

Coccinelle: Tool support for automated CERT C Secure Coding Standard certification[☆]

Mads Chr. Olesen^a, René Rydhof Hansen^{a,*}, Julia L. Lawall^b, Nicolas Palix^b

^a*Department of Computer Science, Aalborg University,
Selma Lagerlöfs Vej 300, DK-9220 Aalborg East, Denmark.*

^b*Department of Computer Science, University of Copenhagen (DIKU),
Universitetsparken 1, DK-2100 Copenhagen East.*

Abstract

Writing correct C programs is well-known to be hard, not least due to the many language features intrinsic to C. Writing secure C programs is even harder and, at times, seemingly impossible. To improve on this situation the US CERT has developed and published a set of coding standards, the “CERT C Secure Coding Standard”, that (in the version currently being worked on) enumerates 122 rules and 180 recommendations with the aim of making C programs (more) secure. The large number of rules and recommendations makes automated tool support essential for certifying that a given system is in compliance with the standard.

In this paper we report on ongoing work on adapting the Coccinelle bug-finder and program transformation tool, into a tool for analysing and certifying C programs according to, e.g., the CERT C Secure Coding standard or the MISRA (the Motor Industry Software Reliability Association) C standard. We argue that such a tool must be highly adaptable and customisable to each software project as well as to the certification rules required by a given standard.

Furthermore, we present current work on integrating Clang (the LLVM C frontend) as a program analysis component into Coccinelle. Program analysis information, e.g., from data-flow or pointer analysis, is necessary both for more precise compliance checking, i.e., with fewer false positives, and also for enabling more complete checking, i.e., with fewer false negatives resulting from pointer aliasing.

Keywords: automated tool support, CERT C Secure Coding, certification

[☆]Supported by the ISIS project (FTP grant number 274-08-0214).

*Corresponding author

Email addresses: `mchro@cs.aau.dk` (Mads Chr. Olesen), `rrh@cs.aau.dk` (René Rydhof Hansen), `julia@diku.dk` (Julia L. Lawall), `npalix@diku.dk` (Nicolas Palix)

1. Introduction

Writing correct C programs is well-known to be hard. This is, in large part, due to the many programming pitfalls inherent in the C language and compilers, such as low-level pointer semantics, a very forgiving type system and few, if any, run time checks. Writing a *secure* C program is even more difficult, as witnessed by the proliferation of published security vulnerabilities in C programs: even seemingly insignificant or “small” bugs may lead to a complete compromise of security.

In an effort to improve the quality of security critical C programs, the US CERT¹ organisation is maintaining and developing a set of rules and recommendations, called the *CERT C Secure Coding Standard* (CCSCS), that programmers should observe and implement in C programs in order to ensure at least a minimal level of security. The version of CCSCS currently under development enumerates 122 rules and 180 recommendations covering topics ranging from proper use of C preprocessor directives and array handling to memory management, error handling and concurrency. The sheer number of rules and recommendations makes it almost impossible for a human programmer to manually guarantee, or even check, compliance with the full standard. Automated tool support for compliance checking is therefore essential.

In this paper we describe work in progress on a prototype tool for automated CCSCS compliance checking. The tool is based on the open source program analysis and program transformation tool *Coccinelle* that has been successfully used to find bugs in the Linux kernel, the OpenSSL cryptographic library [1, 2, 3], and other open source infrastructure software. *Coccinelle* is scriptable using a combination of a domain specific language, called *Semantic Patch Language* (SmPL), as well as in O’Caml and Python. The scripts, called *semantic patches*, specify search patterns partly based on syntax and partly on the control flow of a program. This makes *Coccinelle* easily adaptable to new classes of errors and new codebases with distinct API usage and code style requirements. *Coccinelle* does not perform program analysis in the traditional sense, e.g., data-flow analysis or range analysis. However, for the purposes of program certification and compliance checking such analyses are essential, both to ensure soundness of the certification and to improve precision of the tool. For this reason we are currently working on integrating the *Clang Static Analyzer* with *Coccinelle* in order to enable *Coccinelle* to use the analysis (and other) information found by Clang.

The Clang Static Analyzer is part of the C frontend for the LLVM project². In addition to classic compiler support, it also provides general support for program analysis, using the monotone framework, and a framework for checking source code for (security) bugs. The emphasis in the source code checkers of the Clang project is on minimising false positives (reporting “errors” that are not

¹Formerly known as the US Computer Emergency Response Team (www.cert.org)

²Web: <http://clang.llvm.org>

really errors) and thus is likely to miss some real error cases. To further enhance the program analysis capabilities of Clang, in particular for inter-procedural program analyses, we have integrated a library, called WALi³ for program analysis using weighted push-down systems (WPDS)[4] into Clang. To enable rapid prototyping and development of new or specialised analyses, we have implemented Python bindings for the WALi library.

The rest of the paper is organised as follows. In Section 2 we give an overview of the CERT C Secure Coding Standard including a brief description of the rule categories. Section 3 illustrates how Coccinelle can be used as a compliance checker, by implementing a rule from each category as a semantic patch. Section 4 describes how the Coccinelle rules can benefit from having access to program analysis information. Section 5 discusses current work in progress, including experiments and the integration of Clang and Coccinelle. Finally Section 7 concludes.

2. The CERT C Secure Coding Standard

The CERT C Secure Coding Standard (CCSCS) is a collection of rules and recommendations for developing secure C programs. The first version of the CCSCS was published in 2008 as [5]. However, in this paper we focus on the version currently being developed. The development process is a collaborative effort managed through the CCSCS web site⁴. The current version⁵ of the CCSCS consists of 122 rules and 180 recommendations. The rules and recommendations are divided into 16 categories covering the core aspects of the C programming language. Figure 1 shows an overview of these categories and a summary of the number of rules and recommendations in each category.

2.1. Overview of the CCSCS

Experience shows that when programming in C, certain programming practises and language features, e.g., language features with unspecified (or compiler dependent) behaviour result in insecure programs or, at the very least, in programs that are hard to understand and check for vulnerabilities. This experience is at the heart of the CCSCS. Many of the observed problems arise when programmers rely on a specific compiler’s interpretation of behaviour that is undefined in the ANSI standard for the C programming language (ANSI C99). Other problems are caused, or at least facilitated, by the flexibility of the C language and the almost complete lack of run-time checks.

Based on the observed problems, the US CERT has identified a number of key issues and developed a set of *rules* that specify both how to avoid problematic features and also constructively how to use potentially dangerous constructs

³Web: <http://www.cs.wisc.edu/wpis/wpds/>

⁴Web: <https://www.securecoding.cert.org/confluence/display/seccode/CERT+C+Secure+Coding+Standard>

⁵Last checked 5 May 2011.

