

Model Checking x86 Executables with CodeSurfer/x86 and WPDS++[★]

G. Balakrishnan¹, T. Reps^{1,2}, N. Kidd¹, A. Lal¹, J. Lim¹,
D. Melski², R. Gruian², S. Yong², C.-H. Chen, and T. Teitelbaum²

¹ Comp. Sci. Dept., University of Wisconsin; {bgogul, reps, kidd, akash, junghee}@cs.wisc.edu

² GrammaTech, Inc.; {melski, radu, suan, chi-hua, tt}@grammatech.com

Abstract. This paper presents a tool set for model checking x86 executables. The members of the tool set are *CodeSurfer/x86*, *WPDS++*, and the *Path Inspector*. CodeSurfer/x86 is used to extract a model from an executable in the form of a *weighted pushdown system*. WPDS++ is a library for answering generalized reachability queries on weighted pushdown systems. The Path Inspector is a software model checker built on top of CodeSurfer and WPDS++ that supports safety queries about the program’s possible control configurations.

1 Introduction

This paper presents a tool set for model checking x86 executables. The tool set builds on (i) recent advances in static analysis of program executables [2], and (ii) new techniques for software model checking and dataflow analysis [5, 26, 27, 21]. In our approach, CodeSurfer/x86 is used to extract a model from an x86 executable, and the reachability algorithms of the WPDS++ library [20] are used to check properties of the model. The Path Inspector is a software model checker that automates this process for safety queries involving the program’s possible control configurations (but not the data state). The tools are capable of answering more queries than are currently supported by the Path Inspector (and involve data state); we illustrate this by describing two custom analyses that analyze an executable’s use of the run-time stack.

Our work has three distinguishing features:

- The program model is extracted from the executable code that is run on the machine. This means that it automatically takes into account platform-specific aspects of the code, such as memory-layout details (i.e., offsets of variables in the run-time stack’s activation records and padding between fields of a struct), register usage, execution order, optimizations, and artifacts of compiler bugs. Such information is hidden from tools that work on intermediate representations (IRs) that are built directly from the source code.
- The entire program is analyzed—including libraries that are linked to the program.
- The IR-construction and model-extraction processes do not assume that they have access to symbol-table or debugging information.

Because of the first two properties, our approach provides a “higher fidelity” tool than most software model checkers that analyze source code. This can be important for certain kinds of analyses; for instance, many security exploits depend on platform-specific features, such as the structure of activation records. Vulnerabilities can escape notice when a tool does not have information about adjacency relationships among variables.

Although the present tool set is targeted to x86 executables, the techniques used [2, 27, 21] are language-independent and could be applied to other types of executables.

The remainder of the paper is organized as follows: Sect. 2 illustrates some of the advantages of analyzing executables. Sect. 3 sketches the methods used in CodeSurfer/x86 for IR recovery. Sect. 4 gives an overview of the model-checking facilities that the tool set provides. Sect. 5 discusses related work.

2 Advantages of Analyzing Executables

This section presents some examples that show why analysis of an executable can provide more accurate information than an analysis that works on source code.

One example of a mismatch between the program’s intent and the compiler’s generated code can be seen in the following lines of source code taken from a login program [18]:

```
memset(password, '\0', len);  
free(password);
```

[★] Portions of this paper are based on a tool-demonstration paper of the same title that will appear at CAV 2005.

The login program (temporarily) stores the user’s password—in clear text—in a dynamically allocated buffer pointed to by pointer variable `password`. To minimize the lifetime of sensitive information (in RAM, in swap space, etc.), the source-code fragment shown above zeroes-out the buffer pointed to by `password` before returning it to the heap. Unfortunately, a compiler that performs useless-code elimination may reason that the program never uses the values written by the call on `memset` and therefore the call on `memset` can be removed, thereby leaving sensitive information exposed in the heap [18]. This vulnerability is invisible in the source code; it can only be detected by examining the low-level code emitted by the optimizing compiler.

