

Bamberger Beiträge zur  
Wirtschaftsinformatik und Angewandten Informatik

ISSN 0937-3349

Nr. 85

# Symbolic Object Code Analysis

Jan Tobias Mühlberg und Gerald Lüttgen

Februar 2010

Fakultät Wirtschaftsinformatik und Angewandte Informatik  
Otto-Friedrich-Universität Bamberg

# Symbolic Object Code Analysis

Jan Tobias Mühlberg and Gerald Lüttgen

Software Technologies Research Group  
University of Bamberg, 96045 Bamberg, Germany.  
{jan-tobias.muehlberg, gerald.luetttgen}@swt-bamberg.de

**Abstract** Current software model checkers quickly reach their limit when being applied to verifying pointer safety properties in source code that includes function pointers and inlined assembly. This paper introduces an alternative technique for checking pointer safety violations, called *Symbolic Object Code Analysis* (SOCA), which is based on bounded symbolic execution, incorporates path-sensitive slicing, and employs the SMT solver Yices as its execution and verification engine. Extensive experimental results of a prototypic SOCA Verifier, using the Verisec suite and almost 10,000 Linux device driver functions as benchmarks, show that SOCA performs competitively to current source-code model checkers and that it also scales well when applied to real operating systems code and pointer safety issues. SOCA effectively explores semantic niches of software that current software verifiers do not reach.

**Keywords:** model checking; symbolic execution; program slicing; object code analysis; linux device drivers

## 1 Introduction

One challenge in verifying complex software is the proper analysis of pointer operations. A recent study shows that a majority of errors found in device drivers involve *memory safety* [9]. Writing software that is free of memory safety concerns, e.g., free of errors caused by pointers to invalid memory cells, is difficult since many such issues result in program crashes at later points in execution. Hence, a statement causing a memory corruption may not be easily identifiable using conventional validation and testing tools, e.g., *Purify* [36] and *Valgrind* [33].

Current static verification tools, including *software model checkers* such as [4,10,12,17], are also not of much help: they either assume that programs do “not have wild pointers” [3], perform poorly in the presence of pointers [31], or simply cannot handle certain software. A particular challenging kind of software are operating systems (OS) components such as device drivers, which are usually written in C code involving function pointers, pointer arithmetic and inlined assembly. Further issues arise because of platform-specific and compiler-specific details concerning memory layout, padding and offsets [2]. In addition, several approaches to model checking *compiled programs* given in assembly or bytecode

[6,29,38,40], and also to integrating *symbolic execution* [23] with model checking [15,16,22,35,39] have recently been presented. However, these are tailored to exploit specific characteristics of certain programming paradigms such as object-oriented programming, or lack support for data structures, function pointers and computed jumps, or require substantial manual modelling effort (cf. Sec. 5).

Our contributions presented in this paper are as follows. We introduce a novel, automated technique to identifying memory safety violations, called *Symbolic Object Code Analysis* (SOCA), which is based on the *symbolic execution* [23] of compiled and linked programs (cf. Sec. 2). In contrast to other verification techniques, SOCA requires only a minimum of manual modelling effort, namely the abstract, symbolic specification of a program’s execution context in terms of function inputs and initial heap content. The SOCA technique traverses the program’s *object code* in a systematic fashion up to a certain depth and width, and calculates at each assembly instruction a *slice* [41] required for checking the relevant pointer safety properties. It translates such a slice and properties into a bit-vector constraint problem and executes the property checks by invoking the *Yices* SMT solver [14] (cf. Sec. 3). To the best of our knowledge, SOCA is the only program verification technique available that features full support for pointer arithmetics, function pointers and computed jumps. While SOCA is based on existing and well-known techniques, combining and implementing these for object code analysis is challenging. Much engineering effort went into our SOCA implementation, so that it scales to complex real-world OS code such as Linux device drivers. We believe that our experience gained will be helpful for future developers of program analysis and program verification tools (cf. Sec. 3).

The particular combination of techniques in SOCA is well suited for checking memory safety. Analysing object code is beneficial in that it inherently considers compiler specifics such as results of code optimisations, makes memory layout obvious, and does away with the challenge of handling mixed input languages involving assembly code. Symbolic execution, rather than the concrete execution adopted in testing, can handle software functions with many input parameters, whose values are typically not known at compile time. It is the existence of efficient SMT solvers that makes the symbolic approach feasible. Symbolic execution also implies a path-wise exploration, thus reducing the aliasing problem and allowing us to handle even complex pointer operations and computed jumps. In addition, slicing can now be conducted at path-level instead of at program-level, resulting in drastically smaller slices to the extent that abstraction is not necessary for achieving scalability. However, the price of symbolic execution is that it must be bounded and can thus only analyse code up to a finite depth and width.

Interesting questions regarding the SOCA technique are whether it is competitive to state-of-the-art model checking on programs with “well-behaved” pointers such as [4,12,17], and whether it scales when applied to “dirty” programs such as device drivers, which cannot be properly analysed with source-code model checkers. To answer these questions, we have implemented a prototypic SOCA tool for programs compiled for the 32-bit Intel Architecture (IA32), the *SOCA*

*Verifier*, and performed extensive experiments (cf. Sec. 4). Using the *Verisec* benchmark [27] we show that the SOCA Verifier performs on par with the model checkers *LoopFrog* [25] and *SatAbs* [12] with regards to performance, error detection and false-positive rates. We then apply the SOCA Verifier to 9296 functions taken from 250 Linux device drivers. Our tool is able to successfully analyse 95% of these functions and, despite the fact that SOCA performs a bounded analysis, 28% of the functions are analysed exhaustively. Therefore, SOCA proves itself to be a capable technique when being confronted with checking pointer-complex software such as OS components. It effectively explores semantic niches that neither current testing tools nor current software model checkers reach.

## 2 Pointers, Aliasing & Intermediate Representation

The verification technique developed in this paper aims at ensuring that every pointer in a given program is valid in the sense that it *(i)* never references a memory location outside the address space allocated by or for that program, and *(ii)* respects the usage rules determined by the Application Programming Interfaces (APIs) employed by the program. There exist several categories of memory safety properties — *(1) dereferencing invalid pointers*: a pointer may not be NULL, shall be initialised, and shall not point to a memory location outside the address space allocated by or for the program; *(2) uninitialised reads*: memory cells shall be initialised before they are read; *(3) violation of memory permissions*: when the program is loaded into memory, the segments of the program file are assigned with permissions that determine whether a segment can be read, written or executed; *(4) buffer overflows*: out-of-bounds read and write operations to objects on the heap and stack, which may lead to memory corruption and give way to various security problems; *(5) memory leaks*: when a program dynamically allocates memory but loses the handle to it, the memory cannot be deallocated anymore; *(6) proper handling of allocation and deallocation*: OSs usually provide several APIs for the dynamic (de)allocation of memory, whose documentation specifies precisely what pairs of functions are to be employed how.

```

01 #include <stdio.h>                                08 for (*p1=0; *p1<10; (*p1)++)
02 #include <sys/types.h>                            09 { *p2=0; }
03                                                    10
04 int main (void) {                                11 printf ("%08x: %d\n", p1, *p1);
05  int32_t i, *p2=&i;                                12 printf ("%08x: %d\n", p2, *p2);
06  int16_t *p1=&((int16_t*) &i)[0];                 13 printf ("%08x: %d\n", &i, i);
07                                                    14 return (0); }

```

**Figure 1.** Example of pointer aliasing in C: `e_loop.c`.

***Aliasing in source code and object code.*** A major issue for analysing pointer programs is aliasing. Aliasing means that a data location in memory may be ac-

cessed through different symbolic names. Since aliasing relations between symbolic names and data locations often arise unexpectedly during program execution, they may result in erroneous program behaviours that are particularly hard to trace and debug. To illustrate this, the C program given in Fig. 1 shows a complicated way of implementing an endless loop.

\$ gcc -O2 e_loop.c	\$ gcc -O2 e_loop.c	\$ gcc -O1 e_loop.c
\$ ./a.out	\$ ./a.out	\$ ./a.out
bfc76f2c: 10	bfc7428c: 10	-> does not terminate
bfc76f2c: 0	bfc7428c: 10	
bfc76f2c: 0	bfc7428c: 10	

**Figure 2.** Output of the program given in Fig. 1 when compiled with (a) gcc version 4.1.2 (left) and (b,c) gcc version 4.3.1 (middle and right).

```

80483ba: xor    %eax,%eax    ;; eax := 0;
80483c4: lea    -0xc(%ebp),%ebx ;; ebx := ebp - 0xc
80483c8: add    $0x1,%eax    ;; eax := eax + 0x00000001
80483cb: cmp    $0x9,%eax    ;; (ax = 9)?
80483cf: movl   $0x0,-0xc(%ebp) ;; *p2 (= ebp - 0xc) := 0
80483d6: mov    %ax,(%ebx)    ;; *p1 (= ebx = ebp - 0xc) := ax
80483d9: jle    80483c8        ;; if (ax <= 9) goto 80483c8

```

**Figure 3.** Excerpt of the disassembled program compiled in Fig. 2(b).

As depicted in Fig. 2, three different outcomes of the program’s execution can occur as a result of varying assumptions made about pointer aliasing by the developer and the compiler, as well as compiler optimisations applied to the code. More surprises are revealed when studying the excerpt of the corresponding assembly code displayed in Fig. 3, which was obtained by disassembling the program that produced the output shown in Fig. 2(b). One can see at instructions 80483cf and 80483d6 that *p1* and *p2* are pointing to the same location in memory, and that *\*p2* is actually written before *\*p1*. This is unexpected when looking at the program’s source code but valid from the compiler’s point of view, since it assumes that the two pointers are pointing to different data objects. As another consequence of this assumption, register *eax* is never reloaded from the memory location to which *p1* and *p2* point.

