtrEndText` to point to the end of the string.

A contract for `SkipLine` is shown in Fig. 4. The precondition demands that upon entry: `*PtrEndText` points to within the bounds of a buffer; the allocation size from the

```

void SkipLine(int NbLine, char** PtrEndText) {
  int indice;
  char* PtrEndLoc;
[1] indice=0;
[2] begin_loop:
[3] if (indice>=NbLine) goto end_loop;
[4] PtrEndLoc = *PtrEndText;
[5] *PtrEndLoc = '\n';
[6] *PtrEndText = PtrEndLoc + 1;
[7] indice = indice + 1;
[8] goto begin_loop;
[9] end_loop:
[10] PtrEndLoc = *PtrEndText
[11] *PtrEndLoc = '\0'; }

void main() {
  char buf[SIZE];
  char *r, *s;
[1] r = buf;
[2] SkipLine(1,&r);
[3] fgets(r,SIZE-1,stdin);
[4] s = r + strlen(r);
[5] SkipLine(1,&s); }

```

Figure 3: SkipLine, a string manipulation function from EADS Airbus with a toy main function.

location `*PtrEndText` is greater than `NbLine`; and, `NbLine` is at greater or equal to zero. The function may modify the `*PtrEndText` pointer and the buffer pointed-to by `*PtrEndText`. The postcondition indicates that `*PtrEndText` points to a null-terminated string of length zero, and its value is advanced by `NbLine` bytes.

Due to multi-level pointer indirections, destructive updates, and pointer arithmetic, it is rather challenging to verify the absence of errors in this function. CSSV is able to statically verify the absence of string errors in this function, without reporting any false alarm.

The toy `main` procedure, shown in Fig. 3, calls `SkipLine` to insert a newline character, reads input from the standard input, and concatenates an additional newline by calling `SkipLine` again. This procedure has an off-by-one error. In the case of a user input of length `SIZE-1`, `buf` is full and there is no space for the additional newline. CSSV detects this error in `main` without reporting any false alarm.

There is a strong correlation between the provided set of contracts and the messages reported. However, errors do not go undetected. For example, omitting `NbLine >= 0` from the precondition of `SkipLine` yields an error message during the analysis of the procedure. The message indicates that the postcondition

$$*PtrEndText == [*PtrEndText]_{pre} + NbLine$$

may not hold. Interestingly, the counter-example produced by CSSV for this message shows that this postcondition does not hold when the value of `NbLine` is negative.

Providing a precondition which is stronger than the weakest precondition can yield error messages on a procedure invocation. For example, requiring in the precondition of

```

void SkipLine(int NbLine, char** PtrEndText)
  requires is_within_bounds(*PtrEndText) &&
           *PtrEndText.alloc > NbLine && NbLine >= 0
  modifies *PtrEndText
           *PtrEndText.is_nullt *PtrEndText.strlen
  ensures *PtrEndText.is_nullt &&
           *PtrEndText.strlen == 0 &&
           *PtrEndText == [*PtrEndText]_{pre} + NbLine ;

```

Figure 4: A contract for SkipLine.

`SkipLine` that `*PtrEndText` points-to a null-terminated string will cause an error message regarding the call to `SkipLine` at line [2] of `main`.

2.4 Instrumented Concrete Semantics

The C programming language does not define semantics for C programs. In the ANSI-C standard there is an informal notion of defined and undefined behaviors. However, the exact behavior can change, and often does, from one implementation of a compiler to another. Due to the following features of the language, it is not trivial to define semantics for C:

- **Address-of** operation enables changing the value of a variable without assigning to the variable. It also permits pointers to invisible variables.
- **Allocation** library functions (e.g., `malloc`) provide an unformatted contiguous memory locations, while from the logic point of view there is a “hierarchy” of objects where one object may contain objects of different types. Moreover, objects are type-less, thus providing flexibility, and allowing accesses to a location according to different types. Therefore, it is difficult to define and check the legitimacy of an access.
- **Pointer arithmetic** is frequently used and has a defined result. However, checking its validity is impossible without additional instrumented information.
- **Cast** operation exposes the internal memory layout e.g., by allowing casting from integer to pointer type.

In this section, we sketch an instrumented operational semantics for C that verifies the absence of out-of-bound violations while allowing pointer arithmetic, destructive updates and casting. The general idea is to define a non-standard low-level semantics that explicitly represents the base address of every memory location and the allocated size starting from the base address. This semantics is rigorous. It forbids programs with undefined ANSI-C behavior but it also checks additional requirements reflecting good programming styles such as dereferences beyond the null-termination byte. This semantics provides the foundation of CSSV’s abstract interpretation, i.e., the abstract interpretation conservatively represents the states of this semantics. In addition, CSSV statically verifies the absence of string errors by conservatively checking the preconditions of this semantics. The reader is referred to [12] for a discussion on this semantics and refinements for checking the validity of




Figure 5: A concrete state arising at entry to `SkipLine` invoked by the second call from `main`. For clarity, the allocation size of `buf` is only shown symbolically.

accesses to general arrays and multi-level structures. There, the soundness of CSSV is proved with respects to the operational semantics.

DEFINITION 2.1. A **concrete state at a procedure P** is a tuple: $state^{\sharp} = (\mathcal{L}^{\sharp}, \mathcal{BA}^{\sharp}, aSize^{\sharp}, loc^{\sharp}, st^{\sharp}, numBytes^{\sharp}, base^{\sharp})$ where:

- \mathcal{L}^{\sharp} is a finite set of all static, stack, and dynamically allocated locations.
- $\mathcal{BA}^{\sharp} \subseteq \mathcal{L}^{\sharp}$ is the set of base addresses in \mathcal{L}^{\sharp} .
- $aSize^{\sharp}: \mathcal{BA}^{\sharp} \rightarrow \mathbb{N}$ defines the allocation size in bytes of the memory region starting at a base address.
- $loc^{\sharp}: visvarp \rightarrow \mathcal{BA}^{\sharp}$ maps visible variables into their assigned global or stack locations (which is always a base address).
- $st^{\sharp}: \mathcal{L}^{\sharp} \rightarrow val$ defines the memory content, where $val = \{uninit, undefined\} \cup primitive \cup \mathcal{L}^{\sharp}$ is the set of possible values. The value `uninit` represents uninitialized values; `undefined` represents results from illegal memory access; `primitive` refers to the set of C primitive type (`char`, `int`, etc.) values.
- $numBytes^{\sharp}: \mathcal{L}^{\sharp} \rightarrow \mathbb{N}$ defines for each location the number of bytes of the value stored starting at the location.
- $base^{\sharp}: \mathcal{L}^{\sharp} \rightarrow \mathcal{BA}^{\sharp}$ maps every location to its base address.