A second example where analysis of an executable does better than typical source-level analyses involves pointer arithmetic and function pointers. Consider the following piece of code:

```
int (*f)(void);
int diff = (char*)&f2 - (char*)&f1; // The offset between f1 and f2
f = &f1;
f = (int (*)(void))(char*)f + diff; // f now points to f2
(*f)(); // indirect call;
```

Existing source-level analyses (that we know of) are ill-prepared to handle the above code. The common assumption is that pointer arithmetic on function pointers leads to undefined behavior, so either (a) they assume that the indirect function call might call any function; or (b) they simply ignore the arithmetic and assume that the indirect function call calls `f1` (on the assumption that the code is ANSI-C compliant). In contrast, the *value-set analysis* (VSA) algorithm [2] used in CodeSurfer/x86 correctly identifies `f2` as the invoked function. Furthermore, VSA can detect when pointer arithmetic results in a function pointer that does not point to the beginning of a function; the use of such a function pointer to perform a function “call” is likely to be a bug (or else a very subtle, deliberately introduced security vulnerability).

A third example involves a function call that passes fewer arguments than the procedure expects as parameters. (Many compilers accept such (unsafe) code as an easy way of implementing functions that take a variable number of parameters.) With most compilers, this effectively means that the call-site passes some parts of one or more local variables of the calling procedure as the remaining parameters (and, in effect, these are passed by reference—an assignment to such a parameter in the callee will overwrite the value of the corresponding local in the caller.) An analysis that works on executables can be created that is capable of determining what the extra parameters are [2], whereas a source-level analysis must either make a cruder over-approximation or an unsound under-approximation.

A final example is shown in Fig. 1. The C code on the left uses an uninitialized variable (which triggers a compiler warning, but compiles successfully). A source-code analyzer must assume that `local` can have any value, and therefore the value of `v` in `main` is either 1 or 2. The assembly listings on the right show how the C code could be compiled, including two variants for the prolog of function `callee`. The Microsoft compiler (`cl`) uses the second variant, which includes the following strength reduction:

The instruction `sub esp,4` that allocates space for `local` is replaced by a `push` instruction of an arbitrary register (in this case, `ecx`).

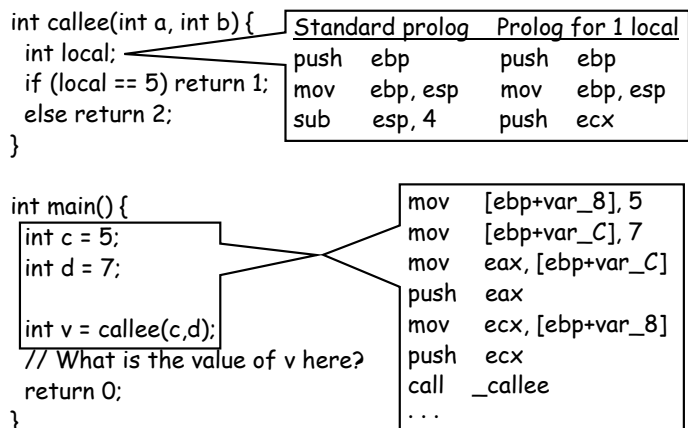


Fig. 1. Example of unexpected behavior due to the application of an optimization. The box at the top right shows two variants of code generated by an optimizing compiler for the prolog of `callee`. Analysis of the second of these reveals that the variable `local` always contains the value 5.

In contrast to an analysis based on source code, an analysis of an executable can determine that this optimization results in `local` being initialized to 5, and therefore `v` in `main` can only have the value 1.

3 Recovering Intermediate Representations from x86 Executables

To recover IRs from x86 executables, CodeSurfer/x86 makes use of both IDAPro [19], a disassembly toolkit, and GrammaTech’s CodeSurfer system [12], a toolkit for building program-analysis and inspection tools. Fig. 2 shows how the components of CodeSurfer/x86 fit together.

An x86 executable is first disassembled using IDAPro. In addition to the disassembly listing, IDAPro also provides access to the following information: (1) procedure boundaries, (2) calls to library functions using an algorithm called the Fast Library Identification and Recognition Technology (FLIRT) [15], and (3) statically known memory addresses and offsets. IDAPro provides access to its internal resources via an API that allows users to create plug-ins to be executed by IDAPro. We created a plug-in to IDAPro, called the Connector, that creates data structures