This example shows that source-code-based analysis has to decide for a particular semantics of the source language, which may not be the one that is used by a compiler. Hence, results obtained by analysing the source code may not meet a program’s runtime behaviour. While this motivates the analysis of compiled programs, doing so does not provide a generic solution for dealing with pointer aliasing, as aliasing relationships may depend on runtime conditions.

**Intermediate representation.** A program under analysis is stored by us in an intermediate representation (IR) borrowed from *Valgrind* [33], a framework for dynamic binary instrumentation. The IR consists of a set of *basic blocks* containing a group of statements such that all transfers of control to the block

are to the first statement in the group. Once the block has been entered, all of its statements are executed sequentially until an *exit* statement is reached. An exit is always denoted as `goto <target>`, where `<target>` is either a constant or a temporary register that determines the next program location to be executed. Guarded jumps are written as `if (<condition>) goto <target>`, where `<condition>` is a temporary register of type boolean, which has previously been assigned within the basic block.

IA32 Assembly	IR Instructions
<code>xor %eax,%eax</code>	<code>t9 = GET:I32(0) ;; t9 := eax</code> <code>t8 = GET:I32(0) ;; t8 := eax</code> <code>t7 = Xor32(t9,t8) ;; t7 := t9 xor t8</code> <code>PUT(0) = t7 ;; eax := t7</code>
<code>lea -0xc(%ebp),%ebx</code>	<code>t42 = GET:I32(20)</code> <code>t41 = Add32(t42,0xFFFFFFFF4:I32)</code> <code>PUT(12) = t41</code>

**Figure 4.** First two instructions of Fig. 3 and their respective IR instructions.

Fig. 4 depicts an example for assembly statements and their corresponding IR statements. It shows how, e.g., the `xor` statement is decomposed into explicitly loading (GET) the source register 0 into the temporary registers `t8` and `t9`, performing the `xor` operation into the temporary register `t7`, followed by storing (PUT) the result back to the guest state. All operands used in the first block of the example are 4 bytes, or 32 bits, in size.

As can be seen, the IR is essentially a typed assembly language in static-single-assignment form [28], and employs *temporary registers*, which are denoted as  $t\langle n \rangle$ , and the *guest state*. The guest state consists of the contents of the registers that are available in the architecture for which the program under analysis is compiled. While machine registers are always 8 bits long, temporary registers may be 1, 8, 16, 32 or 64 bits in length. As a result of this, statement `t9 = GET:I32(0)` means that  $t9$  is generated by concatenating machine registers 0 to 3. Since each IR block is in static-single-assignment form with respect to the temporary registers,  $t9$  is assigned only once within a single IR block. As a valuable feature for analysing pointer safety, Valgrind’s IR makes all load and store operations to memory cells explicit.

### 3 SOCA – Symbolic Object Code Analysis

This section introduces our novel approach to verifying memory safety in compiled and linked programs, to which we refer as *Symbolic Object Code Analysis* (SOCA). The basic idea behind our approach employs well-known techniques including *symbolic execution* [23], *SMT solving* [26] and *program slicing* [41]. However, combining these ideas and implementing them in a way that scales

to real applications, such as Linux device drivers, is challenging and the main contribution of this paper.

Starting from a program’s given entry point, we automatically translate each instruction of the program’s *object code* into Valgrind’s IR language. This is done lazily, i.e., as needed, by iteratively following each program path in a depth-first fashion and resolving target addresses of computed jumps and return statements. We then generate systems of bit-vector constraints for the path under analysis, which reflect the path-relevant register content and heap content of the program. In this process we employ a form of program slicing, called *path-sensitive and heap-aware program slicing* (cf. p. 8), which is key to SOCA’s scalability and makes program abstraction unnecessary. Finally, we invoke the SMT solver *Yices* [14] to check the satisfiability of the resulting constraint systems and thus the validity of the path. This approach allows us to instrument the constraint systems on-the-fly as necessary, by adding constraints that express, e.g., whether a pointer points to an allocated address.

SOCA leaves most of a program’s input and initial heap content unspecified in order to allow the SMT solver to search for subtle inputs that may reveal pointer errors. Obviously, our analysis by symbolic execution cannot be complete: the search space has to be bounded since the total number of execution paths and the number of instructions per path may be infinite. Our experimental results (cf. Sec 4) show that this boundedness is not a restriction in practice: many interesting programs, such as Linux device driver functions, are relatively “shallow” and may still be analysed either exhaustively or to an acceptable extent.

**Translating IR into Yices constraints.** To translate IR statements into bit-vector constraint systems for Yices, we have defined a simple operational semantics for Valgrind’s IR language. Due to space constraints we cannot present this semantics here and refer the reader to [30] instead. Instead, we focus directly on examples illustrating this translation.

As a first example we consider the  $PUT(0) = t7$  statement from Fig. 4. Intuitively, the semantics of *PUT* is to store the value held by *t7* to the guest state, in register 0 to 3 (i.e., *r0* to *r3* below):

IR Instruction	Constraint Representation
$PUT(0) = t7$	<pre>(define r0::(bitvector 8)(bv-extract 31 24 t7)) (define r1::(bitvector 8)(bv-extract 23 16 t7)) (define r2::(bitvector 8)(bv-extract 15 8 t7)) (define r3::(bitvector 8)(bv-extract 7 0 t7))</pre>

Here, the **bv-extract** operation denotes bit-vector extraction. Note that the IA32 CPU registers are assigned in reverse byte order, while arithmetic expressions in Yices are implemented for bit-vectors that have their most significant bit at position 0. Since access operations to the guest state may be 8, 16, 32 or 64 bit aligned, we have to translate the contents of temporary register when accessing the guest state.

Similar to the *PUT* instruction, we can express *GET*, i.e., loading a value from the guest state, as the concatenation of bit-vectors, and the *Xor* and *Add* instructions in terms of bit-vector arithmetic:

IR Instruction	Constraint Representation
t9 = GET:I32(0)	(define t9::(bitvector 32) (bv-concat (bv-concat r3 r2) (bv-concat r1 r0)))
t7 = Xor32(t9,t8)	(define t7::(bitvector 32) (bv-xor t9 t8))
t41 = Add32(t42, 0xFFFFFFFF4:I32)	(define t88::(bitvector 32) (bv-add t87 (mk-bv 32 4294967284))

More challenging to implement are the IR instructions *ST* (store) and *LD* (load) which facilitate memory access. The main difference of these instructions to *PUT* and *GET* is that the target of *ST* and the source of *LD* are variable and may only be computed at runtime. To include these statements in our framework we have to express them in a flexible way, so that the SMT solver can identify cases in which safety properties are violated. In Yices we declare a function *heap* as our representation of the program's memory. An exemplary *ST* statement *ST*(t5) = t32 can be expressed in terms of updates of that function:

IR Instruction	Constraint Representation
ST(t5) = t32	(define heap::(-> (bitvector 32) (bitvector 8)))  (define heap.0::(-> (bitvector 32) (bitvector 8)) (update heap ((bv-add t5 (mk-bv 32 3))) (bv-extract 7 0 t32))) (define heap.1::(-> (bitvector 32) (bitvector 8)) (update heap.0 ((bv-add t5 (mk-bv 32 2))) (bv-extract 15 8 t32))) (define heap.2::(-> (bitvector 32) (bitvector 8)) (update heap.1 ((bv-add t5 (mk-bv 32 1))) (bv-extract 23 16 t32))) (define heap.3::(-> (bitvector 32) (bitvector 8)) (update heap.2 ((bv-add t5 (mk-bv 32 0))) (bv-extract 31 24 t32)))

Since the above *ST* instruction stores the content of a 32-bit variable in four separate 8-bit memory cells, we have to perform four updates of *heap*. Byte-ordering conventions apply in the same way as explained for *PUT*. Constraints for the *LD* instruction are generated analogous to *GET*.

**Encoding pointer safety assertions.** Being able to translate each object code instruction into constraints allows us to express the safety pointer properties given in Sec. 2 in terms of assertions within the constraint systems. The simplest case of such an assertion is a null-pointer check. For the *ST* instruction in the above example, we state this assertion as (assert (= t5 (mk-bv 32 0))).