Code	Long name	# of Rules/Recomm.	
01-PRE	Preprocessor	3	14
02-DCL	Declarations and Initialization	11	21
03-EXP	Expressions	12	22
04-INT	Integers	6	18
05-FLT	Floating Point	7	7
06-ARR	Arrays	7	3
07-STR	Characters and Strings	9	11
08-MEM	Memory Management	6	13
09-FIO	Input Output	16	19
10-ENV	Environment	3	5
11-SIG	Signals	6	3
12-ERR	Error Handling	4	8
13-API	Application Programming Interfaces	N/A	8
14-CON	Concurrency	9	2
49-MSC	Miscellaneous	11	22
50-POS	POSIX	12	4

Figure 1: Categories in the CERT C Secure Coding Standard

in a secure way, e.g., programming patterns for securely handling dynamic allocation and de-allocation of memory. The rules in the CCSCS are almost all unambiguous, universal and generally applicable in the sense that they do not depend on the specific application being developed. Furthermore the rules are, for the most part, formulated at the level of individual source files or even parts of source files and thus require little or no knowledge of the surrounding application or the context in which it is used. This makes the rules ideally suited for automated checking, although see Section 3 for a more detailed discussion of this.

In addition to the above mentioned rules, the CCSCS also contains an even larger number of *recommendations*. The recommendations often represent the *best practise* for programming secure systems. In contrast to the rules, the recommendations are not limited to constructs that are local to a single file or function, but often also cover more global issues as well as design issues, such as how to handle sensitive information, how to use and implement APIs, how to declare and access arrays, and so forth. While most of the recommendations are still amenable to automated analysis, it may take more work and, in particular, it will require configuring and specialising the automated tool to the specific project being checked, e.g., by specifying which data in the program may contain sensitive information or which macros are considered safe or how to canonicalize file names. A programmer is not required to follow the recommendations in order to be compliant with the CCSCS.

The CCSCS is much too large to cover in detail here, instead we give a brief overview of the different categories and the kind of (potential) errors they are designed to catch. In Section 3 we focus on how to check the rules using

Coccinelle and discuss one rule for every category in detail.

2.2. Categories of the CCSCS

Preprocessor (01-PRE). The rules and recommendations in this category are concerned with proper use of the C preprocessor. Most (large) C projects use preprocessor directives, especially macro definitions, extensively. Since these can dramatically change the “look” of a program, it is very important at least to avoid the many common pitfalls enumerated in this category.

Many static analysis tools are not very good at checking these rules since they typically work on the expanded code and thus do not even see the macros. This is unfortunate since a lot of semantic information can be gleaned from well-designed macros and their use.

Declarations and Initialization (02-DCL). The rules and recommendations in this category mostly cover tricky semantics of the type system and variable declarations such as implicit types, scopes, and conflicting linkage classifications.

The recommendations in this category codify good programming practises, e.g., using visually distinct identifiers (DCL02-C) and using typedefs to improve code readability (DCL05-C). While many of the recommendations can be automatically verified others (like DCL05-C) require human interaction.

Expressions (03-EXP). The rules and recommendations in this category are concerned with issues related to expressions, including (unspecified) evaluation order, type conversions, sizes of data types, general use of pointers, and so forth.

Below we show how rule EXP34-C (do not dereference null pointers) can be checked using the Coccinelle tool.

Integers (04-INT). The rules and recommendations in this category are concerned with issues related to proper handling of integers. The main emphasis for the rules is on avoiding overflows and wrap-around for very large or very small integer values. Automated checking for these rules can be difficult since that may require sophisticated data flow or interval analysis. Alternatively, a tool can instead check that a program includes sufficient checking in the program itself to avoid the dangerous situations. In some cases it is possible to use Coccinelle to automatically insert the proper checks. However, inserting such checks automatically would seem to violate the point of a security certification.

The recommendations are similarly concerned with conversions, limits and sizes of the integer types. Like the rules in this category, automated checking of the recommendations can be difficult and may require sophisticated program analysis.

Floating Point (05-FLP). the rules and recommendations in this category are concerned with issues relating to proper handling of floating point types: loss of precision, proper use of mathematical functions, and type conversion. Automated checking is at least as difficult as for the integer case.

Arrays (06-ARR). The rules and recommendations in this category focus on avoiding out of bounds array indexing and pointer access to arrays. Automated checking is likely to require pointer analysis in order to ensure correctness and to minimise false positives.

Characters and Strings (07-STR). The rules and recommendations in this category are concerned with: ensuring that strings are null terminated, proper size calculation of strings, and bounds checking for strings.

Memory Management (08-MEM). The rules and recommendations in this category cover some of the many pitfalls surrounding dynamic memory allocation, including not accessing freed memory, do not “double free” memory, only freeing dynamically allocated memory and so forth. Implementing memory management correctly is notoriously difficult and even small bugs in this category are likely to result in a security vulnerability, e.g., a buffer overflow or a null pointer dereference. Below we discuss rule MEM30-C (do not access freed memory) in more detail and show how it can be checked using Coccinelle.

Input Output (09-FIO). The rules and recommendations in this category are mainly concerned with the proper use of library functions for (file) input and output, including proper opening and closing of files, creation of temporary files, as well as secure creation of format strings.

Environment (10-ENV). The rules and recommendations in this category are concerned with proper handling of the execution environment, i.e., environment variables, and calls to external command processors are covered by the rules and recommendations in the ENV category.

Signals (11-SIG). The rules and recommendations in this category are concerned with raising and handling signals in a secure manner, including ensuring that signal handlers do not call `longjmp()` and do not modify or access shared objects.

Error Handling (12-ERR). The rules and recommendations in this category are concerned with detecting and handling errors and proper handling of the `errno` variable. Examples include not modifying the `errno` variable and not relying on indeterminate values of `errno`. Below we discuss the rule ERR33-C (detect and handle errors) in more detail and examine how Coccinelle can be used to check this rule. Note that this rule is different from most other rules in that it is actually *application dependent* since errors are detected and handled differently in different applications. Consequently, in order for an automated tool to support checking of this rule, it must be possible to customise and adapt the tool to a specific project’s error handling strategy.

Application Programming Interface (13-API). In the version of CCSCS currently under development, this category has no rules, only recommendations, since proper API design is highly application specific. Similar to the error handling (ERR) category above, automated tool support requires a very adaptable tool.

Concurrency (14-CON). The rules and recommendations in this category are general observations concerning concurrent programming such as avoiding race conditions and deadlocks (by locking in a predefined order).

Miscellaneous (49-MS). The rules and recommendations in this category are those that do not fit into any other category, e.g., it is recommended to compile cleanly at high warning levels (MSC00-C) and it is a rule that a non-void function's flow of control never reaches the end of the function (MSC37-C). Below we discuss rule MSC37-C (ensure that control never reaches the end of a non-void function) in more detail and show how this rule can be checked using Coccinelle.

POSIX (50-POS). The rules and recommendations in this category cover compliance with and proper use of POSIX. In particular things to avoid doing with POSIX, such as calling `vfork()` and not using signals to terminate threads.