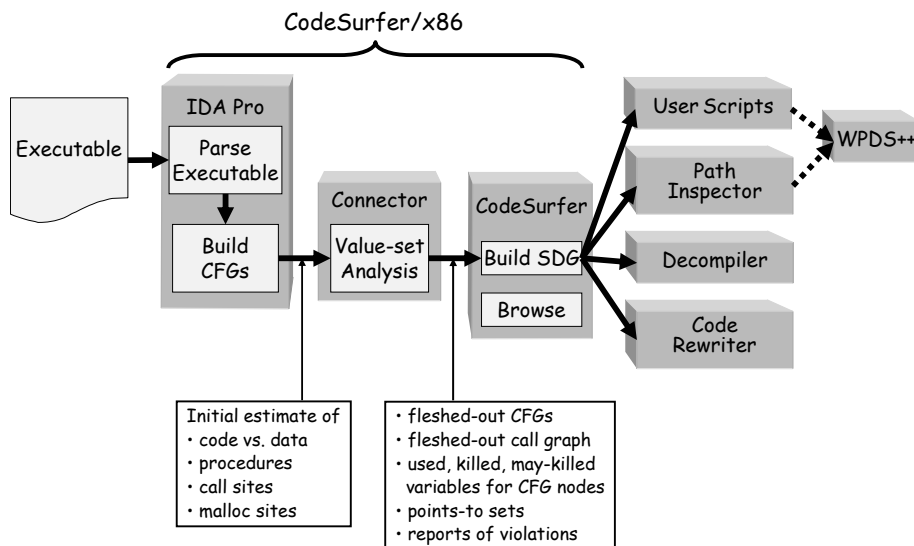


Fig. 2. Organization of CodeSurfer/x86 and companion tools.

to represent the information that it obtains from IDAPro. The IDAPro/Connector combination is also able to create the same data structures for dynamically linked libraries, and to link them into the data structures that represent the program itself. This infrastructure permits whole-program analysis to be carried out—including analysis of the code for all library functions that are called.

Using the data structures in the Connector, we implemented a static-analysis algorithm called *value-set analysis* (VSA) [2]. VSA does not assume the presence of symbol-table or debugging information. Hence, as a first step, a set of data objects called a-locs (for “abstract locations”) is determined based on the static memory addresses and offsets provided by IDAPro. VSA is a combined numeric and pointer-analysis algorithm that determines an over-approximation of the set of numeric values and addresses (or *value-set*) that each a-loc holds at each program point. A key feature of VSA is that it tracks integer-valued and address-valued quantities simultaneously. This is crucial for analyzing executables because numeric values and addresses are indistinguishable at execution time.

IDAPro does not identify the targets of all indirect jumps and indirect calls, and therefore the call graph and control-flow graphs that it constructs are not complete. However, the information computed during VSA can be used to augment the call graph and control-flow graphs on-the-fly to account for indirect jumps and indirect calls.

VSA also checks whether the executable conforms to a “standard” compilation model—i.e., a runtime stack is maintained; activation records (ARs) are pushed onto the stack on procedure entry and popped from the stack on procedure exit; a procedure does not modify the return address on stack; the program’s instructions occupy a fixed area of memory, are not self-modifying, and are separate from the program’s data. If it cannot be confirmed that the executable conforms to the model, then the IR is possibly incorrect. For example, the call-graph can be incorrect if a procedure modifies the return address on the stack. Consequently, VSA issues an error report whenever it finds a possible violation of the standard compilation model; these represent possible memory-safety violations. The analyst can go over these reports and determine whether they are false alarms or real violations.

Once VSA completes, the value-sets for the a-locs at each program point are used to determine each point’s sets of used, killed, and possibly-killed a-locs; these are emitted in a format that is suitable for input to CodeSurfer. CodeSurfer then builds a collection of IRs, consisting of abstract-syntax trees, control-flow graphs (CFGs), a call graph, a system dependence graph (SDG) [17], VSA results, the sets of used, killed, and possibly killed a-locs at each instruction, and information about the structure and layout of global memory, activation records, and dynamically allocated storage. CodeSurfer supports both a graphical user interface (GUI) and an API (as well as a scripting language) to provide access to these structures.

4 Model-Checking Facilities

Model checking [11] involves the use of sophisticated pattern-matching techniques to answer questions about the flow of execution in a program: a model of the program’s possible behavior is created and checked for conformance with a model of expected behavior (as specified by a user query). In essence, model-checking algorithms explore the program’s state-space and answer questions about whether a bad state can be reached during an execution of the program.