If the resulting constraint system is satisfiable, Yices will return a possible assignment to the constraint system variables representing the program’s input. This input is constructed such that it will drive the program into a state in which  $t5$  holds the value null at the above program point.

However, many memory safety properties require additional information to be known about the program’s current execution context. In particular, answering the question whether a pointer may point to an “invalid” memory area requires us to know which cells are currently allocated. We retain this information by adding a function named *heaploc* to our memory representation:

```
(define heaploc::(-> (bitvector 32) (record alloc::bool init::bool
  start::(bitvector 32) size::(bitvector 32))))
```

This allows us to express assertions stating that, e.g., pointer  $t5$  has to point to an allocated address at the program location where it is dereferenced, as:

```
(assert (= (select (heaploc t5) alloc) false))
```

All other pointer safety properties mentioned in Sec. 2 may be encoded along the lines of those two examples. Most of them require further additional information to be added to the *heaploc* function. To reduce the size and search space of the resulting constraint systems we check assertions one-by-one with a specialised *heaploc* function for each property. The full details on our generation of constraint systems can be found in [30].

**Path-sensitive slicing.** To ensure scalability of our SOCA technique, we do not run Yices on an entire path’s constraint system. Instead we compute a slice [41] of the constraint system containing only those constraints that are relevant to the property to be checked at a particular program location.

The approach to path-sensitive program slicing in SOCA employs an algorithm based on system dependence graphs as introduced in [19]. Our slices are computed using conventional slicing criteria  $(L, var)$  denoting a variable  $var$  that is used at program location  $L$ , but over the single path currently being analysed instead of the program’s entire control flow. The slice is then computed by collecting all statements on which  $var$  is data dependent by tracing the path backwards, starting from  $L$  up to the program entry point. While collecting flow dependencies is relatively easy for programs that do only use CPU registers and temporary registers, it becomes difficult when dependencies to the heap and stack are involved.

**Handling memory access in slicing.** Consider the following two IR statements: 01 ST( $t5$ ) =  $t32$ ; 02  $t31$  = LD:I32( $t7$ ). To compute a slice for the slicing criterion (02,  $t31$ ) we have to know whether the store statement ST may affect the value of  $t31$ , i.e., whether  $t5$  and  $t7$  may alias. We obtain this information by using Yices to iteratively compute the potential address range that can be accessed via  $t5$ . This is done by making Yices find a satisfying model  $e$  for  $t5$ . When reading a model, which is represented by Yices by a bit-vector, we compute its integer representation. We then compute further satisfying models

$e'$  such that  $e > e'$  or  $e < e'$  holds, until the range is explored. The computation is done by stepwise adding or retracting constraints so as to use Yices as efficiently as possible. However, sometimes complex constraint systems arise that require the power and efficiency of modern SAT solvers. Since we remember only the maximal and minimal satisfying models for a given pointer, this is an over-approximation as not the entire address range may be addressable by that pointer. However, using this abstraction presents a trade-off concerning only the size of the computed slices and not their correctness, and helps us to keep the number of Yices runs and the amount of data to be stored small.

By computing the potential address range accessed by a pointer used in a load statement, e.g., `t7` in our case, and looking for memory intervals overlapping with the range of `t7`, we can now determine which store operations may affect the result of the load operation. Despite being conservative when computing address ranges, our experience shows that most memory access operations end up having few dependencies; this is because most pointers evaluate to a concrete value, i.e., the constraint system has exactly one satisfying model, rather than a symbolic value.

**Handling computed jumps.** A major challenge when analysing compiled programs arises from the extensive use of function pointers and jump target computations. While most source-code-based approaches simply ignore function pointers [4,12,17], this cannot be done when analysing object code since jump computations are too widely deployed here. The most common example for a computed jump is the *return* statement after a function call. To perform a return, the bottom element of the stack is loaded into a temporary register, e.g., `t1`, followed by a `goto t1` statement, which effectively sets the value of the program counter to `t1`. Further examples for computed jumps are jump tables and function pointers.

In our approach, jump target addresses are determined in the same way as addresses for load and store operations, i.e., by computing a slice for each jump target and then using Yices to determine satisfying models for the target register.

**Optimising GET and PUT statements.** A problem with respect to the scalability of our approach arises from the vast number of *GET* and *PUT* statements in IR code. In particular, the frequent de-/re-composing of word-aligned temporary registers into guest registers and back into temporary registers introduces lots of additional variables in the SMT solver. These *GET* and *PUT* statements are introduced into our IR in order to make the IR block generated for a single CPU instruction reentrant with respect to the guest state. Thereby we avoid the need to repeat the translation from object code to IR whenever an instruction is used in a different execution context, at the expense of having to deal with larger constraint systems.

An efficient way around this issue is to optimise unnecessary *GET* and *PUT* operations away, based on a *reaching definition* analysis for a given register and path. Practical results show that this simple optimisation greatly reduces the memory consumption of Yices for large constraint systems. We can apply the

same optimisations to memory accesses in cases where the address arguments to *LD* and *ST* evaluate to constant values. From our experience, dealing with unnecessary *GET*, *PUT*, *LD* and *ST* statements, by performing the above optimisations on IR level for an entire execution path, results in more efficient constraint systems and shorter runtimes of SOCA and Yices than by allowing Valgrind to perform similar optimisations at basic-block level.

***Determining a valid initial memory state.*** Another challenge when implementing symbolic execution as an SMT problem is given by the enormous search space that may result from leaving the program’s initial memory state undefined. OS components, including functions taken from device drivers, make regularly use of an external data environment consisting of heap objects allocated and initialised by other modules of the OS. Hence, this data environment cannot be inferred from the information available in the program binary. In practice, data environments can often be embedded into our analysis without much effort, by adding a few lines of C code as a preamble, as is shown in [32].

## 4 Experimental Results

To evaluate our SOCA technique regarding its ability to identify pointer safety issues and to judge its performance when analysing OS components, we have implemented SOCA in a prototypic tool, the *SOCA Verifier*. The tool comprises 15,000 lines of C code and took about one person-year to build; details of its architecture can be found in [30]. This section reports on the extensive experiments we conducted in applying the SOCA Verifier to a benchmark suite for software model checkers and to a large set of Linux device drivers. All experiments were carried out on a 16-core PC with 2.3 GHz clock speed and 256 GB of RAM, running 16 instances of the SOCA Verifier in parallel. However, an off-the-shelf PC with 4 GB of RAM is sufficient for everyday use, when one must not verify thousands of programs concurrently to meet a deadline.

### 4.1 Experiments I: The Verisec Benchmark

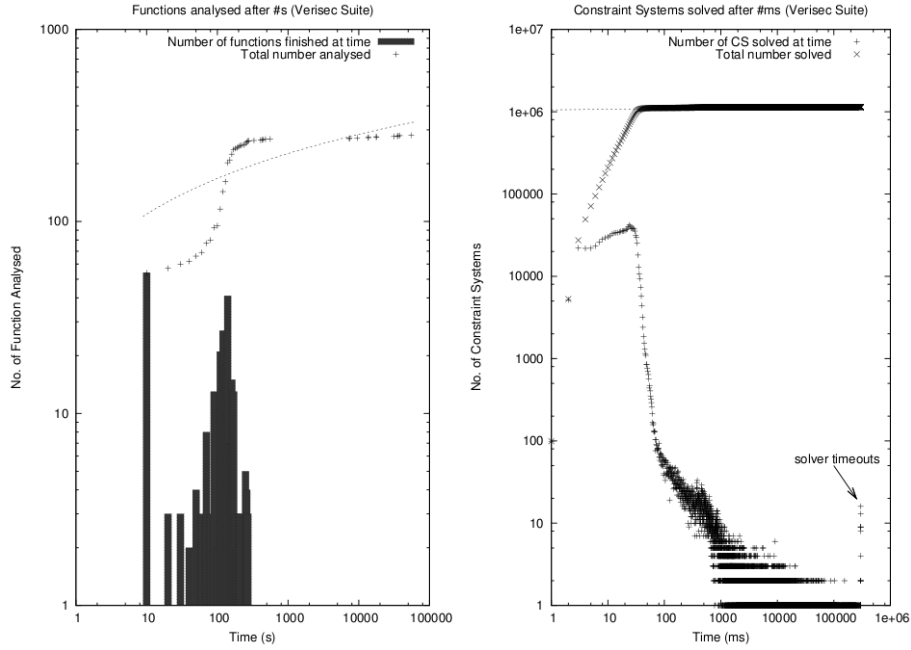
To enable a qualitative comparison of the SOCA Verifier to other tools, we applied it to the Verisec benchmark [27]. Verisec consists of 298 test programs (149 faulty programs – positive test programs – and 149 corresponding fixed programs – negative test programs) for buffer overflow vulnerabilities, taken from various open source programs. These test cases are given in terms of C source code which we compiled into object code using *gcc*, and provide a configurable buffer size which we set to 4. The bounds for the SOCA Verifier were set to a maximum of 100 paths to be analysed, where a single instruction may appear at most 500 times per path. Yices was configured to a timeout of 300 seconds per invocation. Of these bounds, only the timeout for Yices was ever reached.