3. Compliance Checking with Coccinelle

In this section we exemplify how Coccinelle can be used as an automated checker for CCSCS rules by implementing a Coccinelle semantic patch, each checking one rule, for each of the 16 rule categories in the CCSCS. Before going into the details of the individual rules, we briefly introduce Coccinelle; for lack of space we cannot give a thorough introduction to Coccinelle and the languages used to script it, instead we refer to previous work [1, 6, 7] and to each of the example implementations in the following where Coccinelle features are introduced as needed and discussed in more detail.

The Coccinelle tool was originally developed to provide support for documenting and automating updates to Linux device drivers necessitated by a change in the underlying API, the so-called *collateral evolutions* [7]. Finding the right place to perform collateral evolutions in a large code base requires a highly configurable and efficient engine for code searching. In Coccinelle this engine is based on model checking of a specialised modal logic, called CTL-VW, over program models [6] enabling search not only for specific syntactic patterns but also for control flow patterns. Individual program searches (and transformations) are specified in a domain specific language, called SmPL (for Semantic Patch Language), designed to be similar to the unified patch format widely used by Linux kernel developers and other open source developers. Such program searches are called *semantic patches* or even Coccinelle scripts. The combination of easy configurability and efficient search capabilities makes Coccinelle an excellent tool for searching for code patterns that may lead to potential

bugs or violations of coding standards. It has been successfully used to search for bugs in open source infrastructure software such as the Linux kernel and the OpenSSL cryptographic library [1, 2]. The Coccinelle tool is released under the GNU GPLv2 open source license.

3.1. PRE31-C: Avoid side-effects in arguments to unsafe macros

In contrast to most other analysis tools, Coccinelle generally works on the *unexpanded* source code, making it possible to search for code patterns that involve preprocessor directives. It also has (experimental) support for preprocessor directives for conditional compilation (`#ifdef`'s). The tool can automatically turn “well-behaved” `#ifdef`'s into corresponding C conditionals and thereby make the control flow structure visible for source code analysis. In this context, a “well-behaved” `#ifdef`, is one that is wrapped cleanly around an entire block of statements and does not, e.g., break or occur in the middle of a syntactical entity, e.g., an identifier or a function argument.

The CCSCS defines an *unsafe macro* to be a macro whose expansion is not guaranteed to evaluate its arguments only once. If a programmer is unaware either that a given macro is unsafe or even that it is a macro, since it may look and otherwise behave like a function, the programmer may use an expression with a side-effect as argument to the (unsafe) macro and thereby potentially crash the program or leave it vulnerable to attack.

In general it is recommended that unsafe macros are clearly labelled as such, e.g., by prefixing or suffixing the macro name with “UNSAFE”. For this simple case, we can use the following Coccinelle semantic patch to check that rule PRE31-C is followed:

```

1 @ unsafe @
2 expression E;
3 identifier x;
4 identifier mname ~="UNSAFE";
5 @@
6     mname(<+... \ ( E++ \ | E-- \ | ++E \ | --E \ | x = E \ |
7                 x += ... \ | x -= ... \ | x *= ... \ |
8                 x /= ... \ | x |= ... \ | x &= ... \ |
9                 x(...) \ ) ...+>

```

The semantic patch works by looking for identifiers that contain the string “UNSAFE” (declared on line 4) and are used in a function call with an argument that (potentially) has a side-effect (lines 6 to 9). We are not restricted to looking for specific constant strings in identifiers: regular expressions are fully supported. Note the explicit enumeration of all the possible forms such an argument can take, enclosed in ‘\ (’ and ‘\)’, with alternatives separated by ‘\ |’. This is called a *disjunction* in Coccinelle and it works by looking for the *first* disjunct that matches and then skipping the rest of the disjuncts. This “short circuit” evaluation of disjuncts is very useful to find (and ignore) non-violations, as shown in Section 3.6.

An alternative to looking for macros with special names is to maintain a database of names of unsafe macros and then perform a lookup in this database whenever a macro is used in a potentially unsafe way. The semantic patch developed in Section 3.2 in part illustrates such an approach, where O’Caml scripting is used to maintain and use a hashtable of identifiers for checking rule DCL32-C.

3.2. DCL32-C: Guarantee that mutually visible identifiers are unique

The ANSI C99 standard for the C programming language specifies that *at least* 63 initial characters are significant in an identifier. Thus, identifiers that share a 63 character long prefix may be considered identical by the compiler. The DCL32-C rule requires that all (in scope) identifiers are unique, i.e., must differ within the first 63 characters.

Below a Coccinelle semantic patch is shown that simply searches for all variable declarations. This seemingly simple search forms the heart of the semantic patch used to search for potential violations of the DCL32-C rule:

```
1 @@
2 type T;
3 identifier id;
4 @@
5     T id;
```

Observe that this is very similar to what a variable declaration looks like in a C program.

In Figure 2 the full semantic patch is shown. It collects all identifiers of length 63 or more into a hashtable and warns if there are (potential) violations of the rule. The rule does not take the scope of the declared identifiers into account and thus may give rise to unnecessary warnings (false positives). However, since identifiers of length 63 or more are rarely used this is unlikely to be a problem in practise. If, for a specific project, it turns out to be a problem, the semantic patch can be extended to take more scope information into account. The semantic patch includes a simple O’Caml script (lines 11 to 22) that collects all the found identifiers (of length 63 or more) and adds them to a hash table. Before adding an identifier to the hash table, it is checked for collisions, and thus potential violations, and a warning is printed if there are (potential) collisions (line 18) otherwise nothing is printed (line 20).

The basic semantic patch searching for declarations has been augmented with a *position meta-variable* denoted `@pos` (line 9). The position meta-variable is bound to the position (line and column number) of each match. Position meta-variables are useful both for reporting purposes, but also as “anchors” that can be used in further searches. Here, it is used only for reporting purposes: the position found in the code search is *inherited* by the O’Caml script (line 12), meaning that in the O’Caml script the variable ‘p’ now contains an O’Caml representation of the position where a long identifier was found.

```

1  @ initialize:ocaml @
2  let idhash = Hashtbl.create 128
3
4  @ decl @
5  type T;
6  identifier id;
7  position pos;
8  @@
9   T id@pos;
10
11 @ script:ocaml @
12 p << decl.pos;
13 x << decl.id;
14 @@
15 if (String.length(x)) >= 63 then
16   let sid = String.sub x 0 63 in
17   let _ = if (Hashtbl.mem idhash sid) then
18     (* print warning *)
19     else
20     (* do nothing *)
21     in
22   Hashtbl.add idhash sid (x,p)

```

Figure 2: Coccinelle script to find “long” identifiers.

3.3. EXP34-C: Do not dereference null pointers

In the CCSCS, the rationale for this rule is that attempts to dereference null pointers result in undefined behaviour. In recent years, attackers and vulnerability researchers have had great success at leveraging null pointer dereferences into full blown security vulnerabilities, making this rule very important for application security. The current version of the CCSCS contains an example involving the Linux kernel and the `tun` virtual network driver.