A concrete state that arises at entry to `SkipLine` when invoked by the second call in `main` is shown in Fig. 5. We draw contiguous memory locations as boxes and display their allocation sizes underneath the boxes. Here, we assume that integers and pointers are four bytes long and a character is one byte long. We draw a variable v above a box whose base address is $loc^{\sharp}(v)$. The value inside each box shows the corresponding store content. Pointer values are drawn as edges.

Intuitively, the state keeps track of the set of allocated locations (\mathcal{L}^{\sharp}). The origin location of each memory region that is guaranteed to be contiguous is in \mathcal{BA}^{\sharp} . In order to handle destructive update to a variable via the address-of operation, loc^{\sharp} represents the address of variables, and st^{\sharp}

maps locations into their values. For example, the pointer to `s` is described as a pointer to a location which is $loc^{\sharp}(s)$. Our concrete example contains, among others, the following interesting mappings:

$$\begin{aligned} st^{\sharp}(loc^{\sharp}(PtrEndText)) &= loc^{\sharp}(s) \\ numBytes^{\sharp}(loc^{\sharp}(PtrEndText)) &= 4 \\ st^{\sharp}(loc^{\sharp}(buf) + 1) &= 'h' \\ numBytes^{\sharp}(loc^{\sharp}(buf) + 1) &= 1 \end{aligned}$$

indicating that `PtrEndText` points-to the stack location of `s` which is a four-byte value, and that the second byte of `buf` contains the character 'h'.

The association of the number of bytes with locations enables us to handle cases where a location is accessed through different types. Specifically, writing a location as one type and later reading it as a different size type results in the *undefined* value.

L-value and *R-value* of C expressions can be defined by straightforward structural induction. In particular for a variable v , we define the L-value and the R-value, denoted as lv_v^{\sharp} and rv_v^{\sharp} , respectively, as follows:

$$\begin{aligned} lv_v^{\sharp} &\stackrel{\text{def}}{=} loc^{\sharp}(v) \\ rv_v^{\sharp} &\stackrel{\text{def}}{=} st^{\sharp}(lv_v^{\sharp}) = st^{\sharp}(loc^{\sharp}(v)) \end{aligned}$$

We define a function, $index^{\sharp}$, to reason about the displacement of a location from its base. Formally,

$$\begin{aligned} index^{\sharp}: \mathcal{L}^{\sharp} &\rightarrow \mathbb{N} \\ index^{\sharp}(l^{\sharp}) &\stackrel{\text{def}}{=} l^{\sharp} - base^{\sharp}(l^{\sharp}) \end{aligned}$$

With the additional information of $aSize^{\sharp}$ and $base^{\sharp}$ the R-value of an attribute `p.offset` is $index^{\sharp}(rv_p^{\sharp})$. In addition, the use of the instrumented mappings allows the semantics to validate that pointer arithmetic and dereferences are within bounds. For example, a pointer expression `p + i` is within the bounds of the buffer pointed-to by `p` when the following condition holds:

$$0 \leq index^{\sharp}(rv_p^{\sharp}) + rv_i^{\sharp} \leq aSize^{\sharp}(base^{\sharp}(rv_p^{\sharp}))$$

We follow [20, pp.205] and check that the result of pointer arithmetic is either before or at the first location beyond the upper bound.

3. CSSV

CSSV analyzes each procedure separately. We refer to the analyzed procedure as P . CSSV checks for three kinds of errors: (i) ANSI-C violations related to strings, such an access out of bounds. (ii) Violations of pre- and post-conditions of procedures as required by the provided contracts. When a procedure is invoked, the callee's precondition is checked. At the end of P , the postcondition of P is checked. (iii) Our analysis checks certain *cleanliness* conditions that correspond to good programming style. In particular, it validates that all accesses are before the null-termination byte, if it exists.

3.1 Technical Overview

Pointers and integers interact in a non-trivial way, especially in the C programming language. For example, it is non-trivial to check the safety of the expression

```
*PtrEndText = '\n'
```

in line [5] of `SkipLine`, i.e., that the pointer `*PtrEndText` is within bounds. CSSV infers the relationships between the offset of `*PtrEndText`, the allocation size of its base address, and the integer variables `indice` and `NbLine` needed to verify the safety of this destructive update. As we shall see, our algorithm statically verifies such inequalities by combining a pointer-analysis algorithm that detects pointers to the same base address, with an integer-analysis algorithm that detects *offset* relationships among pointers. The offset of a pointer is the index of the location it points to. Of course, in contrast to the concrete semantics, the abstract semantics summarizes many concrete locations by a single abstract location. It also maintains the potential points-to relationships between these addresses.

CSSV applies a whole program flow-insensitive pointer analysis to detect statically which pointers may point to the same base address. In particular, for every function, it provides a summary of all of its calling contexts. In principle, a conservative analysis can utilize this information and analyze a function with all possible calling contexts. However, this can yield many false alarms. For example, the whole-program analysis of `SkipLine` yields that `PtrEndText` may point to either `s` or `r`. Conservatively analyzing the function’s body with the two calling contexts, requires treating updates to integer properties such as the offset of `*PtrEndText` as weak updates. Therefore, the analysis will fail to show that the postcondition holds. As a result, a false alarm will be issued. CSSV avoids this false alarm by performing strong updates in certain cases. The main idea is to precompute *procedural points-to information* that guarantees that strong updates to the offset of `*PtrEndText` can be performed. In general, it guarantees that in well-behaved programs direct updates through the formal parameters can be interpreted as strong updates.

The procedural points-to information is used by C2IP to generate an integer program. Integer constraint variables summarize the semantic properties (e.g., allocated size) of the represented locations. Finally, a conservative integer analysis determines potential values of the semantic properties and verifies the constraints upon them.

The rest of this section is organized as follows: Section 3.2 describes the procedure that inlines contracts in P . Section 3.3 formalizes the procedural points-to information for P . Section 3.4 describes the C2IP transformation applied to P . Section 3.5 sketches the integer-analysis algorithm.

3.2 Exposing the Behavior of Procedures

The first step of CSSV takes as input the C program and the provided set of contracts, and generates a new C procedure $inline(P)$ by exposing the contracts of P and of the invoked procedures. Since $inline(P)$ contains `assert` statements that verify contracts, the behavior of $inline(P)$ differs from the behavior of P on inputs which violate the contracts.