In previous work [25,27], Verisec was used to evaluate the C-code model checkers *SatAbs* [12] and *LoopFrog* [25]. To enable a transparent comparison, we

**Table 1.** Comparison of SatAbs, LoopFrog and SOCA

	$R(d)$	$R(f)$	$R(\neg f d)$
SatAbs (from [27])	0.36	0.08	n/a
LoopFrog (from [25])	1.0	0.26	0.74
SOCA	0.66	0.23	0.81

adopt the metrics proposed in [45]: in Table 1 we report the *detection rate*  $R(d)$ , the *false-positive rate*  $R(f)$ , and the *discrimination rate*  $R(\neg f|d)$ . The latter is defined as the ratio of positive test cases for which an error is correctly reported, plus the negative test case for which the error is correctly not reported, to all test cases. Hence, tools are penalised for not finding bugs and for not reporting a sound program as safe.



**Figure 5.** Performance results for the Verisec benchmark. (a) Numbers of test cases verified by time (left). (b) Numbers of constraint systems solved by time (right).

As Table 1 testifies, the SOCA Verifier reliably detects the majority of buffer overflow errors in the benchmark, and has a lower false-positive rate and a better discrimination rate than the other tools. Remarkable is also that the SOCA Verifier failed for only 4 cases of the Verisec suite: once due to memory exhaustion and three times due to missing support for certain IR instructions in our tool. Only, our detection rate is lower than the one reported for LoopFrog. An explanation for this is the nature of Verisec’s test cases where static arrays are declared globally. This program setup renders Verisec easily comprehensible for

source-code verification tools since the bounds of data objects are clearly identifiable in source code. In object code, however, the boundaries of data objects are not visible anymore. This makes the SOCA Verifier less effective when analysing programs with small, statically declared buffers.

Hence, despite having used a benchmark providing examples which are in favour of source code analysis, our results show that object code analysis, as implemented in the SOCA Verifier, can compete with state-of-the-art source-code model checkers. However, as our tool analyses object code, it can be employed in a much wider application domain. Unfortunately, benchmarks that include dynamic allocation and provide examples of pointer safety errors other than buffer overflows are, to the best of our knowledge, not available.

**Table 2.** Performance statistics for the Verisec suite

	average	standard deviation	min	max	total
<b>per test case</b>					
total runtime	18m30s	1h33m	162ms	15h21m	91h54m
slicing time	28s150ms	41s808ms	28ms	5m15s	2h19m
Yices time	17m59s	1h33m	110ms	15h20m	89h19m
no. of CS	4025.11	173.76	11	8609	11994834
pointer operations	8.73	37.74	4	242	2603
<b>per Yices invocation</b>					
runtime	267ms	4s986ms	1ms	5m	88h59m
CS size	891.64	7707.95	0	368087	
memory usage	6.82MB	46.54MB	3.81MB	2504.36MB	

Fig. 5(a) gives details on run times; it shows the CPU times consumed for analysing each test cases in the Verisec benchmark. The vast majority of test cases is analysed by the SOCA Verifier within less than three minutes per case. As presented in Table 2, the average computation time consumed per test case is 18.5 minutes. In total, about 92 CPU hours were used. The memory consumption of both, the SOCA Verifier and Yices together, amounts to an average of only 140 MBytes and a maximum of about 3 GBytes, which is a memory capacity that is typically available in today’s PCs. Notably, Ku reported in [27] that the SatAbs tool crashed in 73 cases and timed out in another 87 cases with a timeout of 30 minutes. The runtime of the SOCA Verifier exceeds this time in only 7 cases.

In Fig. 5(b) we show the behaviour of Yices for solving the constraint systems generated by the SOCA Verifier. For the Verisec suite, a total of 11,994,834 constraint systems were solved in 89 hours. 2,250,878 (19%) of these constraint systems express verification properties, while the others were required for computing control flow, e.g., for deciding branching conditions and resolving computed jumps. With the timeout for Yices set to 5 minutes, the solver timed out on 34 constraint systems, and 96% of the constraint systems were solved in less than one second. Thus, the SOCA Verifier’s performance is on par with state-of-the-art software model checkers. Especially, it is efficient enough to be used

as an automated debugging tool by software developers, both regarding time efficiency and space efficiency.

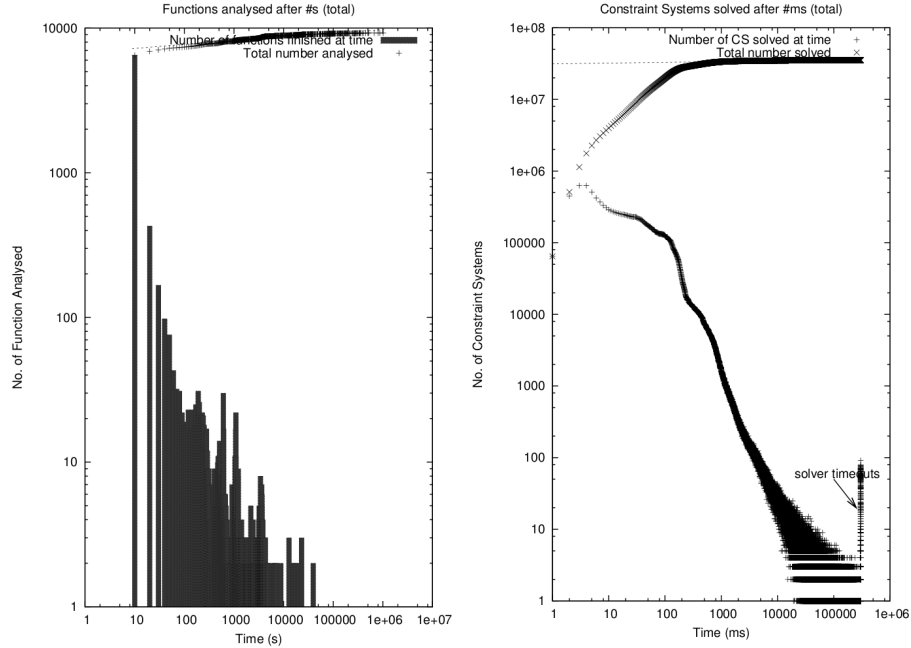
## 4.2 Experiments II: Linux Device Drivers

To evaluate the scalability of the SOCA Verifier, a large set of 9296 functions originating from 250 Linux device drivers of version 2.6.26 of the Linux kernel compiled for IA32 was analysed by us. Our experiments employ the Linux utility `nm` to obtain a list of function symbols present in a device driver. By statically linking the driver to the Linux kernel we resolved undefined symbols in the driver, i.e., functions provided by the OS kernel that are called by the driver’s functions. The SOCA technique was then applied on the resulting binary file to analyse each of the driver’s functions separately. The bounds for the SOCA Verifier were set to a maximum of 1000 paths to be analysed, where a single instruction may appear at most 1000 times per path, thereby effectively bounding the number of loop iterations or recursions to that depth. Moreover, Yices was configured to a timeout of 300 seconds per invocation.

Table 3. Performance statistics for the Linux device drivers					
	average	standard	min	max	total
	deviation				
per test case					
total runtime	58m28s	7h56m	21ms	280h48m	9058h32m
slicing time	8m35s	2h13m	0	95h39m	1329h46m
Yices time	48m36s	7h28m	0	280h30m	7531h51m
no. of CS	3591.14	9253.73	0	53449	33383239
pointer operations	99.53	312.64	0	4436	925277
no. of paths	67.50	221.17	1	1000	627524
max path lengths	727.22	1819.28	1	22577	
per Yices invocation					
runtime	845ms	8s765ms	1ms	5m2s	8295h56m
CS size	4860.20	20256.77	0	7583410	
Memory usage	5.75MB	14.76MB	3.81MB	3690.00MB	

Our obtained results are summarised in Table 3 and show that 94.4% of the functions in our sample could be analysed by the SOCA Verifier. In 67.5% of the functions the exhaustion of execution bounds led to an early termination of the analysis. However, the analysis reached a considerable depth even in those cases, analysing paths of lengths of up to 22,577 CPU instructions. Interestingly, 27.8% of the functions could be analysed exhaustively, where none of the bounds regarding the number of paths, the path lengths, or the SMT solver’s timeout were reached. As depicted in Fig. 6(a), the SOCA Verifier returns a result in less than 10 mins in the majority of cases, while the generated constraint systems were usually solved in less than 500 ms. The timeout for Yices was hardly ever reached (cf. Fig. 6(b)).

As an aside, it should be mentioned that in 0.98% (91 functions) of the sample Linux driver functions, the SOCA Verifier may have produced unsound results



**Figure 6.** Performance results for the Linux device drivers. (a) Numbers of test cases verified by time (left). (b) Numbers of constraint systems solved by time (right).

due to non-linear arithmetic within the generated constraint systems, which is not decidable by Yices. In addition, our verifier failed in 5.6% of the cases (522 functions) due to either memory exhaustion, missing support for particular assembly instructions in our tool or Valgrind, or crashes of Yices.