One potential source of null pointers, as noted in the CCSCS examples, is when memory allocation functions, e.g., `malloc()`, `calloc()`, and `realloc()`, fail and return null. If the return value from allocation functions is not properly checked for failure, and handled accordingly, there is a high risk that a program will eventually, or can be made to, dereference a null pointer.

Using Coccinelle to find such code patterns is straightforward. In Figure 3 the corresponding semantic patch is shown: we first look for calls to the relevant allocation functions (lines 8 to 14). The possible allocation functions are specified using the *disjunction* pattern (denoted by ‘(’, ‘|’, and ‘)’) that succeeds if either of the alternatives (separated by ‘|’) match. Following that, the script looks for a *control flow* path, represented by ‘...’, where the identifier (`x`) is *not* assigned to, i.e., a path where it is not modified (line 15), and where the identifier is not tested for “null-ness” (line 16). The latter is in order to cut down on the number false positives. Here the ‘... WHEN != x = E’ and the ‘WHEN != if(E == NULL) S1 else S2’ means along *any* control flow path where assignment to `x` does not occur, i.e., any control flow path where `x` is not modified and which contains no null test on `x`. Finally, we look for a dereference of `x` (lines 17 to 23), again using the disjunction pattern to specify three common

```

1 @@
2 identifier x;
3 expression E,E1;
4 type T1;
5 identifier fld;
6 statement S1, S2;
7 @@
8 (
9   x = (T1) malloc(...)
10  |
11   x = (T1) calloc(...)
12  |
13   x = (T1) realloc(...)
14 )
15   ... WHEN != x = E
16       WHEN != if(E == NULL) S1 else S2
17 (
18   *x
19  |
20   x[E1]
21  |
22   x->fld
23 )

```

Figure 3: Coccinelle script to find dereferencing of null pointers due to insufficient checks on memory allocation.

ways to dereference a pointer: as a pointer (line 18), as an array (line 20), or for field member access (line 22).

Note that, even though the semantic patch specifies that there can be no conditionals with a condition on the form ‘E==NULL’ (in line 16), Coccinelle will automatically also match variations of this condition such as ‘NULL==E’, and ‘!E’. This feature is called *isomorphisms* and is a general, customisable, and scriptable feature of Coccinelle designed to handle syntactic variations of the same semantic concept, in this case, comparing a variable to the NULL pointer. Isomorphisms, while not strictly necessary, represent a large reduction in the amount of work a programmer has to do when developing a semantic patch. Isomorphisms are also useful in developing patches that are more complete (cover more cases) since corner and special cases need only be handled once.

While the semantic patch in Figure 3 will catch many common violations of rule EXP34-C, it cannot catch all possible violations. First of all, null pointers may come from many other places than the memory allocation functions, e.g., user defined functions and library functions. In principle it is of course possible to manually extend the semantic patch with all the functions possibly returning a null pointer, however, this quickly becomes unwieldy. Another drawback of the semantic patch, as shown, is that it currently overlooks violations occurring *after* a null test. It is possible to manually refine the semantic patch to take more tests into account in a proper way. In [1] a more comprehensive Coccinelle approach to dereferencing of null pointers is described. This approach covers not only standard allocations functions, but basically any function returning null. In addition, care is taken to consider null tests and handle them properly.

Another alternative would be if the semantic patch could make use of information from a data-flow analysis. That way it would not be necessary to explicitly cover all syntactic possibilities for null testing or dereferencing. Many of the rules in this category (03-EXP) would likewise benefit from having access to program analysis information such pointer-analysis, e.g., rule EXP32-C, or data-flow analysis information, e.g., rule EXP41-C. In Section 4 we describe our current work on integrating analysis information into Coccinelle scripts.

3.4. INT30-C: Ensure that unsigned integer operations do not wrap

While it is not currently possible to perform complete checking of rules for secure handling of integers, Coccinelle can still be used to check that integer operations are suitably protected as suggested in the CCSCS, e.g., by a pre- or post-condition on the integer operation in question ensuring that potential violations will at least be caught at run-time. The semantic patch below implements a simple, and very specific, search to find indexing into an array, formed by the addition of integers that may potentially overflow:

```

1 @@
2   identifier arr;
3   unsigned int ui1, ui2;
4 @@
5   (
6     if(UINT_MAX - ui1 < ui2) {
7       ...
8     } else {
9       ...
10      arr[i1 + i2]
11      ...
12    }
13   |
14   arr[ui1 + ui2]
15  )

```

It is straightforward to generalise the above semantic patch to cover more operations and other potentially dangerous uses besides array indexing. By using the program transformation capabilities of Coccinelle, such pre- or post-conditions can be *inserted* automatically into the source code of an application. While this is useful for program development or re-engineering, it is less useful for checking that an application is compliant with the CCSCS and therefore we do not go into further details here.

Most of the rules in the 04-INT category require data-flow analysis to fully check them. In particular, such information could be used to statically guarantee that certain integer expressions could not possibly overflow and thus eliminate the need for an explicit check.

```

1  @ for @
2    float x;
3    double y;
4    statement S;
5  @@
6  (
7    for(<+... x ...+>; <+... x ...+>; <+... x ...+>) S
8  |
9    for(<+... y ...+>; <+... y ...+>; <+... y ...+>) S
10 )
11
12 @ while @
13   float x;
14   double y;
15 @@
16 (
17   while(<+... x ...+>) {
18     <+... x ...+>
19   }
20 |
21   while(<+... y ...+>) {
22     <+... y ...+>
23   }
24 )

```

Figure 4: Coccinelle script to find floating point loop counters.

3.5. FLP30-C: Do not use floating point variables as loop counters

Since the precision limits of floating point values are implementation/platform dependent, it could lead to non-portable code to use floating point variables as loop counters. The general problem of automatically deciding which variables in a loop are actually loop counters requires extensive program analysis. Instead of trying to catch every case, the semantic patch in Figure 4 looks for “suspicious” loops, i.e., loops where a floating point variable is referenced in the loop condition and for also in the loop body or the loop update expression for `while`-loops and `for`-loops respectively. Unfortunately, Coccinelle does not currently support `do/while` loops.

Similar to the rules for handling integers (Section 3.4), access to data-flow analysis information would enable more precise checking of the rules of this category (05-FLP) as well as checks for a wider range of potential violations.