For model checking, the CodeSurfer/x86 IRs are used to build a *weighted pushdown system* (WPDS) [5, 26, 27, 21] that models possible program behaviors. Weighted pushdown systems are a model-checking technology that is similar to what is used in MOPS, a software model checker that has demonstrated the ability to find subtle security vulnerabilities in large code bases [6]. However, in contrast to the ordinary (unweighted) pushdown systems used in MOPS, the techniques available in the CodeSurfer/x86 tool set [27, 21] are capable of representing the (logically) infinite set of data valuations that may arise during program execution. This capability allows the CodeSurfer/x86 tool set to address certain kinds of security queries that cannot be answered by MOPS.

4.1 WPDS++ and Stack-Qualified Dataflow Queries

WPDS++ [20] is a library that implements the symbolic reachability algorithms from [27, 21] on weighted pushdown systems. We follow the standard approach of using a pushdown system (PDS) to model the interprocedural control-flow graph (one of CodeSurfer/x86’s IRs). The stack symbols correspond to program locations; there is only a single PDS state; and PDS rules encode control flow as follows:

Rule	Control flow modeled
$q\langle u \rangle \hookrightarrow q\langle v \rangle$	Intraprocedural CFG edge $u \rightarrow v$
$q\langle c \rangle \hookrightarrow q\langle \text{entry}_P \ r \rangle$	Call to P from c that returns to r
$q\langle x \rangle \hookrightarrow q\langle \rangle$	Return from a procedure at exit node x

Given a configuration of the PDS, the top stack symbol corresponds to the current program location, and the rest of the stack holds return-site locations—much like a standard run-time execution stack.

Encoding the interprocedural control-flow as a pushdown system is sufficient for answering queries about reachable control states (as the Path Inspector does; see Sect. 4.2): the reachability algorithms of WPDS++ can determine if an undesirable PDS configuration is reachable [6]. However, WPDS++ also supports *weighted* PDSs. These are PDSs in which each rule is weighted with an element of a (user-defined) semiring. The use of weights allows WPDS++ to perform interprocedural dataflow analysis by using the semiring’s *extend* operator to compute weights for sequences of rule firings and using the semiring’s *combine* operator to take the meet of weights generated by different paths [27, 21]. (When the weights on rules are conservative abstract data transformers, an over-approximation to the set of reachable concrete configurations is obtained, which means that counterexamples reported by WPDS++ may actually be infeasible.)

The CodeSurfer/x86 IRs are a rich source of opportunities to check properties of interest using WPDS++. For instance, WPDS++ has been used to implement an illegal-stack-manipulation check: for each node n in procedure P , this checks whether the net change in stack height is the same along all paths from entry_P to n that have perfectly matched calls and returns (i.e., along “same-level valid paths”). In this analysis, a weight is a function that represents a stack-height change. For instance, `push ecx` and `sub esp, 4` both have the weight $\lambda \text{height. height} - 4$. `Extend` is (the reversal of) function composition; `combine` performs a meet of stack-height-change functions. (The analysis is similar to linear constant propagation [29].) When a memory access performed relative to r ’s activation record (AR) is out-of-bounds, stack-height-change values can be used to identify which a-locs could be accessed in ARs of other procedures.

VSA is an interprocedural dataflow-analysis algorithm that uses the “call-strings” approach [30] to obtain a degree of context sensitivity. Each dataflow fact is tagged with a call-stack suffix (or *call-string*) to form (call-string, dataflow-fact) pairs; the call-string is used at the exit node of each procedure to determine to which call site a (call-string, dataflow-fact) pair should be propagated. The call-strings that arise at a given node n provide an opportunity to perform stack-qualified dataflow queries [27] using WPDS++. CodeSurfer/x86 identifies induction-variable relationships by using the affine-relation domain of Müller-Olm and Seidl [24] as a weight domain. A *post** query builds an automaton that is then used to find the affine relations that hold in a given calling context—given by call-string cs —by querying the *post**-automaton with respect to a regular language constructed from cs and the program’s call graph.

4.2 The Path Inspector

The Path Inspector provides a user interface for automating safety queries that are only concerned with the possible control configurations that an executable can reach. It uses an automaton-based approach to model checking: the query is specified as a finite automaton that captures forbidden sequences of program locations. This “query automaton” is combined with the program model (a WPDS) using a cross-product construction, and the reachability algorithms of WPDS++ are used to determine if an error configuration is reachable. If an error configuration is reachable, then *witnesses* (see [27]) can be used to produce a program path that drives the query automaton to an error state.