Our evaluation shows that the SOCA Verifier scales to real-world OS software while delivering very good performance. Being automatic and not restricted to analysing programs available in source code only, the SOCA Verifier is an efficient tool that is capable of aiding a practitioner in debugging pointer-complex software such as OS components. The application of the SOCA Verifier is, however, not restricted to verifying memory safety. In [32] we presented a case study on retrospective verification of the Linux Virtual File System (VFS) using the SOCA Verifier for checking violations of API usage rules such as deadlocks caused by misuse of the Linux kernel’s `spinlock` API.

## 5 Related Work

A survey on automated techniques for formal software verification can be found in [13]. By having the potential of being exhaustive and fully automatic, *model checking*, in combination with *abstraction* and *refinement*, is a successful technique used in software verification [11].

**Model checking bytecode and assembly languages.** In recent years, several approaches to model checking *compiled programs* by analysing bytecode

and assembly code have been presented. In [40], *Java PathFinder (JPF)* for model checking Java bytecode is introduced. *JPF* generates the state space of a program by monitoring a virtual machine. Model checking is then conducted on the states explored by the virtual machine, employing collapsing techniques and symmetry reduction for efficiently storing states and reducing the size of the state space. These techniques are effective because of the high complexity of *JPF* states and the specific characteristics of the Java memory model. In contrast, the SOCA technique to verifying object code involves relatively simple states and, in difference to Java, the order of data within memory is important in IA32 object code. Similar to *JPF*, *StEAM* [29] model checks bytecode compiled for the Internet C Virtual Machine.

*BTOR* [6] and *[mc]square* [34,38] are tools for model checking assembly code for micro-controllers. They accept assembly code as their input, which may either be obtained during compilation of a program or, as suggested in [38], by disassembling a binary program. As shown in [18], the problem of disassembling a binary program is undecidable in general. The SOCA technique focuses on the verification of binary programs without the requirement of disassembling a program at once.

All the above tools are explicit model checkers that require a program’s entire control flow to be known in advance of the analysis. As we have explained above, this is not feasible in the presence of computed jumps. The SOCA technique has been especially designed to deal with OS components that make extensive use of jump computations.

***Combining model checking with symbolic execution.*** *Symbolic execution* was introduced by King [23] as a means of improving program testing by covering a large class of normal executions with one execution, in which symbols representing arbitrary values are used as input to the program. This is exactly what our SOCA technique does, albeit not for testing but for systematic, powerful memory safety analysis. A recent approach using symbolic execution to derive inputs that make a given program crash has been proposed in *EXE* [7]. In contrast to our work, *EXE* relies on manual annotations, is not focused on memory safety, and works at source code level.

Several frameworks for integrating symbolic execution with model checking have recently been presented, including *Symbolic JPF* [35] and *DART* [16]. *Symbolic JPF* is a successor of the previously mentioned *JPF*. *DART* implements directed and automated random testing to generate test drivers and harness code to simulate a program’s environment. The tool accepts C programs and automatically extracts function interfaces from source code. Such an interface is used to seed the analysis with a well-formed random input, which is then mutated by collecting and negating path constraints while symbolically executing the program. Unlike the SOCA Verifier, *DART* handles constraints on integer types only and does not support pointers and data structures.

A language agnostic tool in the spirit of *DART* is *SAGE* [15], which is used internally at Microsoft. *SAGE* works at IA32 instruction level, tracks integer constraints as bit-vectors, and employs machine-code instrumentation in a sim-



ilar fashion as we do in [32]. *SAGE* is seeded with a well-formed program input and explores the program space with respect to that input. Branches in the control flow are explored by negating path constraints collected during the initial execution. This differs from our approach since SOCA does not require seeding but explores the program space automatically from a given starting point. The SOCA technique effectively computes program inputs for all paths explored during symbolic execution.

A bounded model checker for C source code based on symbolic execution and SAT solving is SATURN [43]. This tool is specialised on checking locking properties and null-pointer de-references and is thus not as general as SOCA. The authors show that their tool scales for analysing the entire Linux kernel. Unlike the SOCA Verifier, the approach in [43] computes function summaries instead of adding the respective code to the control flow, unwinds loops a fixed number of times and does not handle recursion.

**Concolic testing.** An area of research closely related to ours is that of *concolic testing* [22,39]. This technique relies on performing concrete execution on random inputs while collecting path constraints along executed paths. The constraints are then used to compute new inputs driving the program along alternative paths. In difference to this approach, SOCA uses symbolic execution to explore all paths or, at least orders of magnitude more paths, and concretises only for resolving computed jumps.

**Program slicing.** An important SOCA ingredient other than symbolic execution is *path-sensitive* slicing. *Program slicing* was introduced by Weiser [41] as a technique for automatically selecting only those parts of a program that may affect the values of interest computed at some point of interest. Different to conventional slicing, our slices are computed over a single path instead of an entire program, similar to what has been introduced as *dynamic slicing* in [24] and *path slicing* in [20]. In contrast to those approaches, we use conventional slicing criteria and leave a program’s input initially unspecified. In addition, while collecting program dependencies is relatively easy at source code level, it becomes difficult at object code level when dependencies to the heap and stack are involved. The technique employed by SOCA for dealing with the program’s heap and stack is an variation of the *recency abstraction* described in [1].

**Alternative approaches.** Alternative, recent approaches to proving memory safety are *shape analysis* [42] and *separation logic* [37]. All recent work in this area [21,8] is based on analysing the source code of a program, and calls to library functions and programming constructs such as function pointers are simply abstracted using non-deterministic assignments.

Techniques applying theorem proving to verify object code and assembly code are presented in [5,44]. In [5] the Nqthm prover is employed for reasoning about the functional correctness of implementations of well-known algorithms. [44] proposes a logic-based type system for concurrent assembly code and uses

the Coq proof assistant to verify programs. In contrast to our work, both techniques do not support “higher-order code pointers”, including return pointers in procedure calls.

**Testing pointer safety.** Finally, validation and testing tools such as *Purify* [36] and *Valgrind* [33] must be mentioned since they have been applied successfully to identifying memory safety problems in application software, rather than OS software. They execute an instrumented version of a given program in a protected environment in which various invalid pointer operations can be detected. However, they are meant for manual testing and do not provide means for automatically, and potentially exhaustively, exploring all program behaviour.

## 6 Conclusions and Future Work

This paper presented the novel SOCA technique for automatically checking memory safety of pointer-complex software. Analysing object code allows us to handle software, e.g., OS software, which is written in a mix of C and inlined assembly. Together with SOCA’s symbolic execution, this simplifies pointer analysis when being confronted with function pointers, computed jumps and pointer aliasing. SOCA achieves scalability by adopting path-sensitive slicing and the efficient SMT solver Yices. While the SOCA ingredients are well-known, the way in which we integrated these for automated object code analysis is novel. Much effort went into engineering our SOCA Verifier, and extensive benchmarking showed that it performs on par with state-of-the-art software model checkers and scales well when applied to Linux device driver functions. Our verifier explores semantic niches of software, especially OS software, which current model checkers and testing tools do not reach.

Future work shall be pursued along several orthogonal lines. Firstly, since device driver functions may be invoked concurrently, we plan to extend SOCA to handle concurrency. To the best of our knowledge, the verification of concurrent programs with full pointer arithmetic and computed jumps is currently not supported by any automated verification tool. Secondly, we intend to evaluate different search strategies for exploring the paths of a program, employing heuristics based on, e.g., coverage criteria. Thirdly, as some inputs of device drivers involve pointered data structures, we wish to explore whether *shape analysis* can inform SOCA in a way that reduces the number of false positives raised. Fourthly, the SOCA Verifier shall be interfaced to the *gnu debugger* so that error traces can be played back in a user-friendly form, at source code level.

**Acknowledgements** We would thank Jim Woodcock and Daniel Kroening for their insightful comments made at the first author’s PhD examination.

## References

1. Balakrishnan, G. and Reps, T. Recency-abstraction for heap-allocated storage. In *SAS ’06*, vol. 4134 of *LNCS*, pp. 221–239. Springer, 2006.