We extend C as follows:

- The construct `assume($\langle e \rangle$)` that indicates that $\langle e \rangle$ holds after this statement, i.e., if $\langle e \rangle$ does not hold the execution is aborted without any message. It is used to reflect commitments of other procedures.
- Additional temporary variables named “ $\langle e \rangle$ ” used to store the value of a subexpression $\llbracket \langle e \rangle \rrbracket_{pre}$ at the procedure entry.

- The contract-language attributes which have a well-defined meaning in our instrumented concrete semantics.

Most of the C statements remain intact. Table 2 shows the scheme for translating the affected statements. Procedure entry is encountered before the first executable statement. In this case, the additional variables are initialized and the precondition of P is assumed to hold. The designated variable `return_value` is set at every return statement. At every exit point (including `return`), the postcondition of P is verified.

On a call to g we verify that g ’s precondition holds and assume that the postcondition holds. The original call to g is in the emitted code. This is essential for $inline(P)$ to have the same interpretation as P .

3.3 Pointer Analysis

The second step of CSSV computes an abstraction of all potential pointer relationships between locations in concrete states that may occur during the execution of P . However, only locations that can be accessed during the execution of P are of interest. Therefore, we define the notion of *reachable* locations.

DEFINITION 3.1. *In a concrete state, a location l^b is **reachable** if there exists a visible variable whose content can (indirectly) include l^b (i.e., there is an expression whose L-value is l^b).*

Computing procedural pointer information allows us to infer the pointer relationships among reachable locations of P . Moreover, the procedural pointer information aims at representing the location a formal points to at the procedure entry as a *single* location. This section describes the abstract state representing pointer relationships and an algorithm to compute this state.

3.3.1 Procedural Points-to Information

We formalize an abstract state that regards pointer relationships among reachable locations of P as follows:

DEFINITION 3.2. *A **procedural abstract points-to state** of P (PPT) is a quadruple $state_P = (\mathcal{BA}_P, loc_P, pt_P, sm_P)$ where:*

- \mathcal{BA}_P is a set of abstract locations that represent all reachable concrete base addresses.
- $loc_P: visvar_P \rightarrow 2^{\mathcal{BA}_P}$ maps variables into set of abstract locations representing the variable’s global or stack locations.
- $pt_P: \mathcal{BA}_P \rightarrow 2^{\mathcal{BA}_P}$ abstract the possible pointers. A concrete pointer is represented by a pt_P relationship between the abstract locations representing the base addresses of the source and target locations of the pointer.
- $sm_P: \mathcal{BA}_P \rightarrow \{1, \infty\}$ is an abstract count of the number of concrete base addresses represented by an abstract location, i.e., $sm(ba) = \infty$ when ba may represent more than one base address in a given concrete store, and 1 when it is guaranteed to represent at most one base address. An abstract location having $sm = \infty$ is a **summary abstract location**. Summary abstract locations can be used to represent unbounded sets of base addresses.

Event	Emitted Code
entry of $P(f_1, f_2, \dots, f_n)$	“ $\langle e_i \rangle$ ” = $\langle e_i \rangle$; for every $\lceil \langle e_i \rangle \rceil_{pre}$ in $post[P]$ <code>assume(pre[P](f₁, f₂, ..., f_n));</code>
return $\langle e \rangle$	<code>return_value_P = $\langle e \rangle$;</code>
exit P	<code>assert(post[P](f₁, f₂, ..., f_n));</code>
$\langle e \rangle = g(a_1, a_2, \dots, a_m)$	{ “ $\langle e_i \rangle$ ” = $\langle e_i \rangle$; for every $\lceil \langle e_i \rangle \rceil_{pre}$ in $post[g]$ <code>assert(pre[g](a₁, a₂, ..., a_m));</code> <code>return_value_g = g(a₁, a₂, ..., a_m);</code> <code>assume(post[g](a₁, a₂, ..., a_m));</code> $\langle e \rangle = \text{return_value}_g$; }

Table 2: The emitted C code for affected statements. The notation $pre[x](e_1, e_2, \dots, e_m)$ stands for the precondition of procedure x where formal f_i is replaced with the expression e_i . The expression $post[x]$ is obtained in a similar way, however each of the $\lceil \langle e_i \rangle \rceil_{pre}$ expression is replaced with the variable “ $\langle e_i \rangle$ ”. $return_value_x$ is a designated variable representing the return value in the postcondition of procedure x .

We say that a PPT $(\mathcal{BA}_P, loc_P, pt_P, sm_P)$ is a *sound approximation of a concrete state* $(\mathcal{L}^\natural, \mathcal{BA}^\natural, aSize^\natural, loc^\natural, st^\natural, numBytes^\natural, base^\natural)$ in a procedure P if there exists a function $\alpha: \mathcal{BA}^\natural \rightarrow \mathcal{BA}_P$ satisfying the following requirements:

Base For all reachable $b^\natural \in \mathcal{BA}^\natural$: $\alpha(b^\natural) \in \mathcal{BA}_P$.

Stack For all $v \in visvar_P$: $\alpha(loc^\natural(v)) \in loc_P(v)$.

Pointer For all $l_1^\natural, l_2^\natural \in \mathcal{L}^\natural$ s.t., l_1^\natural and l_2^\natural are reachable, and satisfying $st^\natural(l_1^\natural) = l_2^\natural$: $\alpha(base^\natural(l_2^\natural)) \in pt_P(\alpha(base^\natural(l_1^\natural)))$.

Summary For all $b \in \mathcal{BA}_P$, s.t., $sm_P(b) = 1$, and $b_1^\natural, b_2^\natural \in \mathcal{BA}^\natural$ having $\alpha(b_1^\natural) = \alpha(b_2^\natural) = b$: $b_1^\natural = b_2^\natural$.

DEFINITION 3.3. A $state_P$ is a **sound approximation of P** if it is a sound approximation of all the concrete states that may arise during the execution of P .

L-values and R-values are generalized to return sets of abstract locations. In particular for a visible pointer variable q :

$$lw_q \stackrel{\text{def}}{=} loc_P(q)$$

$$rv_q \stackrel{\text{def}}{=} \bigcup_{l \in lw_q} pt_P(l) = \bigcup_{l \in loc_P(q)} pt_P(l)$$

3.3.2 Constructing Procedural Information

CSSV computes a sound approximation $state_P$ in two stages. First, a whole-program analysis is applied to compute a global abstract points-to state of the whole program $Gstate = (\mathcal{BA}, loc, pt, sm)$ where:

- \mathcal{BA} includes all abstract locations.
- $loc: var \rightarrow 2^{\mathcal{BA}}$.
- $pt: \mathcal{BA} \rightarrow 2^{\mathcal{BA}}$.
- $sm: \mathcal{BA} \rightarrow \{1, \infty\}$.