The Path Inspector includes a GUI for instantiating many common reachability queries [14], and for displaying counterexample paths in the disassembly listing.³ In the current implementation, transitions in the query automaton are triggered by program points that the user specifies either manually, or using result sets from CodeSurfer queries. Future versions of the Path Inspector will support more sophisticated queries in which transitions are triggered by matching an AST pattern against a program location, and query states can be instantiated based on pattern bindings. Future versions will also eliminate (many) infeasible counterexamples by using transition weights to represent abstract data transformers (similar to those used for interprocedural dataflow analysis).

5 Related Work

Several others have proposed techniques to obtain information from executables by means of static analysis [22, 13, 9, 8, 10, 4, 3]. However, previous techniques deal with memory accesses very conservatively; e.g., if a register is assigned a value from memory, it is assumed to take on any value. VSA does a much better job than previous work because it tracks the integer-valued and address-valued quantities that the program’s data objects can hold; in particular, VSA tracks the values of data objects *other than just the hardware registers*, and thus is not forced to give up all precision when a load from memory is encountered. This is a fundamental issue; the absence of such information places severe limitations on what previously developed tools can be applied to.

The basic goal of the algorithm proposed by Debray et al. [13] is similar to that of VSA: for them, it is to find an over-approximation of the set of values that each *register* can hold at each program point; for us, it is to find an over-approximation of the set of values that each (abstract) data object can hold at each program point, where data objects include *memory locations* in addition to registers. In their analysis, a set of addresses is approximated by a set of congruence values: they keep track of only the low-order bits of addresses. However, unlike VSA, their algorithm does not make any effort to track values that are not in registers. Consequently, they lose a great deal of precision whenever there is a load from memory.

Cifuentes and Fraboulet [9] give an algorithm to identify an intraprocedural slice of an executable by following the program’s use-def chains. However, their algorithm also makes no attempt to track values that are not in registers, and hence cuts short the slice when a load from memory is encountered.

The work that is most closely related to VSA is the data-dependence algorithm for executables proposed by Amme et al. [1]. However, that algorithm performs only an intraprocedural analysis, and it is not clear that the algorithm fully accounts for dependences between memory locations.

Several people have developed techniques to analyze executables in the presence of additional information, such as the source code, symbol-table information, or debugging information [22, 28]. For many security-related applications, these are not appropriate assumptions.

Christodorescu and Jha used model-checking techniques to detect malicious-code variants [7]. Given a sample of malicious code, they extract a parameterized state machine that will accept variants of the code. They use CodeSurfer/x86 (with VSA turned off) to extract a model of each program procedure, and determine potential matches between the program’s code and fragments of the malicious code. Their technique is intraprocedural, and does not analyze data state.

Other groups have used run-time program monitoring and checkpointing to perform a systematic search of a program’s dynamic state space [16, 23, 25]. Like our approach, this allows for model checking properties of the low-level code that is actually run on the machine. However, because the dynamic state space can be unbounded, these approaches cannot perform an exhaustive search. In contrast, we use static analysis to perform a (conservative) exhaustive search of an abstract state space.

³ We assume that source code is not available, but the techniques extend naturally if it is: one can treat the executable code as just another IR in the collection of IRs obtainable from source code. The mapping of information back to the source code would be similar to what C source-code tools already have to perform because of the use of the C preprocessor.