3.6. ARR37-C: Do not add or subtract an integer to a pointer to a non-array object

The semantic patch shown below uses the Coccinelle disjunction operator (also discussed in Section 3.1). In particular, the semantic patch relies on the “short circuit” semantics of the disjunction, i.e., a disjunction matches if any of the disjuncts match and disjuncts are tried from top to bottom. Thus, any pointer arithmetic performed on an array pointer will be matched by an early disjunct (and ignored) in line 8. As a result, any expression with pointer arithmetic matched by a later disjunct (lines 9–10) will involve a non-array pointer since those have already been “filtered out” by earlier disjuncts. The location

(in code) of potential violations are recorded using the position meta-variable ‘pos’.

```
1 @@
2 type T;
3 T[] arr;
4 T *p;
5 expression E;
6 position pos;
7 @@
8 \(\ arr + E \ | \ arr - E \ | \ arr += E \ | \ arr -= E \ |
9   p@pos + E \ | \ p@pos - E \ |
10  p@pos += E \ | \ p@pos -= E
11 \)
```

Also for this category (06-ARR) access to program analysis information would be very useful: data-flow analysis to guarantee that array indices are within bounds, e.g., rules ARR30-C and ARR32-C, and pointer analysis to handle aliased pointers, e.g., rule ARR36-C and ARR37-C.

3.7. STR36-C: Do not specify the bound of a character array initialized with a string literal

Modifying a string literal may lead to undefined behaviour or even an access violation when the string literal is stored in read-only memory. Consequently, rule STR36-C forbids any modification of string literals. A simple semantic patch, shown below, can be used to find violations of rule STR36-C. The semantic patch looks for declarations of character arrays with an explicit bound (as represented by the Coccinelle meta-variable `E` of type `expression`) as well as an initializer (represented by the ‘...’ Coccinelle operator).

```
1 @@
2   identifier str;
3   expression E;
4 @@
5   char str[E] = ...;
```

Many of the rules in this category (07-STR) require both data-flow and pointer analysis to be checked in a complete manner, e.g., to determine that all access to a string is within bounds.

3.8. MEM30-C: Do not access freed memory

In the C programming language, as in most programming languages, using the value of a pointer to memory that has been deallocated, with the `free()` function, results in undefined behaviour. In practise, reading from deallocated memory may result in crashes, leaks of information, and exploitable security vulnerabilities. Rule MEM30-C ensures that deallocated memory will not be accessed. The problem underlying this rule is very similar to that described in

rule EXP34-C (do not dereference null pointers): instead of focusing on null pointers, this rule covers all pointers that have been freed.

Below a Coccinelle script is shown, covering some of the simple(r) cases of this rule. The script first looks for any identifier (declared in line 2) that occurs as an argument to the `free()` function (line 7). Following that, the script looks for a *control flow* path where the identifier (`x`) is *not* assigned to, i.e., a path where it is not modified (line 8). Finally, using the *disjunction* search pattern (denoted by ‘(’, ‘|’, and ‘)’) that succeeds if either of the alternatives (separated by ‘|’) match, the script looks for a *use* of the identifier that results in the actual violation. Here four common uses are covered: used as an argument to a function (line 10), dereferenced as a pointer (line 12) or an array (line 14), and dereferenced for member field access (line 16).

```
1 @@
2 identifier x;
3 expression E,E1;
4 function f;
5 identifier fld;
6 @@
7   free(x);
8   ... WHEN != x = E
9   (
10  f(...,x,...)
11  |
12  *x
13  |
14  x[E1]
15  |
16  x->fld
17  )
```

Proper memory management is essential for security critical C programs and is enforced by the rules in this category (08-MEM). Using Coccinelle to find various types of bugs caused by flawed memory management have been studied extensively in the context of the Linux kernel [8, 9].

Pointer analysis can help catch rule violations that occur through the use of an aliased pointer.

3.9. FIO34-C: Use *int* to capture the return value of character IO functions

Exploiting a security vulnerability most often requires an attacker to present some special input to the program under attack, e.g., specially crafted strings that overflow a buffer and overwrite important system information. The 15 rules of the input/output rule category (09-FIO) illustrate (some of) the difficulty of handling input/output in a secure manner.

Some of the character input/output function in the C standard library return an integer representing the character read/written. As noted in the CCSCS,

casting such an integer to a character (of type `char`) may result in a character being misinterpreted as the end-of-file (EOF) marker with potentially devastating results. Rule FIO34-C requires programmers to use integer variables to hold such return values. The semantic patch for checking rule FIO34-C is quite simple: first look for “non-violations” (assignments to a variable of type `int`) and ignore them (the first disjunct), next look for potential violations (assignments to variables of type other than `int`):

```

1 @@
2   int i;
3   identifier c;
4   position pos;
5 @@
6   (
7     i = \( fgetc(...) \| getc(...) \| getchar(...) \|
8           fputc(...) \| putc(...) \| putchar(...) \|
9           ungetc(...) \|
10    |
11    c@pos = \( fgetc(...) \| getc(...) \| getchar(...) \|
12              fputc(...) \| putc(...) \| putchar(...) \|
13              ungetc(...) \|
14   )

```

3.10. ENV32-C: All `atexit` handlers must return normally

In order to provide programmers with an opportunity to execute cleanup code upon program exit, the C programming language allows programmers to register exit handlers, i.e., functions that are executed immediately before the program terminates. If an exit handler is not allowed to run to completion it may lead to undefined behaviour. In particular, exit handlers must not be terminated by a call to `exit()` or by invoking a `longjmp`.

The semantic patch below first looks for functions that are registered as exit handlers using the `atexit()` function (lines 1–4). The *inheritance* mechanism of Coccinelle is used in line 7 to transfer the names of functions registered as exit handlers to the subsequent rule (lines 6–18) that checks exit handlers for calls to `exit()` or `longjmp()`; the position of any such call is recorded in the position meta-variable named `pos` for later use (lines 13 and 15).

```

1 @ atexit @
2 identifier fn;
3 @@
4   atexit(fn)
5
6 @@
7   identifier atexit.fn;
8   position pos;
9   @@

```

```

10 void fn(void) {
11     ...
12 (
13     exit@pos(...);
14 |
15     longjmp@pos(...);
16 )
17     ...
18 }

```

3.11. SIG30-C: Call only asynchronous-safe functions within signal handlers

Like environment interaction, discussed in Section 3.10, signals can be used by an attacker to manipulate the control flow of a program. Proper use of signals and signal handlers is the topic of category 11-SIG.

The CCSCS defines an *asynchronous-safe* function as a function that can be safely called without sideeffects from within a signal handler context. Calling a function that is not asynchronous-safe from a signal handler may result in undefined behaviour and should be avoided.

The set of asynchronous-safe function is application and platform specific. Therefore the semantic patch shown below, starts by initialising a hash table with the names of functions that are asynchronous-safe (line 2). In a manner similar to that used for rule ENV32-C, as shown in Section 3.10, the names of functions registered as signal handlers are recorded (lines 4–7) and used to find the definition of any signal handler (lines 9–16) and check that all functions called in the signal handler are in the hash table of allowed asynchronous-safe functions. The latter is handled by an O’Caml script that performs a simple lookup in the hash table (line 21).

```

1 @ initialize:ocaml @
2 let idhash = (* read hash table from file *)
3
4 @ signal @
5 identifier fn;
6 @@
7     signal(...,fn)
8
9 @ handler @
10 identifier atexit.fn;
11 identifier fnc;
12 position pos;
13 @@
14 void fn(void) {
15     <+... fnc@pos(...) ...+>
16 }
17

```

```

18 @ script:ocaml @
19 func << handler.fnc
20 @@
21 if not (Hashtbl.mem idhash func) then
22   (* record potential violation *)
23 else
24   (* ignore *)

```

3.12. ERR33-C: Detect and handle errors

The lack of proper exceptions in the C programming language means that error conditions have to be explicitly encoded and communicated to other parts of the program. Most often a run-time error in a given C function will be communicated by returning an *error value*, frequently -1 or NULL. Ignoring an error condition is highly likely to lead to unexpected and/or undefined behaviour, it is therefore essential that the return value is always checked for all calls to a function that may return an error value and that any error condition is handled properly. Rule ERR33-C formalises this requirement.