2. Balakrishnan, G., Reps, T., Melski, D., and Teitelbaum, T. WYSINWYX: What You See Is Not What You eXecute. In *VSTTE '08*, vol. 4171 of *LNCS*, pp. 202–213. Springer, 2008.
3. Ball, T., Bounimova, E., Cook, B., Levin, V., Lichtenberg, J., McGarvey, C., Ondrusek, B., Rajamani, S. K., and Ustuner, A. Thorough static analysis of device drivers. *SIGOPS Oper. Syst. Rev.*, 40(4):73–85, 2006.
4. Ball, T. and Rajamani, S. K. Automatically validating temporal safety properties of interfaces. In *SPIN '01*, vol. 2057 of *LNCS*, pp. 103–122. Springer, 2001.
5. Boyer, R. S. and Yu, Y. Automated proofs of object code for a widely used microprocessor. *J. ACM*, 43(1):166–192, 1996.
6. Brummayer, R., Biere, A., and Lonsing, F. BTOR: Bit-precise modelling of word-level problems for model checking. In *SMT '08/BPR '08*, pp. 33–38. ACM, 2008.
7. Cadar, C., Ganesh, V., Pawlowski, P. M., Dill, D. L., and Engler, D. R. EXE: Automatically generating inputs of death. In *CCS '06*, pp. 322–335. ACM, 2006.
8. Calcagno, C., Distefano, D., O'Hearn, P., and Yang, H. Compositional shape analysis by means of bi-abduction. *SIGPLAN Not.*, 44(1):289–300, 2009.
9. Chou, A., Yang, J., Chelf, B., Hallem, S., and Engler, D. R. An empirical study of operating system errors. In *SOSP '01*, pp. 73–88. ACM, 2001.
10. Clarke, E., Kroening, D., and Lerda, F. A tool for checking ANSI-C programs. In *TACAS '04*, vol. 2988 of *LNCS*, pp. 168–176. Springer, 2004.
11. Clarke, E. M., Grumberg, O., and Peled, D. A. *Model checking*. MIT Press, 2000.
12. Clarke, E. M., Kroening, D., Sharygina, N., and Yorav, K. SATABS: SAT-based predicate abstraction for ANSI-C. In *TACAS '05*, vol. 3440 of *LNCS*, pp. 570–574. Springer, 2005.
13. D'Silva, V., Kroening, D., and Weissenbacher, G. A survey of automated techniques for formal software verification. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 27(7):1165–1178, 2008.
14. Dutertre, B. and de Moura, L. The Yices SMT solver. Technical Report 01/2006, SRI, 2006. Available at <http://yices.cs1.sri.com/tool-paper.pdf>.
15. Godefroid, P., de Halleux, P., Nori, A. V., Rajamani, S. K., Schulte, W., Tillmann, N., and Levin, M. Y. Automating software testing using program analysis. *IEEE Software*, 25(5):30–37, 2008.
16. Godefroid, P., Klarlund, N., and Sen, K. DART: Directed automated random testing. In *PLDI '05*, pp. 213–223. ACM, 2005.
17. Henzinger, T. A., Jhala, R., Majumdar, R., Nacula, G. C., Sutre, G., and Weimer, W. Temporal-safety proofs for systems code. In *CAV '02*, vol. 2402 of *LNCS*, pp. 526–538. Springer, 2002.
18. Horspool, R. N. and Marovac, N. An approach to the problem of detranslation of computer programs. *Computer J.*, 23(3):223–229, 1980.
19. Horwitz, S., Reps, T., and Binkley, D. Interprocedural slicing using dependence graphs. *ACM TOPLAS*, 12(1):26–60, 1990.
20. Jhala, R. and Majumdar, R. Path slicing. *SIGPLAN Not.*, 40(6):38–47, 2005.
21. Josh Berdine, C. C. and O'Hearn, P. W. Symbolic execution with separation logic. In *APLAS '05*, vol. 3780 of *LNCS*, pp. 52–68. Springer, 2005.
22. Kim, M. and Kim, Y. Concolic testing of the multi-sector read operation for flash memory file system. In *SBMF '09*, vol. 5902 of *LNCS*, pp. 251–265. Springer, 2009.
23. King, J. C. Symbolic execution and program testing. *Commun. ACM*, 19(7):385–394, 1976.
24. Korel, B. and Laski, J. Dynamic slicing of computer programs. *J. Syst. Softw.*, 13(3):187–195, 1990.

25. Kroening, D., Sharygina, N., Tonetta, S., Tsitovich, A., and Wintersteiger, C. M. Loop summarization using abstract transformers. In *ATVA '08*, vol. 5311 of *LNCS*, pp. 111–125. Springer, 2008.
26. Kroening, D. and Strichman, O. *Decision Procedures*. Springer, 2008.
27. Ku, K. Software model-checking: Benchmarking and techniques for buffer overflow analysis. Master's thesis, University of Toronto, 2008.
28. Leung, A. and George, L. Static single assignment form for machine code. In *PLDI '99*, pp. 204–214. ACM, 1999.
29. Leven, P., Mehler, T., and Edelkamp, S. Directed error detection in C++ with the assembly-level model checker StEAM. In *Model Checking Software*, vol. 2989 of *LNCS*, pp. 39–56. Springer, 2004.
30. Mühlberg, J. T. *Model Checking Pointer Safety in Compiled Programs*. PhD thesis, Department of Computer Science, University of York, 2010. Examined, available at [http://swt.uni-bamberg.de/soca/soca\\_thesis.pdf](http://swt.uni-bamberg.de/soca/soca_thesis.pdf).
31. Mühlberg, J. T. and Lüttgen, G. BLASTing Linux code. In *FMICS '06*, vol. 4346 of *LNCS*, pp. 211–226. Springer, 2006.
32. Mühlberg, J. T. and Lüttgen, G. Verifying compiled file system code. In *SBMF '09*, vol. 5902 of *LNCS*, pp. 306–320. Springer, 2009.
33. Nethercote, N. and Seward, J. Valgrind: A framework for heavyweight dynamic binary instrumentation. *SIGPLAN Not.*, 42(6):89–100, 2007.
34. Noll, T. and Schlich, B. Delayed nondeterminism in model checking embedded systems assembly code. In *Hardware and Software: Verification and Testing*, vol. 4899 of *LNCS*, pp. 185–201. Springer, 2008.
35. Păsăreanu, C. S., Mehltz, P. C., Bushnell, D. H., Gundy-Burlet, K., Lowry, M., Person, S., and Pape, M. Combining unit-level symbolic execution and system-level concrete execution for testing NASA software. In *ISSTA '08*, pp. 15–26. ACM, 2008.
36. Rational Purify. IBM Corp., <http://www.ibm.com/software/awdtools/purify/>.
37. Reynolds, J. C. Separation logic: A logic for shared mutable data structures. In *LICS '02*, pp. 55–74. IEEE, 2002.
38. Schlich, B. and Kowalewski, S. [mc]square: A model checker for microcontroller code. In *ISOLA '06*, pp. 466–473. IEEE, 2006.
39. Sen, K., Marinov, D., and Agha, G. CUTE: A concolic unit testing engine for C. In *ESEC/FSE-13*, pp. 263–272. ACM, 2005.
40. Visser, W., Havelund, K., Brat, G., Park, S. J., and Lerda, F. Model checking programs. *FMSD*, 10(2):203–232, 2003.
41. Weiser, M. Program slicing. In *ICSE '81*, pp. 439–449. IEEE, 1981.
42. Wilhelm, R., Sagiv, M., and Reps, T. Shape analysis. In *CC '00*, vol. 1781 of *LNCS*, pp. 1–16. Springer, 2000.
43. Xie, Y. and Aiken, A. SATURN: A scalable framework for error detection using boolean satisfiability. *ACM TOPLAS*, 29(3):16, 2007.
44. Yu, D. and Shao, Z. Verification of safety properties for concurrent assembly code. In *ICFP '04*, pp. 175–188. ACM, 2004.
45. Zitser, M., Lippmann, R., and Leek, T. Testing static analysis tools using exploitable buffer overflows from open source code. *SIGSOFT Softw. Eng. Notes*, 29(6):97–106, 2004.

<b>Bamberger Beiträge zur Wirtschaftsinformatik</b>
---

- |               |   |
|---------------|---|
| Nr. 1 (1989)  | Augsburger W., Bartmann D., Sinz E.J.: Das Bamberger Modell: Der Diplom-Studiengang Wirtschaftsinformatik an der Universität Bamberg (Nachdruck Dez. 1990)  |
| Nr. 2 (1990)  | Esswein W.: Definition, Implementierung und Einsatz einer kompatiblen Datenbankschnittstelle für PROLOG   |
| Nr. 3 (1990)  | Augsburger W., Rieder H., Schwab J.: Endbenutzerorientierte Informationsgewinnung aus numerischen Daten am Beispiel von Unternehmenskennzahlen  |
| Nr. 4 (1990)  | Ferstl O.K., Sinz E.J.: Objektmodellierung betrieblicher Informationsmodelle im Semantischen Objektmodell (SOM) (Nachdruck Nov. 1990)   |
| Nr. 5 (1990)  | Ferstl O.K., Sinz E.J.: Ein Vorgehensmodell zur Objektmodellierung betrieblicher Informationssysteme im Semantischen Objektmodell (SOM)   |
| Nr. 6 (1991)  | Augsburger W., Rieder H., Schwab J.: Systemtheoretische Repräsentation von Strukturen und Bewertungsfunktionen über zeitabhängigen betrieblichen numerischen Daten  |
| Nr. 7 (1991)  | Augsburger W., Rieder H., Schwab J.: Wissensbasiertes, inhaltsorientiertes Retrieval statistischer Daten mit EISREVU / Ein Verarbeitungsmodell für eine modulare Bewertung von Kennzahlenwerten für den Endanwender |
| Nr. 8 (1991)  | Schwab J.: Ein computergestütztes Modellierungssystem zur Kennzahlenbewertung   |
| Nr. 9 (1992)  | Gross H.-P.: Eine semantiktreue Transformation vom Entity-Relationship-Modell in das Strukturierte Entity-Relationship-Modell   |
| Nr. 10 (1992) | Sinz E.J.: Datenmodellierung im Strukturierten Entity-Relationship-Modell (SERM)  |
| Nr. 11 (1992) | Ferstl O.K., Sinz E. J.: Glossar zum Begriffssystem des Semantischen Objektmodells  |
| Nr. 12 (1992) | Sinz E. J., Popp K.M.: Zur Ableitung der Grobstruktur des konzeptuellen Schemas aus dem Modell der betrieblichen Diskurswelt  |
| Nr. 13 (1992) | Esswein W., Locarek H.: Objektorientierte Programmierung mit dem Objekt-Rollenmodell  |
| Nr. 14 (1992) | Esswein W.: Das Rollenmodell der Organisation: Die Berücksichtigung aufbauorganisatorische Regelungen in Unternehmensmodellen   |
| Nr. 15 (1992) | Schwab H. J.: EISREVU-Modellierungssystem. Benutzerhandbuch   |
| Nr. 16 (1992) | Schwab K.: Die Implementierung eines relationalen DBMS nach dem Client/Server-Prinzip   |
| Nr. 17 (1993) | Schwab K.: Konzeption, Entwicklung und Implementierung eines computergestützten Bürovorgangssystems zur Modellierung von Vorgangsklassen und Abwicklung und Überwachung von Vorgängen. Dissertation                 |