References

1. W. Amme, P. Braun, E. Zehendner, and F. Thomasset. Data dependence analysis of assembly code. *Int. J. on Parallel Processing*, 2000.
2. G. Balakrishnan and T. Reps. Analyzing memory accesses in x86 executables. In *Comp. Construct.*, pages 5–23, 2004.
3. J. Bergeron, M. Debbabi, J. Desharnais, M.M. Erhioui, Y. Lavoie, and N. Tawbi. Static detection of malicious code in executable programs. *Int. J. of Req. Eng.*, 2001.
4. J. Bergeron, M. Debbabi, M.M. Erhioui, and B. Ktari. Static analysis of binary code to isolate malicious behaviors. In *WETICE*, pages 184–189, 1999.
5. A. Bouajjani, J. Esparza, and T. Touili. A generic approach to the static analysis of concurrent programs with procedures. In *Princ. of Prog. Lang.*, pages 62–73, 2003.
6. H. Chen, D. Dean, and D. Wagner. Model checking one million lines of C code. In *Network and Dist. Syst. Security*, 2004.
7. M. Christodorescu and S. Jha. Static analysis of executables to detect malicious patterns. In *USENIX Security Symposium*, 2003.
8. C. Cifuentes and A. Fraboulet. Interprocedural data flow recovery of high-level language code from assembly. Technical Report 421, Univ. Queensland, 1997.
9. C. Cifuentes and A. Fraboulet. Intraprocedural static slicing of binary executables. In *Int. Conf. on Softw. Maint.*, pages 188–195, 1997.
10. C. Cifuentes, D. Simon, and A. Fraboulet. Assembly to high-level language translation. In *Int. Conf. on Softw. Maint.*, pages 228–237, 1998.
11. E.M. Clarke, Jr., O. Grumberg, and D.A. Peled. *Model Checking*. The M.I.T. Press, 1999.
12. CodeSurfer, GrammaTech, Inc., <http://www.grammatech.com/products/codesurfer/>.
13. S.K. Debray, R. Muth, and M. Weippert. Alias analysis of executable code. In *Princ. of Prog. Lang.*, pages 12–24, 1998.
14. M. Dwyer, G. Avrunin, and J. Corbett. Patterns in property specifications for finite-state verification. In *Int. Conf. on Softw. Eng.*, 1999.
15. Fast library identification and recognition technology, DataRescue sa/nv, Liège, Belgium, <http://www.datarescue.com/idabase/flrt.htm>.
16. P. Godefroid. Model checking for programming languages using VeriSoft. In ACM, editor, *Princ. of Prog. Lang.*, pages 174–186. ACM Press, 1997.
17. S. Horwitz, T. Reps, and D. Binkley. Interprocedural slicing using dependence graphs. *Trans. on Prog. Lang. and Syst.*, 12(1):26–60, January 1990.
18. M. Howard. Some bad news and some good news, October 2002. <http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dncode/html/secure10102002.asp>.
19. IDAPro disassembler, <http://www.datarescue.com/idabase/>.
20. N. Kidd, T. Reps, D. Melski, and A. Lal. WPDS++: A C++ library for weighted pushdown systems. <http://www.cs.wisc.edu/wpis/wpds++/>.
21. A. Lal, T. Reps, and G. Balakrishnan. Extended weighted pushdown systems. In *Computer Aided Verif.*, 2005.
22. J.R. Larus and E. Schnarr. EEL: Machine-independent executable editing. In *Prog. Lang. Design and Impl.*, pages 291–300, 1995.
23. P. Leven, T. Mehler, and S. Edelkamp. Directed error detection in C++ with the assembly-level model checker StEAM. In *Spin Workshop*, 2004.
24. M. Müller-Olm and H. Seidl. Analysis of modular arithmetic. In *European Symp. on Programming*, 2005.
25. M. Musuvathi, D. Park, A. Chou, D. Engler, and D. Dill. CMC: A pragmatic approach to model checking real code. In *Op. Syst. Design and Impl.*, 2002.
26. T. Reps, S. Schwoon, and S. Jha. Weighted pushdown systems and their application to interprocedural dataflow analysis. In *Static Analysis Symp.*, 2003.
27. T. Reps, S. Schwoon, S. Jha, and D. Melski. Weighted pushdown systems and their application to interprocedural dataflow analysis. *Sci. of Comp. Prog.* To appear.
28. X. Rival. Abstract interpretation based certification of assembly code. In *Int. Conf. on Verif., Model Checking, and Abs. Int.*, 2003.
29. M. Sagiv, T. Reps, and S. Horwitz. Precise interprocedural dataflow analysis with applications to constant propagation. *Theor. Comp. Sci.*, 167:131–170, 1996.
30. M. Sharir and A. Pnueli. Two approaches to interprocedural data flow analysis. In S.S. Muchnick and N.D. Jones, editors, *Program Flow Analysis: Theory and Applications*, chapter 7, pages 189–234. Prentice-Hall, Englewood Cliffs, NJ, 1981.