This rule differs from most of the other rules in the CCSCS in that it is almost entirely application dependent, since it is up to each application or software project to decide how, specifically, error conditions are signalled, what error values are used, what they mean, and how they must be checked and handled. It is therefore impossible to come up with a single, or even a few, rules that will cover the entire spectrum of possibilities. Thus, for a tool to be useful and effective it *must* be very customisable in order to adapt it to project specific code styles and policies. We believe that the specialised semantic patch language (SmPL) used in Coccinelle provides an excellent, and highly adaptable, platform for developing project specific rule checkers.

As an example of how Coccinelle can be customised for project specific error handling standards, we show in [2] how Coccinelle was used to find several bugs in some error handling code in the OpenSSL cryptographic library. Coccinelle has also been used to find flaws in the error handling of the Linux kernel [1].

3.13. Application Programming Interfaces (13-API)

At the time of writing, the latest version of CCSCS does not define any rules for the *Application Programming Interface* category.

3.14. CON33-C: Avoid race conditions when using library functions

Several of the functions defined in the C standard library, including `getenv()` and `rand()`, are not guaranteed to be thread-safe and their use in multi-threaded programs may lead to race conditions that can potentially be leveraged to break security.

Using a semantic patch almost identical to that described in Section 3.11 can be used to find and flag uses of such library functions. For this rule the hash table would contain the names of the potentially *unsafe* library functions instead of the names of asynchronous-safe functions.

3.15. *MSC37-C: Ensure that control never reaches the end of a non-void function*

Non-void functions are required to return a value, using the `return` statement. It results in undefined behaviour to use the return value of a non-void function where control flow reaches the end of the function, i.e., without having explicitly returned a value. For this reason the CCSCS requires that all control flows in a non-void function *must* end in a (non-empty) return statement.

Below a semantic patch is shown, that finds non-void functions with a control flow path not ending in a non-empty return statement. The overall strategy for this search is to first find all `void` functions (line 1 to 7), i.e., functions that are not supposed to return a value, in order to rule them out in our search. Next, we find all function declarations *except* for the functions we have earlier identified as `void` functions (line 13). Once such a function is found, we start looking for a control flow path that does *not* contain a `return` statement (line 16).

Observe that the head of the latter search pattern (line 9) not only contains the name of the search pattern (`func`) but also a directive to Coccinelle that it should disable the use of the ‘`ret`’-isomorphism (cf. the discussion of isomorphisms in Section 3.3) in order to avoid unwanted, potential interference from the isomorphism system. The header also specifies that the current rule should look for the *existence* of a control flow path with the required property, rather than checking for the property along *all* control flow paths, since we have a potential violation if there is even a single control flow path without a `return` statement.

```
1 @ voidfunc @
2 function FN;
3 position voidpos;
4 @@
5     void  FN@voidpos(...) {
6         ...
7     }
8
9 @ func disable ret exists @
10 type T;
11 expression E;
12 function FN;
13 position pos != voidfunc.voidpos;
14 @@
15     T  FN@pos(...) {
16         ... WHEN != return E;
17     }
```

The problem caught by the above semantic patch is inherently syntactic and control flow based, and thus very well suited for Coccinelle searches. Furthermore, checking for violations can be done in a universal and application independent way.

3.16. POSIX-C: Do not use `vfork()`

The Posix category (50-POS) is concerned with guidelines for proper use of POSIX functions and it is *not* part of the core CCSCS. It is included to show how other coding standards can be integrated with the CCSCS.

The `vfork()` function must not be used, according to the CCSCS, due to numerous potential problems, including security issues and undefined behaviour. A trivial semantic patch searches for uses of the `vfork()` function and records the position for later use:

```
1 @@
2   position pos;
3 @@
4   vfork@pos(...)
```

The CCSCS explicitly recommends always using `fork()` function rather than the `vfork()` function. As already mentioned in Section 3.4, Coccinelle can be used to automatically perform program transformations. In this case it is possible to replace all calls to `vfork()` with calls to `fork()`:

```
1 @@
2   param list[n] params;
3 @@
4 -   vfork(params)
5 +   fork(params)
```

In general automatic program transformation is not useful for verifying that a program is compliant with the CCSCS, but in this particular case where problem identification and solution is unambiguous and context independent it can be a very useful and effective tool.

4. Adding Program Analysis Information

From the discussion in the previous section of the categories and how specific rules can be checked using Coccinelle, it should be clear that while Coccinelle is useful for compliance checking it would benefit greatly from having access to proper program analysis information, e.g., for more precise and comprehensive tracking of potential null pointers. Such information could also be used to make checkers more succinct and efficient because fewer syntactic cases need to be covered.

In this section we first discuss how SmPL, the language used for writing semantic patches, may be extended to make program analysis information readily available in an easy to use manner. We also describe an experimental integration of an external program analysis engine into Coccinelle, namely the Clang Static Analyzer⁶.

⁶Web: <http://clang.llvm.org>

4.1. Pointer Analysis: Tracking NULL Pointers and Aliases

Consider the rule EXP34-C (do not dereference null pointers). Here the problem is to find all expressions that may potentially dereference a null pointer. With access to pointer analysis information, every expression that may result in a null pointer can be found and tagged. Note that this is independent of how an expression may result in a null pointer, i.e., it is no longer necessary to explicitly track information only from allocation functions in the semantic patch, since this is handled by the analysis.

Below we show how such analysis information could be incorporated into a semantic patch. The following semantic patch is intended to illustrate one possible way to make analysis information available to semantic patches:

```
1 @@
2 identifier x, fld;
3 expression E1;
4 analysis[null] NINF;
5 @@
6 ( *x@NINF
7 | x@NINF[E1]
8 | x@NINF->fld
9 )
```

The main thing to note in the above semantic patch is the ‘analysis’ declaration (line 4) that declares a meta-variable, called NINF. This meta-variable is then used in much the same way as position meta-variables: by “tagging” an expression with the ‘NINF’ meta-variable, e.g., like ‘x’ in line 7, only expressions that match the syntax (in this case an array) and that may also result in a null pointer are matched by the semantic patch.