- Nr. 18 (1993) Ferstl O.K., Sinz E.J.: Der Modellierungsansatz des Semantischen Objektmodells
- Nr. 19 (1994) Ferstl O.K., Sinz E.J., Amberg M., Hagemann U., Malischewski C.: Tool-Based Business Process Modeling Using the SOM Approach
- Nr. 20 (1994) Ferstl O.K., Sinz E.J.: From Business Process Modeling to the Specification of Distributed Business Application Systems - An Object-Oriented Approach -. 1<sup>st</sup> edition, June 1994
- Ferstl O.K., Sinz E.J. : Multi-Layered Development of Business Process Models and Distributed Business Application Systems - An Object-Oriented Approach -. 2<sup>nd</sup> edition, November 1994
- Nr. 21 (1994) Ferstl O.K., Sinz E.J.: Der Ansatz des Semantischen Objektmodells zur Modellierung von Geschäftsprozessen
- Nr. 22 (1994) Augsburg W., Schwab K.: Using Formalism and Semi-Formal Constructs for Modeling Information Systems
- Nr. 23 (1994) Ferstl O.K., Hagemann U.: Simulation hierarischer objekt- und transaktionsorientierter Modelle
- Nr. 24 (1994) Sinz E.J.: Das Informationssystem der Universität als Instrument zur zielgerichteten Lenkung von Universitätsprozessen
- Nr. 25 (1994) Wittke M., Mekinich, G.: Kooperierende Informationsräume. Ein Ansatz für verteilte Führungsinformationssysteme
- Nr. 26 (1995) Ferstl O.K., Sinz E.J.: Re-Engineering von Geschäftsprozessen auf der Grundlage des SOM-Ansatzes
- Nr. 27 (1995) Ferstl, O.K., Mannmeusel, Th.: Dezentrale Produktionslenkung. Erscheint in CIM-Management 3/1995
- Nr. 28 (1995) Ludwig, H., Schwab, K.: Integrating cooperation systems: an event-based approach
- Nr. 30 (1995) Augsburg W., Ludwig H., Schwab K.: Koordinationsmethoden und -werkzeuge bei der computergestützten kooperativen Arbeit
- Nr. 31 (1995) Ferstl O.K., Mannmeusel T.: Gestaltung industrieller Geschäftsprozesse
- Nr. 32 (1995) Gunzenhäuser R., Duske A., Ferstl O.K., Ludwig H., Mekinich G., Rieder H., Schwab H.-J., Schwab K., Sinz E.J., Wittke M: Festschrift zum 60. Geburtstag von Walter Augsburg
- Nr. 33 (1995) Sinz, E.J.: Kann das Geschäftsprozeßmodell der Unternehmung das unternehmensweite Datenschema ablösen?
- Nr. 34 (1995) Sinz E.J.: Ansätze zur fachlichen Modellierung betrieblicher Informationssysteme - Entwicklung, aktueller Stand und Trends -
- Nr. 35 (1995) Sinz E.J.: Serviceorientierung der Hochschulverwaltung und ihre Unterstützung durch workflow-orientierte Anwendungssysteme
- Nr. 36 (1996) Ferstl O.K., Sinz, E.J., Amberg M.: Stichwörter zum Fachgebiet Wirtschaftsinformatik. Erscheint in: Broy M., Spaniol O. (Hrsg.): Lexikon Informatik und Kommunikationstechnik, 2. Auflage, VDI-Verlag, Düsseldorf 1996

- Nr. 37 (1996) Ferstl O.K., Sinz E.J.: Flexible Organizations Through Object-oriented and Transaction-oriented Information Systems, July 1996
- Nr. 38 (1996) Ferstl O.K., Schäfer R.: Eine Lernumgebung für die betriebliche Aus- und Weiterbildung on demand, Juli 1996
- Nr. 39 (1996) Hazebrouck J.-P.: Einsatzpotentiale von Fuzzy-Logic im Strategischen Management dargestellt an Fuzzy-System-Konzepten für Portfolio-Ansätze
- Nr. 40 (1997) Sinz E.J.: Architektur betrieblicher Informationssysteme. In: Rechenberg P., Pomberger G. (Hrsg.): Handbuch der Informatik, Hanser-Verlag, München 1997
- Nr. 41 (1997) Sinz E.J.: Analyse und Gestaltung universitärer Geschäftsprozesse und Anwendungssysteme. Angenommen für: Informatik '97. Informatik als Innovationsmotor. 27. Jahrestagung der Gesellschaft für Informatik, Aachen 24.-26.9.1997
- Nr. 42 (1997) Ferstl O.K., Sinz E.J., Hammel C., Schlitt M., Wolf S.: Application Objects – fachliche Bausteine für die Entwicklung komponentenbasierter Anwendungssysteme. Angenommen für: HMD – Theorie und Praxis der Wirtschaftsinformatik. Schwerpunkt Heft ComponentWare, 1997
- Nr. 43 (1997): Ferstl O.K., Sinz E.J.: Modeling of Business Systems Using the Semantic Object Model (SOM) – A Methodological Framework - . Accepted for: P. Bernus, K. Mertins, and G. Schmidt (ed.): Handbook on Architectures of Information Systems. International Handbook on Information Systems, edited by Bernus P., Blazewicz J., Schmidt G., and Shaw M., Volume I, Springer 1997
- Ferstl O.K., Sinz E.J.: Modeling of Business Systems Using (SOM), 2<sup>nd</sup> Edition. Appears in: P. Bernus, K. Mertins, and G. Schmidt (ed.): Handbook on Architectures of Information Systems. International Handbook on Information Systems, edited by Bernus P., Blazewicz J., Schmidt G., and Shaw M., Volume I, Springer 1998
- Nr. 44 (1997) Ferstl O.K., Schmitz K.: Zur Nutzung von Hypertextkonzepten in Lernumgebungen. In: Conradi H., Kreutz R., Spitzer K. (Hrsg.): CBT in der Medizin – Methoden, Techniken, Anwendungen -. Proceedings zum Workshop in Aachen 6. – 7. Juni 1997. 1. Auflage Aachen: Verlag der Augustinus Buchhandlung
- Nr. 45 (1998) Ferstl O.K.: Datenkommunikation. In: Schulte Ch. (Hrsg.): Lexikon der Logistik, Oldenbourg-Verlag, München 1998
- Nr. 46 (1998) Sinz E.J.: Prozeßgestaltung und Prozeßunterstützung im Prüfungswesen. Erschienen in: Proceedings Workshop „Informationssysteme für das Hochschulmanagement“. Aachen, September 1997
- Nr. 47 (1998) Sinz, E.J., Wismans B.: Das „Elektronische Prüfungsamt“. Erscheint in: Wirtschaftswissenschaftliches Studium WiSt, 1998
- Nr. 48 (1998) Haase, O., Henrich, A.: A Hybrid Representation of Vague Collections for Distributed Object Management Systems. Erscheint in: IEEE Transactions on Knowledge and Data Engineering
- Nr. 49 (1998) Henrich, A.: Applying Document Retrieval Techniques in Software Engineering Environments. In: Proc. International Conference on Database and Expert Systems

Applications. (DEXA 98), Vienna, Austria, Aug. 98, pp. 240-249, Springer, Lecture Notes in Computer Sciences, No. 1460