This global state is guaranteed to be a sound approximation of all procedures. Second, this global state is used to construct a sound approximation for P . It is possible to construct different sound PPTs for P with different abstract locations and points-to relationships. We decided to bias towards precise representation of formal parameters, with the




Figure 6: The whole-program points-to information for the running example (a), and the PPT for SkipLine (b).

intention of conducting strong updates on their properties in many cases.

Fig. 6 demonstrates the process. Fig. 6(a) shows the whole-program points-to information of our running example. Boxes represent abstract locations. When possible we denote abstract locations as either L-value (e.g., lw_s) or R-value of a unique pointer variable. Otherwise, we provide an arbitrary name (e.g., N). Edges represent the pt relationship. There are no summary abstract locations in this example. The final result of computing the PPT for SkipLine is shown in Fig. 6(b). A new abstract location $rv_{PtrEndText}$ represents the (unique) concrete location which holds the value of $*PtrEndText$.

Given a global abstract pointer state of the whole program $Gstate = (\mathcal{BA}, loc, pt, sm)$, let us construct a PPT for P $state_P = (\mathcal{BA}_P, loc_P, pt_P, sm_P)$. The mapping loc_P is computed by projecting loc to the visible variables of P . Similarly, \mathcal{BA}_P and sm_P are computed by including abstract locations reachable from visible variables of P .

An initial value for pt_P is obtained by projection. In our running example, this yields the same state as the global points-to information shown in Fig. 6(a) without the lw_{buf} abstract location. We aim at a potentially more precise representation. A conservative algorithm which checks whether it is sound to merge the nodes l_1, l_2, \dots, l_m pointed-to by a formal f without creating a new summary node, is presented in Fig. 7. This algorithm checks that for every concrete store at most one concrete location is represented by rv_f (the set of abstract locations pointed-to by a formal parameter f). The correctness of the algorithm is established in [12].

Whenever possible, merging is done by (i) replacing the

```

Boolean parameterizable( PPT  $state_P$ , formal  $f$  )
{
  let  $state_P = (\mathcal{B}A_P, loc_P, pt_P, sm_P)$ 
  let  $l_f = loc_P(f)$  // the L-value of  $f$ 
  if  $sm_P(l_f) = \infty$  return false
  let  $\{l_1, l_2, \dots, l_m\} = pt(l_f)$  // the R-values of  $f$ 
  for  $i = 1$  to  $m$ 
    if  $sm_P(l_i) = \infty$  return false
    remove from  $pt$  edges from  $l_f$  to  $l_j$  where  $j \neq i$ ,
    and let  $pt'$  be the resultant points-to map.
    if exists a reachable node  $l_j, j \neq i$  in  $pt'$  then
      return false
  // at most one of the concrete locations pointed-to by  $f$ 
  // is reachable in a concrete state represented by  $state_P$ 
  return true;
}

```

Figure 7: Algorithm to conservatively check that at most one concrete location is represented by the set pointed-to by a formal parameter f .

abstract locations l_1, l_2, \dots, l_m by a single non-summary abstract location rv_f . (ii) setting $pt(rv_f)$ to $\bigcup_{i=1}^m pt(l_i)$. This may improve the precision of destructive updates through f , but may decrease the precision of other updates.

3.4 C2IP

The C2IP transformation takes the $inline(P)$ procedure with its PPT as input, and produces an integer program (IP for short) as output. The generated IP tracks the string and integer manipulations of P and of the invoked procedures. The IP is nondeterministic, reflecting the fact that not all values are known. The symbol `unknown` stands for an undetermined value. In particular, we use the following expressions:

$x := \text{unknown}$; Assigns any value to x .
 if (`unknown`) Either the true or the false branch can be taken.

The semantics of the `assume` construct in the integer program is to restrict the behavior of nondeterministic programs. Finally, for clarity, we use mathematical constructs in the IP.

The IP includes *constraint variables* used to denote semantic properties of interest such as offsets. C2IP generates statements which assign new values to constraint variables, reflecting the changes in the semantic properties. `Assert` statements over the constraint variables are generated. They check for the safety of basic C expressions and for verifying contracts. In addition, C2IP can validate pointer assertions if specified in the contracts. Due to the flow insensitivity of our pointer analysis, this capability is rather weak in terms of precision. When a precondition may not hold, an error message is reported.

3.4.1 Constraint Variables

For every abstract location, l , C2IP generates the following *constraint variables*:

- $l.val$ to represent potential primitive values stored in the locations represented by l .

C Exp.	Generated IP Condition
*p	$lv_p.offset \geq 0 \wedge ((rv_p.is_nullt \wedge lv_p.offset \leq rv_p.len) \vee (\neg rv_p.is_nullt \wedge lv_p.offset < rv_p.aSize))$
p + i	$0 \leq lv_p.offset + lv_i.val \leq rv_p.aSize$

Table 3: Asserted IP conditions for C expressions.

- $l.offset$ to represent potential offsets of the pointers represented by l , i.e., $l.offset$ conservatively represents $index^h(st^h(l^h))$ for every location l^h represented by l .
- $l.aSize$, $l.is_nullt$ and $l.len$ to describe the allocation size, whether the base address contains a null terminated string, and the length of the string (excluding the null byte) of all locations represented by l .

3.4.2 Translating Expressions

Transforming C expressions involves querying the PPT to obtain the abstract locations a pointer *may* point-to. For the sake of simplicity, in this subsection we assume that every pointer may only point to a single non-summary abstract location. Thus, lv_p (representing the global or stack location of p), and rv_p (representing the location pointed-to by p) are both singletons for every pointer p . In Section 3.4.2.3, general PPTs are considered.

3.4.2.1 Safety Checks.

For every C expression, there is a condition that verifies the validity of the expression. Table 3 lists the generated assert expressions. On every dereference to an address, a check that the address is within bounds is generated. The upper bound is checked depending on whether the buffer is null-terminated. If it is, the dereferenced location is checked to be at or before the null-termination byte. For pointer arithmetic, the generated assert statement checks the requirement that the resultant reference is within or at the upper bound of the buffer. The generated assert resembles the requirement defined in the concrete semantics for pointer arithmetic. This emphasizes that CSSV abstracts the properties needed to statically verify pointer arithmetic.

3.4.2.2 Statements.

C2IP generates statements to reflect semantic changes regarding the properties tracked. The core rules for translating C constructs into IP is shown in Table 4.