Taking this a step further, we can also use analysis information to find all (sub-)expressions that are potential dereferences and then simply search for all expressions that are both tagged as potentially dereferencing and also as potentially resulting in a null pointer. Here a dereferencing expression is taken to mean an expression that may in any way do a pointer dereference:

```
1 @@
2 expression E;
3 analysis[null] NINF;
4 @@
5 *(E@NINF)
```

Here we rely on the isomorphisms of Coccinelle (see the discussion in Section 3.3) to expand the pointer dereference in line 5 to automatically cover all possible forms of pointer dereference.

Since pointers in C may be *aliases* for the same location in memory, it is important that the pointer analysis not only tracks potential null pointers but also tracks all potentially aliasing pointers. This is often called an *alias analysis* or a *points-to analysis*. Such analysis information would be useful in

many other situations, e.g., in the rule MEM30-C (do not access freed memory) where access may occur through an alias. The following semantic patch (with alias analysis information available) would capture this situation (see below for an explanation):

```
1 @@
2 identifier x, y;
3 expression E,E1;
4 function f;
5 identifier fld;
6 analysis[alias] xyalias;
7 @@
8   free(x@xyalias);
9   ... WHEN != y@xyalias = E
10  (
11   f(...,y@xyalias,...)
12  |
13   *y@xyalias
14  |
15   y@xyalias[E1]
16  |
17   y@xyalias->fld
18  )
```

The idea in the above semantic patch is that we first declare an analysis meta-variable in line 6 (called ‘xyalias’). Then, in line 8, we match a call to ‘free()’ on an identifier ‘x’ and bind the xyalias meta-variable to any available alias analysis information for x. Following that we match any assignments to and use of *any identifier* y that is an *alias* for x (represented by y@xyalias in lines 9, 11, 13, 15, and 17).

4.2. Integrating Clang and Coccinelle

Clang was chosen as the external program analysis engine for Coccinelle for several reasons: it is open source, it is being (very) actively developed, it has good support for writing new analyses, it provides a robust and proven infrastructure for manipulating C programs, and so forth.

In this prototyping phase, the emphasis has been on tool integration rather than extending SmPL. Consequently, it is not yet possible to use the ‘analysis’ declaration illustrated in the last section. Instead we use positions, as implemented by the ‘position’ meta-variables, to look up relevant analysis information (see below for details).

For the prototype we run Clang as a pre-processor to Coccinelle, computing the relevant program analysis information and storing it in a file, which may subsequently be read by a semantic patch. Below we show how this works using Python scripting in the semantic patch:

```

1 @ initialize:python @
2
3 # read analysis information generated by Clang into
4 # Python dictionaries: NINF indexed by positions
5
6 @ expr @
7 expression E;
8 position pos;
9 @@
10 *(E@pos)
11
12 @ script:python @
13 p << expr.pos
14 @@
15
16 # lookup NINF status in Clang data
17 if not (NINF[p]):
18     # remove the match
19 else:
20     # accept the match and continue

```

4.2.1. Information Exchange Format

To have a flexible and powerful exchange format for the program analysis information, we chose to use the framework of weighted push-down systems (WPDS). The analysis result of a WPDS analysis is a weighted finite automaton (WFA), describing the analysis information associated with each program point. WFAs allow expressing this information in a context-sensitive manner, while very efficient algorithms exist. We extended Clang with the analysis framework of WPDSs, using the library WALi. This enables us to model the control-flow from Clang as a push-down system, and plug-in different weight domains. Weight domains for different analyses have been presented [10], such as affine-relations analysis, generalised gen-kill analysis and may-aliasing pointer analysis. We have used the gen-kill weight domain to implement a reaching definitions analysis within Clang. The analysis result can then be pre-processed in Coccinelle scripts, as illustrated above, e.g., to get maybe-null analysis information. The benefit of using Clang is that the control-flow graph of the program is readily available, with some infeasible paths automatically pruned.

The analysis is written as a special analysis pass that constructs the WPDS, assigns weights, and perform a query for each function. The analysis results (annotated weighted finite automata) are output, and subsequently interpreted by the concrete Coccinelle script when analysis information is needed. Currently the Python scripting interface is used with some additional support code for calling Clang and interpreting the output.

The information that we have integrated at this point is the reaching defini-

tions analysis. The output from Clang is a textual representation of the solved WFA, an example of one line of this output is:

```
( p , ( uninit_use.c , ( 5 , 9 ) ) , accept ) <\S.(S - {NULL}) U
      {(simple:a@uninit_use.c:3:9@uninit_use.c:4:9,1)}>
```

All lines are split into their components:

From state of the WPDS, which will be the state p in most cases.

Symbol in this case “ $(uninit_use.c , (5 , 9))$ ” indicating the program point.

To state which will be the accepting state $accept$.

The weight associated with this transition, which is the program analysis information associated with the program point.

The weight again needs to be parsed, in this case into its gen and kill set. In the above example the kill set is empty, and the gen set adds a definition point of the variable $simple:a@uninit_use.c:3:9$, namely that it can be defined at $uninit_use.c:4:9$.

Variables are named from: the function they are defined in, their identifier and the position they are defined at. All positions are made up of: a file name, line number and column number.

Finally, a dictionary data structure is constructed such that the reaching definitions for a variable at a program point can be looked up.

One use is to look for uninitialised variables being used, where the basic semantic patch is:

```
1 @ uninituse @
2 type T; identifier I;
3 position defloc , useloc;
4 identifier FN;
5 @@
6 // look for declarations with no assignment
7 T@defloc I;
8 ... when any
9 //which are then used
10 (
11     FN@useloc(...,I,...);
12 |
13     I@useloc
14 )
```

Before being able to use a location from Coccinelle we have to account for small differences in how locations are presented in the CFG of Clang and Coccinelle, e.g. precisely where a variable is defined:

```
1 int a;
2     ^ Clang define location
3     ^ Coccinelle define location
```

Another example is that the use found might be part of a larger expression, so we will have to find the location of the entire expression. Currently we simply map a Coccinelle location to the closest Clang location on the same line.

We can then discard false positive matches, based on whether the data can actually flow from the found definition to the found use, in a somewhat cleaner way than specifying all possible ways the variable could have been modified. The approach of course becomes much more powerful when including analysis information from a pointer analysis.

5. Work in Progress

In this section we discuss the current status and work-in-progress for using Coccinelle to check for CCSCS compliance. In particular we discuss an improved approach to integrating Clang and Coccinelle, as well as compliance checking of the full standard for real world software projects.

5.1. Clang in Coccinelle

Our prototype gave us some experiences to guide the future development:

- Not all information fitted nicely into our exchange format. For example sometimes we simply want to know a piece of information about a data type, e.g. which cases an enum contains.
- We would like to advantage of more of Clang's features: for example Clang has a path-based reachability engine that we could use to rule out false positives.
- The fact that we used Clang as a pre-processor meant we would need to find all possibly needed information, before Clang exited, and then read it into Coccinelle. We would not be able to do an incremental analysis, only asking for smaller pieces of information at a time.