- Nr. 50 (1999) Henrich, A., Jamin, S.: On the Optimization of Queries containing Regular Path Expressions. Erscheint in: Proceedings of the Fourth Workshop on Next Generation Information Technologies and Systems (NGITS'99), Zikhron-Yaakov, Israel, July, 1999 (Springer, Lecture Notes)
- Nr. 51 (1999) Haase O., Henrich, A.: A Closed Approach to Vague Collections in Partly Inaccessible Distributed Databases. Erscheint in: Proceedings of the Third East-European Conference on Advances in Databases and Information Systems – ADBIS'99, Maribor, Slovenia, September 1999 (Springer, Lecture Notes in Computer Science)
- Nr. 52 (1999) Sinz E.J., Böhnlein M., Ulbrich-vom Ende A.: Konzeption eines Data Warehouse-Systems für Hochschulen. Angenommen für: Workshop „Unternehmen Hochschule“ im Rahmen der 29. Jahrestagung der Gesellschaft für Informatik, Paderborn, 6. Oktober 1999
- Nr. 53 (1999) Sinz E.J.: Konstruktion von Informationssystemen. Der Beitrag wurde in geringfügig modifizierter Fassung angenommen für: Rechenberg P., Pomberger G. (Hrsg.): Informatik-Handbuch. 2., aktualisierte und erweiterte Auflage, Hanser, München 1999
- Nr. 54 (1999) Herda N., Janson A., Reif M., Schindler T., Augsburg W.: Entwicklung des Intranets SPICE: Erfahrungsbericht einer Praxiskooperation.
- Nr. 55 (2000) Böhnlein M., Ulbrich-vom Ende A.: Grundlagen des Data Warehousing. Modellierung und Architektur
- Nr. 56 (2000) Freitag B., Sinz E.J., Wismans B.: Die informationstechnische Infrastruktur der Virtuellen Hochschule Bayern (vhb). Angenommen für Workshop "Unternehmen Hochschule 2000" im Rahmen der Jahrestagung der Gesellschaft f. Informatik, Berlin 19. - 22. September 2000
- Nr. 57 (2000) Böhnlein M., Ulbrich-vom Ende A.: Developing Data Warehouse Structures from Business Process Models.
- Nr. 58 (2000) Knobloch B.: Der Data-Mining-Ansatz zur Analyse betriebswirtschaftlicher Daten.
- Nr. 59 (2001) Sinz E.J., Böhnlein M., Plaha M., Ulbrich-vom Ende A.: Architekturkonzept eines verteilten Data-Warehouse-Systems für das Hochschulwesen. Angenommen für: WI-IF 2001, Augsburg, 19.-21. September 2001
- Nr. 60 (2001) Sinz E.J., Wismans B.: Anforderungen an die IV-Infrastruktur von Hochschulen. Angenommen für: Workshop „Unternehmen Hochschule 2001“ im Rahmen der Jahrestagung der Gesellschaft für Informatik, Wien 25. – 28. September 2001

Änderung des Titels der Schriftenreihe *Bamberger Beiträge zur Wirtschaftsinformatik* in *Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik* ab Nr. 61

Note: The title of our technical report series has been changed from *Bamberger Beiträge zur Wirtschaftsinformatik* to *Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik* starting with TR No. 61



<h2 style="text-align: center;">Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik</h2>
--

- |               |  |
|---------------|--|
| Nr. 61 (2002) | Goré R., Mendler M., de Paiva V. (Hrsg.): Proceedings of the International Workshop on Intuitionistic Modal Logic and Applications (IMLA 2002), Copenhagen, July 2002.   |
| Nr. 62 (2002) | Sinz E.J., Plaha M., Ulbrich-vom Ende A.: Datenschutz und Datensicherheit in einem landesweiten Data-Warehouse-System für das Hochschulwesen. Erscheint in: Beiträge zur Hochschulforschung, Heft 4-2002, Bayerisches Staatsinstitut für Hochschulforschung und Hochschulplanung, München 2002   |
| Nr. 63 (2005) | Aguado, J., Mendler, M.: Constructive Semantics for Instantaneous Reactions  |
| Nr. 64 (2005) | Ferstl, O.K.: Lebenslanges Lernen und virtuelle Lehre: globale und lokale Verbesserungspotenziale. Erschienen in: Kerres, Michael; Keil-Slawik, Reinhard (Hrsg.); Hochschulen im digitalen Zeitalter: Innovationspotenziale und Strukturwandel, S. 247 – 263; Reihe education quality forum, herausgegeben durch das Centrum für eCompetence in Hochschulen NRW, Band 2, Münster/New York/München/Berlin: Waxmann 2005 |
| Nr. 65 (2006) | Schönberger, Andreas: Modelling and Validating Business Collaborations: A Case Study on RosettaNet   |
| Nr. 66 (2006) | Markus Dorsch, Martin Grote, Knut Hildebrandt, Maximilian Röglinger, Matthias Sehr, Christian Wilms, Karsten Loesing, and Guido Wirtz: Concealing Presence Information in Instant Messaging Systems, April 2006  |
| Nr. 67 (2006) | Marco Fischer, Andreas Grünert, Sebastian Hudert, Stefan König, Kira Lenskaya, Gregor Scheithauer, Sven Kaffille, and Guido Wirtz: Decentralized Reputation Management for Cooperating Software Agents in Open Multi-Agent Systems, April 2006   |
| Nr. 68 (2006) | Michael Mendler, Thomas R. Shiple, Gérard Berry: Constructive Circuits and the Exactness of Ternary Simulation   |
| Nr. 69 (2007) | Sebastian Hudert: A Proposal for a Web Services Agreement Negotiation Protocol Framework . February 2007   |
| Nr. 70 (2007) | Thomas Meins: Integration eines allgemeinen Service-Centers für PC-und Medientechnik an der Universität Bamberg – Analyse und Realisierungs-Szenarien. Februar 2007  |
| Nr. 71 (2007) | Andreas Grünert: Life-cycle assistance capabilities of cooperating Software Agents for Virtual Enterprises. März 2007  |
| Nr. 72 (2007) | Michael Mendler, Gerald Lüttgen: Is Observational Congruence on $\mu$ -Expressions Axiomatisable in Equational Horn Logic?   |
| Nr. 73 (2007) | Martin Schissler: out of print   |
| Nr. 74 (2007) | Sven Kaffille, Karsten Loesing: Open chord version 1.0.4 User's Manual. Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik Nr. 74, Bamberg University, October 2007. ISSN 0937-3349.  |

- Nr. 75 (2008) Karsten Loesing (Hrsg.): Extended Abstracts of the Second *Privacy Enhancing Technologies Convention* (PET-CON 2008.1). Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik Nr. 75, Bamberg University, April 2008. ISSN 0937-3349.
- Nr. 76 (2008) Gregor Scheithauer and Guido Wirtz: Applying Business Process Management Systems? A Case Study. Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik Nr. 76, Bamberg University, May 2008. ISSN 0937-3349.
- Nr. 77 (2008) Michael Mendler, Stephan Scheele: Towards Constructive Description Logics for Abstraction and Refinement. Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik Nr. 77, Bamberg University, September 2008. ISSN 0937-3349.
- Nr. 78 (2008) Gregor Scheithauer and Matthias Winkler: A Service Description Framework for Service Ecosystems. Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik Nr. 78, Bamberg University, October 2008. ISSN 0937-3349.
- Nr. 79 (2008) Christian Wilms: Improving the Tor Hidden Service Protocol Aiming at Better Performances. Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik Nr. 79, Bamberg University, November 2008. ISSN 0937-3349.
- Nr. 80 (2009) Thomas Benker, Stefan Fritzemeier, Matthias Geiger, Simon Harrer, Tristan Kessner, Johannes Schwalb, Andreas Schönberger, Guido Wirtz: QoS Enabled B2B Integration. Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik Nr. 80, Bamberg University, May 2009. ISSN 0937-3349.
- Nr. 81 (2009) Ute Schmid, Emanuel Kitzelmann, Rinus Plasmeijer (Eds.): Proceedings of the ACM SIGPLAN Workshop on *Approaches and Applications of Inductive Programming* (AAIP'09), affiliated with ICFP 2009, Edinburgh, Scotland, September 2009. Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik Nr. 81, Bamberg University, September 2009. ISSN 0937-3349.
- Nr. 82 (2009) Ute Schmid, Marco Ragni, Markus Knauff (Eds.): Proceedings of the KI 2009 Workshop *Complex Cognition*, Paderborn, Germany, September 15, 2009. Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik Nr. 82, Bamberg University, October 2009. ISSN 0937-3349.
- Nr. 83 (2009) Andreas Schönberger, Christian Wilms and Guido Wirtz: A Requirements Analysis of Business-to-Business Integration. Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik Nr. 83, Bamberg University, December 2009. ISSN 0937-3349.
- Nr. 84 (2010) Werner Zirkel and Guido Wirtz: A Process for Identifying Predictive Correlation Patterns in Service Management Systems. Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik Nr. 84, Bamberg University, February 2010. ISSN 0937-3349.
- Nr. 85 (2010) Jan Tobias Mühlberg und Gerald Lüttgen: Symbolic Object Code Analysis. Bamberger Beiträge zur Wirtschaftsinformatik und Angewandten Informatik Nr. 85, Bamberg University, February 2010. ISSN 0937-3349.