On allocation, the resultant pointer always points to a base address. Therefore, its offset is always zero. We set the allocation size of the abstract location that represents the newly allocated location. Destructive updates are separated into two cases: (i) The assignment of the null character, which sets the buffer to a null-terminated string. The length of the string is the location of the first zero byte. C2IP generates a check that all dereferences are before the null-termination byte (if it exists). We can therefore safely assume that when assigning a null-termination byte it is the first one. (ii) In the assignment of a non-zero character, it is checked whether an existing null-termination byte is overwritten.

The generated IP does not contain function calls. Because C2IP transforms the $inline(P)$ procedure, the pre- and post-

C Construct	IP Statements
$p = \text{Alloc}(i);$	$lv_p.offset := 0;$ $rv_p.aSize := lv_i.val;$ $rv_p.is_nullt := false;$
$p = q + i;$	$lv_p.offset := lv_q.offset + lv_i.val;$
$*p = c;$	if $c = 0$ { $rv_p.len := lv_p.offset;$ $rv_p.is_nullt := true;$ } else if $rv_p.is_nullt \wedge lv_p.offset = rv_p.len$ $lv_p.is_nullt := \text{unknown};$
$c = *p;$	if $rv_p.is_nullt \wedge lv_p.offset = rv_p.len$ $lv_c.val := 0;$ else $lv_c.val := \text{unknown};$
$g(a_1, a_2, \dots, a_m);$	$mod[g](a_1, a_2, \dots, a_m);$
$*p == 0$	$rv_p.is_nullt \wedge rv_p.len = lv_p.offset$
$p > q$	$lv_p.offset > lv_q.offset$
$p.alloc$	$rv_p.aSize - lv_p.offset$
$p.offset$	$lv_p.offset$
$p.is_nullt$	$rv_p.is_nullt$
$p.strlen$	$rv_p.len - lv_p.offset$

Table 4: The generated transformation for C statements, conditional expressions and for contract-language attributes. p and q are variables of type pointer to char. i and c are variables of int type. Alloc is a memory allocation routine, e.g., malloc and alloca .

conditions of an invoked procedure g are transformed too. However, the call to a procedure needs to be analyzed conservatively. C2IP converts the call to g with the modification clause of g and substitutes actual for formal parameters. The modification clause is interpreted as assignments of **unknown** to the constraint variables of the abstract locations that represent potentially modified objects.

To increase precision, certain program conditions are interpreted. The second part of Table 4 shows the interpreted conditions. When checking for null-termination, C2IP replaces the condition with a condition over constraint variables that track the existence of a null-character and the length. Pointer comparisons are replaced by expressions over the appropriate **offset** constraint variables.

For convenient use, the contract language allows specifying attributes on pointers instead of on base addresses. For example, $p.alloc$ represents the allocation size starting at the location pointed to by p . The last part of Table 4 lists the transformation of contract’s attributes to constraint variables by referring to the abstract locations pointed to by the specified pointer.

3.4.2.3 Other C Constructs.

In the case that an L-value (R-value) in the abstract points-to state includes more than one abstract location or a summary abstract location, the translation rules of Table 4 need to be changed to guarantee sound results. To reflect the fact that a base address represented by l may or may not be modified, C2IP generates every statement (shown in Table 4) as a nondeterministic assignment, under an **if (unknown)** statement. In addition, the analysis must take into account

all possible values of a pointer, and verify expressions on all possible pointer values. This applies to all generated assert statements and program conditions.

To handle casting and unions, C2IP generates for an assignment to one type of constraint variable assignments of **unknown** to the other constraint variables of the same abstract location. For example, an assignment of an integer to a concrete location represented by abstract location l yields an assignment to $l.val$. In addition, C2IP generates the assignment $l.offset := \text{unknown}$. In particular, a cast to and from pointer type is conservatively handled by an assignment to **unknown**.

The pointer analysis determines which functions may be invoked at a call statement via a function pointer. Then, CSSV generates a non-deterministic statement that selects an arbitrary function call.

It is difficult to write general contracts for the format functions, such as $\text{sprintf}()$ and $\text{printf}()$. Therefore, for the format functions, C2IP generates automatically pre- and post-condition according to the exact calling context. CSSV warns in cases where the format parameter is not a constant.

3.4.2.4 The Complexity of C2IP.

The number of constraint variables in the IP is $O(V)$ where V is the number of variables and allocation sites in the C program. Because a pointer may point to V abstract locations, the translation of a C expression that contains one pointer generates $O(V)$ IP statements. Therefore, the size of the IP is $O(S * V)$, where S is the number of C expressions. This is an order-of-magnitude improvement over the transformation in [13], which generates $O(V^2)$ variables and $O(S * V^2)$ statements.

3.5 Integer Analysis

In the final step, CSSV analyzes the IP and reports potential assert violations. In theory, any sound integer analysis can be used. Because many of the tracked semantic properties are external to the procedure, and sometimes even to the whole application, it is essential to track relationships between constraint variables and not just possible values. Furthermore, many of the conditions to infer involve three and more properties, e.g., the postcondition of **SkipLine** regarding the new offset of ***PtrEndText**.

Given that our goal is to generate as few as possible false messages, we apply the a linear-relation analysis [6, 17] which discovers linear inequalities among numerical variables. This method identifies linear inequalities of the form: $\sum_{i=1}^n c_i x_i + b \geq 0$, where x_i is an integer variable and c_i and b are constants. In our case, x_i are the constraint variables. Upon termination of the integer analysis, the information at every control-flow node conservatively represents the inequalities that are guaranteed to hold whenever the control reaches the respective point. The reader is referred to [6, 17, 13] for information about integer analysis.

3.5.1 Assert Checking

During integer analysis, each **assert** statement is verified. This is done by checking if the asserted integer expression is implied by the linear inequalities that hold at the corresponding control-flow node. If the assertion cannot be verified then a counter-example is generated. The counter-example describes the values of the constraint variables where a string error in the C program may arise.

$ \begin{aligned} rv_{\text{buf}}.aSize &= SIZE \\ rv_{\text{buf}}.len &\geq 1 \\ rv_{\text{buf}}.aSize &\geq rv_{\text{buf}}.len + 1 \\ lv_s.offset &= rv_{\text{buf}}.len \end{aligned} $ <p style="text-align: center;">(a)</p>
<pre> [5] SkipLine(1,&s); require(rv_buf.aSize - lv_s.offset > 1) error: the requirement may be violated when: rv_buf.aSize = rv_buf.len + 1 </pre> <p style="text-align: center;">(b)</p>

Figure 8: A report on the error in line [5] of main. The derived inequalities before execution of line [5] of main (a), and a counter example (b).

Fig. 8 demonstrates how the static integer-analysis algorithm identifies the error in the call to `SkipLine` in line [5] of `main`. The algorithm discovers that the inequalities shown in Fig. 8 (a) hold before the execution of line [5], and that when the equality shown in Fig. 8 (b) holds a violation of `SkipLine`’s precondition occurs.