Based on these experiences we came up with a second approach: Instead of using Clang as a pre-processor we want to create Python bindings for the internal datastructures of Clang, and use these in our semantic patches. In this way instead of having two separate processes, we would have Clang running within the Coccinelle process' address space. With appropriate support code this would solve the problem of defining an exchange format: all the information Clang has will be available in the Python scripting in the semantic patch.

In addition we will be able to make calls to different parts of the modular Clang codebase, for exploiting the internal parts of Clang in specific cases when it makes sense. We would also be able to much more closely query for the only the information we need, and allow for an incremental analysis.

The disadvantage is that the semantic patches would be very dependent on the internal structure of Clang, for which the Clang developers give no guarantees. This could be alleviated to a large degree by defining abstracting interfaces, in Python, that represent the most commonly used analysis information from

Clang in an uniform way. In this way only the interfaces will have to be maintained. However, since we would not have a data exchange format it will be harder to integrate with other tools, since the each tool will have to be handled separately.

We are currently in the process of creating the Python bindings and the library for a common interface. We are also investigating how best to exploit the different features present in the Clang codebase.

The bindings are created using SWIG allowing for a (hopefully) reduced maintenance burden as the Clang codebase evolves. There are a few involved technical difficulties which need to be addressed, due to differences in the handling of objects in C++ and Python, but the solutions are well described in the SWIG documentation.

5.2. Compliance Checking Real World Software

Coccinelle has already been used successfully to find numerous bugs in the Linux kernel, the OpenSSL library, and other open source projects used in the “real world”. Especially the experience with bug finding in the Linux kernel shows that the approach scales well even to very large software projects.

One of the biggest problems when checking such large projects, is the number of *false positives*, i.e., warnings of potential violations that turn out not to be violations. Here the customisability of Coccinelle has turned out to be a great tool for reducing the number of false positives, since it enables a programmer to refine the semantic patches to take the project specific code styles into account that give rise to the most false positives.

The integration of program analysis information, e.g., obtained from Clang, will enable a code search to take (more) semantic information into account and will thus reduce the number of false positives further.

5.3. Implementing Checkers for the Full Standard

While we have only detailed the implementation of Coccinelle checkers for one rule from each category, we plan to implement Coccinelle checkers for all the CCSCS rules that are suitably application independent. For rules that are application dependent, such as rule ERR33-C (discussed in Section 3.12), it seems possible to provide an “abstract” semantic patch that can be instantiated with project specific details similar to the approach taken in [2].

We intend to make the complete set of checkers available for download as open source.

6. Related Work

The past decade has seen the development and release of numerous compile time tools for program navigation, bug finding and code style checking for programs written in C, as well as many other languages. These tools include the MC tool [11] (later used as basis for the commercial tool Coverity Prevent).

Similar to Coccinelle, the MC tool is a bug finder that can be adapted to specific projects, however the source code for MC has never been released.

Splint [12] and Flawfinder [13] are two examples of Open Source bug finders. Both are able to check for a relatively small set of bugs. Both are somewhat adaptable but requires either (light-weight) annotation of the source code or Python programming.

While several commercial static analysis tools support compliance checking⁷ for a wide spectrum of coding standard, including the CCSCS, we are not aware of any Open Source bug finder tools working towards this goal.

7. Conclusion

In this paper we have shown that the Coccinelle tool is very well suited for checking some of the rules comprising the CERT C Secure Coding Standard. We have further argued that integrating program analysis information would facilitate even more comprehensive, more expressible, and even more flexible semantic patches to be written.

References

- [1] J. L. Lawall, J. Brunel, N. Palix, R. R. Hansen, H. Stuart, G. Muller, Wysiwb: A declarative approach to finding api protocols and bugs in linux code, in: Proceedings of the 2009 IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2009, IEEE, Estoril, Lisbon, Portugal, 2009, pp. 43–52.
- [2] J. L. Lawall, B. Laurie, R. R. Hansen, N. Palix, G. Muller, Finding error handling bugs in openssl using coccinelle, in: Eighth European Dependable Computing Conference, EDCC-8, IEEE Computer Society, Valencia, Spain, 2010, pp. 191–196.
- [3] N. Palix, J. L. Lawall, G. Muller, Tracking code patterns over multiple software versions with herodotos, in: J.-M. Jézéquel, M. Südholt (Eds.), Proceedings of the 9th International Conference on Aspect-Oriented Software Development, AOSD 2010, ACM, Rennes and Saint-Malo, France, 2010, pp. 169–180.
- [4] T. Reps, S. Schwoon, S. Jha, D. Melski, Weighted pushdown systems and their application to interprocedural dataflow analysis, *Science of Computer Programming* 58 (1-2) (2005) 206–263.
- [5] R. C. Seacord, *The CERT C Secure Coding Standard*, Addison-Wesley, 2008.

⁷See the CCSCS web page for details on tool support.

- [6] J. Brunel, D. Doligez, R. R. Hansen, J. L. Lawall, G. Muller, A foundation for flow-based program matching: using temporal logic and model checking, in: Z. Shao, B. C. Pierce (Eds.), Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, ACM, Savannah, GA, USA, 2009, pp. 114–126.
- [7] Y. Padioleau, J. L. Lawall, R. R. Hansen, G. Muller, Documenting and automating collateral evolutions in linux device drivers, in: J. S. Sven-tek, S. Hand (Eds.), Proceedings of the 2008 EuroSys Conference, ACM, Glasgow, Scotland, UK, 2008, pp. 247–260.
- [8] H. Stuart, Hunting bugs with Coccinelle, Master’s thesis, Department of Computer Science (DIKU), University of Copenhagen (2008).
- [9] H. Stuart, R. R. Hansen, J. Lawall, J. Andersen, Y. Padioleau, G. Muller, Towards easing the diagnosis of bugs in OS code, in: 4th Workshop on Programming Languages and Operating Systems (PLOS 2007), Stevenson, WA, USA, 2007.
- [10] T. Reps, A. Lal, N. Kidd, Program analysis using weighted pushdown systems, in: Proceedings of the 27th international conference on Foundations of software technology and theoretical computer science, FSTTCS’07, Springer-Verlag, 2007, pp. 23–51.
- [11] D. R. Engler, B. Chelf, A. Chou, S. Hallem, Checking system rules using system-specific, programmer-written compiler extensions, in: Fourth USENIX Symposium on Operating Systems Design and Implementation (OSDI), San Diego, CA, 2000, pp. 1–16.
- [12] D. Larochelle, D. Evans, Statically Detecting Likely Buffer Overflow Vulnerabilities, in: Proc. of the 10th USENIX Security Symposium, USENIX, Washington D.C., USA, 2001, pp. 177–190.
URL <http://lclint.cs.virginia.edu/>
- [13] D. Wheeler, Flawfinder home page, Web page: <http://www.dwheeler.com/flawfinder/> (Oct. 2006).
URL <http://www.dwheeler.com/flawfinder/>