4. DERIVING CONTRACTS

This section presents integer-analysis algorithms to strengthen pre- and post-conditions. The following process is applied to a procedure P :

1. Compute side-effect information for P .
2. Run the inliner and C2IP with vacuous `true` pre- and post-condition which produces an integer program IP_0 .
3. Run ASPost, a forward integer analysis of [6] on IP_0 which computes a safe approximation of the strongest postcondition. Obtain a new IP program IP_1 by strengthening the postcondition with the set of linear inequalities generated by the integer analysis at the procedure exit.
4. Run AWPPre, a backward integer analysis on IP_1 which computes an approximation to the weakest liberal precondition. Obtain a new IP program IP_2 by strengthening the precondition with a set of linear inequalities generated by the analysis at the procedure entry.
5. *Writeback* — by using the PPT, convert the pre- and the post-conditions of IP_2 to C expressions over the formal parameters and global variables of P .

The derivation process can also start with manually given contracts. For applications with acyclic call graphs, the above process can be automatically applied in a bottom-up fashion, starting with the leaf procedures.

4.1 Integer Analysis

The ASPost algorithm is essentially the algorithm of Section 3.5 without issuing messages. It computes linear inequalities that hold at the exit point. Local variables are eliminated. The resulting inequalities are added to the input postcondition.

To improve the effectiveness of the derivation, the inliner phase is allowed to add designated variables to record values

of properties that may be modified by P . For every potentially modified integer property expressed as a C expression $\langle e \rangle$, the *inline*(P) procedure includes a new variable “ $\langle e \rangle$ ” with an additional C statement

$$\text{assume}(\langle e \rangle == \langle e \rangle);$$

During the writeback process, this variable is replaced by an appropriate $\llbracket \langle e \rangle \rrbracket_{pre}$ expression in the postcondition. In this example, since `*PtrEndText` may be modified, variables are used to record all its properties. In particular, a variable “`*PtrEndText.offset`” records the value of the expression `*PtrEndText.offset` at the entry.

The linear relationships obtained by ASPost when applied to `SkipLine` in the running example with a `true` precondition are:

$$\begin{aligned}
N.is_nullt &= \text{true} \\
N.len &= rv_{\text{PtrEndText}}.offset \\
rv_{\text{PtrEndText}}.offset &\geq \text{“*PtrEndText.offset”} + lv_{\text{NbLine}}.val
\end{aligned}
\tag{1}$$

The existence of a null-termination byte and the new length of the base address points to-by `*PtrEndText` is computed by ASPost precisely. ASPost finds a relationship between the old and new offsets of `*PtrEndText`. However, this relationship is weaker than the manually provided one on which the inequality is an equality. Both ASPost and AWPPre may lose information due to joins of control-flow paths and due to the widening operation.

AWPPre is similar to the forward algorithm in the sense that it uses the same abstract domain and abstract operations. The main difference is the treatment of assignments, which are handled by substitutions.

4.2 Write Back

The pre- and post-conditions generated by AWPPre and ASPost are converted into C expressions over the formal parameters and global variables of P . These expressions are added to the input contracts using logical-and operator.

4.2.1 Obtaining Postconditions

Recall that the integer analysis computes properties of abstract locations. Each such abstract location corresponds to a set of L-value expressions over global and formal variables of P . Consider an abstract location l and assume, for simplicity, that there is a unique expression, say e , whose L-value is l . In this case, every constraint variable in the inequalities that occur in the exit are replaced by substituting e for l . Each occurrence of a designated formal parameter “ $\langle e \rangle$ ” is replaced by $\llbracket \langle e \rangle \rrbracket_{pre}$. Finally, the semantic properties are converted to the contract-language attributes.

For the equations in (1), the writeback algorithm yields:

```

**PtrEndText.is_nullt &&
**PtrEndText.strlen = 0 &&
*PtrEndText.offset >=  $\llbracket \text{*PtrEndText.offset} \rrbracket_{pre} + \text{NbLine}$ 

```

When an abstract location corresponds to a set of L-value expressions, we generate a weaker postcondition using logical-or operator. An alternative would be to ignore some of these expressions, which may lead to false alarms when the procedure is analyzed by CSSV.

4.2.2 Obtaining Preconditions

Generating C expressions for preconditions from the entry inequalities is similar to the process of generating post-conditions. The main difference is that we use logical-and instead of logical-or when multiple expressions correspond to the same abstract location.

5. EMPIRICAL RESULTS

Implementing the CSSV tool is non-trivial because of the complicated aspects of C and program analysis. We have implemented a prototype of CSSV with significant help from the Semantics Based Tools group at Microsoft Research and from Greta Yorsh from Tel-Aviv university. The compiler from C to CoreC is built upon the AST-Toolkit. CSSV uses Golf, a flow-insensitive context-sensitive points-to analysis technique, as the underlying whole program pointer analysis. Golf uses flow edges to represent assignments. Partial must information on pointer aliases is extracted from these edges. Both the integer analysis and the automatic derivation of pre- and post-condition were implemented using the Polyhedra library.

We applied CSSV to procedures from the following: (i) A string-manipulation library from EADS Airbus with a total of 400 lines and 11 procedures. (ii) `fixwrites` — part of `web2c` a converter from TeX, Metafont, and other related WEB programs to C. `fixwrites` consist of 460 lines and eight procedures. We have manually written contracts for the analyzed procedures.

Table 5 describes the benchmark characteristics and the analysis results. The column **LOC** displays the number of source lines in the original source. Column **SLOC** displays the number of source lines after the source-to-source transformation. The **Contract** column investigates the difficulty of manually providing a contract. We use the characters ‘S’, ‘B’ and ‘I’ as follows: (S) for simple specification, such as `string` and `is_within_bounds`, (B) for specifying the boundaries of buffers, and (I) for other integer relations. There was no need to provide pointer specification for the analyzed code.

Columns **IP Vars** and **IP size** report the number of variables and statements in the integer program produced by C2IP. Columns **CPU** and **Space** display the running time and total allocated space of CSSV. All experiments were conducted on a 900 MHz Intel Pentium-III CPU with 512MB of memory, running Windows 2000.

The **Msg** columns classify the messages reported by CSSV. Messages are classified as errors for cases where there is an input to the application on which the error occurs. The errors detected are due to unsafe calls to library functions, such as `strcpy()`, unsafe assumptions that an input contains a specific character, or unsafe pointer arithmetic.

There are six false messages on Airbus’s code. The program destructively assigns a non-zero character to a certain place in a buffer. CSSV fails to infer that this character is non zero. The function `skip_balanced` safely assumes that the input parameter contains a balanced number of parentheses. This is verified by the `whole` function which is called prior to `skip_balanced`. This example demonstrates that in some cases it is hard to separate safety from correctness. To show that this function is safe, we need to verify correctness, i.e., that the implementation correctly checks that the input string contains a balanced number of parentheses.

Fortunately, in most of the analyzed examples, this is not the case, i.e., the safety does not depend on correctness.

The **Deriving** columns provide information about the effectiveness of the AWPPre and ASPPost algorithms. It is not trivial to measure the result in terms of precision. A new contract for a function P can change the result of the analysis of P itself and of procedures invoking P . We provide a simple measurement that is independent of the calling context. We run ASPPost to generate a postcondition, AWPPre to generate a precondition, and then run CSSV. Columns **CPU** and **Space** display the running time and total allocated space of both ASPPost and AWPPre. Column **Vacuous** displays the number of false-alarm messages reported by CSSV when a vacuous contract for the analyzed procedure is provided. Column **Auto** displays the number of false alarms reported by CSSV when using the automatically derived contracts. On average, the manually provided contracts reduce the number of false alarms by 93% as compared to the vacuous contracts’ false alarms, while the automatic derivation algorithm reduces the number of messages by 25%. The derived preconditions are in many cases weaker than the manually provided ones. Our initial study indicates that this happens when the integer analysis joins two different procedure behaviors. One potential remedy to this imprecision is by using sets of linear inequalities that allow to precisely represent logical-or.

6. RELATED WORK

Many academic and commercial projects produce practical tools that detect string manipulation errors at runtime, e.g., [31, 1, 26, 7]. The main disadvantage of runtime checking is that its effectiveness strongly depends on the input tested, and it does not ensure against future bugs on other inputs. Our goal is a conservative static tool that detects all string errors and provides an assurance against all such errors.

6.1 Static Detection of String Errors

Although the problem of string manipulation safety checking is to verify that accesses are within bounds [21, 2, 33], the domain of string programs requires that the analysis be capable of tracking the following features of the C programming language: (i) handling standard C functions, such as `strcpy()` and `strlen()`, which perform an unbounded number of loop iterations; (ii) statically estimating the length of strings (in addition to the sizes of allocated base addresses); this length is dynamically changed based on the index of the first null character; and (iii) simultaneously analyzing pointer and integer values is required in order to precisely handle pointer arithmetic and destructive updates.

Many academic projects produce unsound tools to statically detect string manipulation errors. In [23] an extension to LCLint is presented. Unsound lightweight techniques, heuristics, and in-code annotations are employed to check for buffer overflow vulnerabilities. Eau claire [4], a tool based on ESC-Java, checks for security holes in C programs by translating a subset of C to guarded commands. Its annotation language is similar in sense to CSSV. In [37] Wagner et al. present an algorithm that statically identifies string errors by performing a flow insensitive unsound analysis. The main disadvantage of all of these unsound tools is that they miss errors while CSSV does not miss any error. Furthermore, none of them can track effects of pointer arithmetic,

App.	Function	Source Code			CSSV				Msg		Deriving			
		LOC	SLOC	Contract	IP Vars	IP Size	CPU _{sec}	Space MB	False	Errors	CPU _{sec}	Space MB	Vacuous	Auto
EADS Airbus	RTC_Sl_SkipLine	13	260	SBI	39	109	2.6	12	0	0	0.3	3	5	5
	RTC_Se_CopieEtFiltre	66	773	SB	127	812	206	347	6	0	95	433	24	24
	RTC_Sl_FiltrerCarNonImp	19	114	S	13	151	0.3	2	0	0	0.2	2	4	4
	RTC_Sl_Find	26	820	SI	108	476	2.7	24	0	0	1.4	54	4	1
	RTC_Sl_StrNCat	8	299	SBI	54	182	0.9	6	0	0	0.2	3	2	0
	RTC_Sl_CalculerStringTime	33	567	SBI	86	529	76	127	0	0	131	173	21	4
	RTC_Sl_FormatMcduTo- Formatprinter	18	273	SB	58	323	6.9	28	0	0	6.8	27	9	9
	RTC_Sl_StoreIntInBuffer	35	222	SBI	59	346	9.8	43	0	0	3.3	22	15	15
RTC_Se_ComposerEntete	10	550	SI	77	352	3.4	23	0	0	1.3	12	2	0	
fixwrites	remove_newline	12	260	S	35	203	0.1	2	0	0	0.61	1	1	0
	insert_long	14	367	SB	138	571	13	99	0	2	23.4	86	5	0
	join	15	701	SB	95	443	2.1	23	0	2	6.7	15	2	2
	whole	30	423	S	46	352	1.2	20	0	1	0.6	4	9	9
	skip_balanced	20	258	SB	29	215	0.3	5	2	0	0.6	3	6	6
	bare	26	333	S	41	319	0.6	12	0	3	0.4	9	11	11

Table 5: The experimental results.

a widely used method for string manipulation. Sound algorithms for statically detecting string errors are presented in [13, 35]. However, they cannot handle all C, in particular multi-level pointers and structures. As far as we know, CSSV is the first sound tool to handle all C and in a rather precise manner.

6.2 Procedural Points-to Information

Many algorithms compute procedural pointer information to improve the cost and precision of interprocedural analysis, e.g., [27, 22, 10, 5]. In contrast, we focus on the problem of representing procedure points-to information in a way which allows us to perform strong updates in well-behaved programs. In [25] a modular parameterized pointer analysis (MoPPA) is described. MoPPA computes procedural pointer information during the process of computing global pointer information. In contrast, our algorithm utilizes existing whole-program scalable pointer analysis and transforms the global information to the procedural information. Our technique thus is more general since it applies to many pointer analysis algorithms and not just to Steensgaard’s analysis [36] which serves as the basis for MoPPA.

6.3 The Automatic Derivation Process

The Houdini annotation-derivation tool [14] tries ESC/Java with different annotations. Such an approach is inadequate in our case because the number of potential annotations is unbounded. In contrast, we derive a contract by forward and backward analyses of the integer program [16].

7. CONCLUSIONS

Buffer overflow is one of the most harmful sources of defects in C programs. Moreover, it makes software vulnerable to hacker attacks. We believe that CSSV provides evidence that sound analysis can be applied to statically verify the absence of all string errors in realistic applications.

Acknowledgments

We would like to thank Manuvir Das for providing and assisting us with AST-ToolKit and GOLF. Thanks to Bertrand Jeannot and Nicolas Halbwachs for providing us the polyhedra library and for their support. Thanks to Greta Yorsh for her assistance in the prototype implementation and for many technical insights. Thanks to Seth Hallem, Roman Manevich, Tom Reps, Ran Shaham, and Reinhard Wilhelm for their helpful comments.

8. REFERENCES

- [1] T.M. Austin, S.E. Breach, and G.S. Sohi. Efficient detection of all pointer and array access errors. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.* ACM Press, 1994.
- [2] R. Bodik, R. Gupta, and V. Sarkar. ABCD: eliminating array bounds checks on demand. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, 2000.
- [3] D.R. Chase, M. Wegman, and F. Zadeck. Analysis of pointers and structures. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, pages 296–310, New York, NY, 1990. ACM Press.
- [4] B. Chess. Improving computer security using extended static checking. In *IEEE Symposium on Security and Privacy*, 2002.
- [5] J. Choi, M. Gupta, M.J. Serrano, V.C. Sreedhar, and S.P. Midkiff. Escape analysis for java. In *Conf. on Object-Oriented Prog. Syst. Lang. and App.*, pages 1–19, 1999.
- [6] P. Cousot and N. Halbwachs. Automatic discovery of linear constraints among variables of a program. In *Symp. on Princ. of Prog. Lang.*, 1978.
- [7] C. Cowan, P. Wagle, C. Pu, S. Beattie, and J. Walpole. Buffer overflows: attacks and defenses for the vulnerability of the decade. In *In Proc. of the DARPA Information Survivability Conference and*

- Expo*, 1999.
- [8] M. Das. Unification-based pointer analysis with directional assignments. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, 2000.
 - [9] M. Das, B. Liblit, M. Fähndrich, and J. Rehof. Estimating the impact of scalable pointer analysis on optimization. In *Static Analysis Symp.*, 2001.
 - [10] A. Deutsch. Interprocedural may-alias analysis for pointers: Beyond k-limiting. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, pages 230–241, New York, NY, 1994. ACM Press.
 - [11] E.W. Dijkstra. *A Discipline of Programming*. Prentice-Hall, 1976.
 - [12] N. Dor. *Statically Detecting All Buffer Overflows in C*. PhD thesis, Univ. of Tel-Aviv, Israel, 2003. In preparation.
 - [13] N. Dor, M. Rodeh, and M. Sagiv. Cleaness checking of string manipulations in C programs via integer analysis. In *Static Analysis Symp.*, 2001.
 - [14] C. Flanagan, K. Rustan, and M. Leino. Houdini, an annotation assistant for Esc/java. In *Formal Methods for Increasing Software Productivity, volume 2021 of Lecture Notes in Computer Science*, 2001.
 - [15] R. Ghiya, D. Lavery, and D. Sehr. On the importance of points-to analysis and other memory disambiguation methods for c programs. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, 2001.
 - [16] N. Halbwachs. *Static Analysis of Linear Properties Invariantly Satisfied by the Numeric Variables of a program*. PhD thesis, Grenoble University, 1979.
 - [17] N. Halbwachs, Y.E. Proy, and P. Roumanoff. Verification of real-time systems using linear relation analysis. *Formal Methods in System Design*, 11(2):157–185, 1997.
 - [18] N. Heintze and O. Tardieu. Ultra-fast aliasing analysis using cla: A million lines of c code in a second. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, 2001.
 - [19] B. Jeannet. New polka library. Available at “<http://www.irisa.fr/prive/Bertrand.Jeannet/newpolka.html>”.
 - [20] B. W. Kernighan and D. M. Ritchie. *The C programming language*. Prentice-Hall, Englewood Cliffs, NJ 07632, USA, 1988.
 - [21] P. Kolte and M. Wolfe. Elimination of redundant array subscript range checks. *ACM SIGPLAN Notices*, 30(6):270–278, 1995.
 - [22] W. Landi. *Interprocedural Aliasing in the Presence of Pointers*. PhD thesis, Dept. of Comp. Sci., Rutgers Univ., 1991.
 - [23] D. Larochelle and D. Evans. Statically detecting likely buffer overflow vulnerabilities. In *10th USENIX Security Symposium*, 2001.
 - [24] G. Leavens and A. Baker. Enhancing the pre- and postcondition technique for more expressive specifications. In *Formal Methods*, 1999.
 - [25] D. Liang and M. J. Harrold. Efficient computation of parameterized pointer information for interprocedural analyses. In *Static Analysis Symp.*, 2001.
 - [26] A. Loginov, S. Yong, S. Horwitz, and T. Reps. Debugging via run-time type checking. In *Proc. of Fundamental Approaches to Softw. Eng. (FASE)*, April 2001.
 - [27] T.J. Marlowe and B. G. Ryder. An efficient hybrid algorithm for incremental data flow analysis. In *Symp. on Princ. of Prog. Lang.*, 1990.
 - [28] B. Miller, D. Koski, C. Lee, V. Maganty, R. Murthy, A. Natarajan, and J. Steidl. Fuzz revisited: A re-examination of the reliability of Unix utilities and services, 1995. Available at <http://www.cs.wisc.edu/~bart/fuzz/fuzz.html>.
 - [29] C. Morgan. *Programming from Specifications*. Prentice-Hall, Englewood N.J, 1990.
 - [30] E.W. Myers. A precise inter-procedural data flow algorithm. In *Symp. on Princ. of Prog. Lang.*, 1981.
 - [31] Inc. Rational. Purify software. Available at “<http://www.rational.com>”, 1995.
 - [32] Microsoft Research. AST-toolkit. 2002.
 - [33] R. Rugina and M.C. Rinard. Symbolic bounds analysis of pointers, array indices, and accessed memory regions. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, 2000.
 - [34] B. G. Ryder, W. A. Landi, P. A. Stocks, S. Zhang, and R. Altucher. A schema for interprocedural modification side-effect analysis with pointer aliasing. *ACM Transactions on Programming Languages and Systems*, 23(2):105–186, 2001.
 - [35] A. Simon and A. King. Analyzing string buffers in c. In *International Conference on Algebraic Methodology and Software Technology*, 2000.
 - [36] B. Steensgaard. Points-to analysis in almost-linear time. In *Symp. on Princ. of Prog. Lang.*, pages 32–41, 1996.
 - [37] D. Wagner, J. Foster, E. Brewer, and A. Aiken. A first step towards automated detection of buffer overrun vulnerabilities. In *Symp. on Network and Distributed Systems Security*, 2000.
 - [38] G. Yorsh. *CoreC: A Simplifier for C*, 2002. <http://www.cs.tau.ac.il/~gretay/GFC.htm>.