

Binsec/Codex, an abstract interpreter to verify safety and security properties of systems code

Olivier Nicole Matthieu Lemerre Xavier Rival

April, 2021

Abstract

This document describes the internals of BINSEC/CODEX, an analyzer able to verify safety and security properties on machine code, notably used to verify absence of runtime errors and privilege escalation in embedded kernels. After stating our assumptions on the hardware, we give a detailed overview of the abstract domains used in the analysis, and give examples of why each domain is needed.

Contents

1	Introduction	2
2	Concrete semantics of machine code	4
3	Abstract domains used in the analysis	6
3.1	Control flow domain	7
3.2	Numeric domain	9
3.3	Weak shape domain	10
3.3.1	Notations	10
3.3.2	Types and their meaning	10
3.3.3	The single type abstract domain	14
3.3.4	The weak shape abstract domain	15
4	Conclusion	18

1 Introduction

This document describes the internals of BINSEC/CODEX, an analyzer for machine code based on abstract interpretation. BINSEC/CODEX is a module of BINSEC [2]. BINSEC/CODEX was used to analyze and verify absence of runtime errors and privilege escalation on embedded kernels [10].

Abstract interpretation is a practical framework to build program analyzers, as well as a theoretical framework justifying the correctness of such analyzers [1]. An abstract interpreter infers properties of a program in a sound way, i.e. the inferred properties will hold on all possible behaviours of the program. By Rice's theorem [11], such a sound inference cannot be complete, i.e. it may include behaviours that will never appear in the actual program execution.

In the context of BINSEC/CODEX, we will concentrate on state properties. A state property is a predicate on a program state. What exactly constitutes a program state is defined in Section 2, but it can be seen as the state of all variable and memory locations of the program, as well as the program counter (the current code location). The absence of runtime error, for instance, is a state property. The absence of use-after-free bugs in C is not.

Inferring state properties using abstract interpretation consists in propagating an abstract state that represents a superset of all possible states, until a fixpoint is reached. It can be seen as a generalization of data flow analysis [6] with more general domains and widening operators to handle loops.

We will use the code in Figure 1 as a running example. It is an example of memory protection initialization in a simple x86 32-bit embedded OS. Each task is allocated a fixed region of memory describe by the fields `mem_base` and `mem_size` of a structure `context`. A global variable `params` contains the array of all contexts, as well as the number of tasks (assumed constant throughout kernel execution).

Memory protection in 32-bit x86 depends on a global descriptor table (`desc_table`) containing segment descriptors. Each segment descriptor encodes a range of memory addresses, along with the permissions associated with these addresses. Among other things, a segment descriptor contains a base address and a length. What the code in Figure 1 does is creating a descriptor for each task and writing it to the descriptor table. The code of the `create_descriptor` function is omitted for now.

The goal is to verify that all segment descriptor written to the descriptor table describe a memory region that is disjoint from the kernel space.

```

typedef struct context {
    uint32_t mem_base;
    uint32_t mem_size;
    // ...
} context;

struct params {
    unsigned int nb_tasks;
    context *contexts;
    uint64_t *desc_table;
    // ...
};

struct params *params;
unsigned int i;
context *ctx;

for(i = 0; i < params->nb_tasks; i++)
{
    ctx = &params->contexts[i];
    params->desc_table[i] = create_descriptor(ctx->mem_base,
                                             ctx->mem_size);
}

```

Figure 1: Running example

If that is achieved, and that you also verify that there is no run-time error, then it entails that there cannot be a privilege escalation exploit on this kernel —except via means not considered in our hardware model, such as side channels. This verification must be performed on the binary code resulting of the compilation of the C code (see Figure 1 below), firstly because it removes the compiler from the trusted components, and secondly because such system code necessarily comprises hand-written, architecture-specific assembly (e.g. to modify memory protection). In what follows, we will see which challenges such a verification poses.

After expliciting our hardware model (Section 2), we will detail the set of abstract domains we use for the analysis (Section 3), including a novel memory domain.

2 Concrete semantics of machine code

By concrete semantics, we mean the way we model the architecture, and the semantics of analyzed programs on this architecture. A semantics consists in two things: defining what is a program state, and specifying how to transition from one state to the next.

We assume a 32-bit architecture. Values are elements of $V_{32} = [0, 2^{32} - 1]$ or $V_8 = [0, 2^8 - 1]$. The set of memory addresses A is a subset of V_{32} . Memories are maps from addresses to 8-bit values: $M = A \rightarrow V_8$. We denote by R the set of register names. In 32-bit x86, it would contain for example `eax` and `ebx`, but also `CPL` (current privilege level), which is a system register only accessible through some system instructions.

The set of states is $\mathcal{S} = M \times (R \rightarrow V_{32})$. A state includes the current code location, via the `eip` register. For $s \in \mathcal{S}$, we denote by $s[r \leftarrow x]$ the state identical to s , except that register r holds the value x ; and $s[[\alpha]_n \leftarrow x]$, the state identical to s , except that the address range $\alpha, \alpha + 1, \dots, \alpha + n - 1$ holds the value x (in the present case, in little-endian convention).

Figure 2 shows the binary and the assembly produced by `gcc -O3` on the running example. All control flow and type information is stripped away. The `for` loop is replaced by a simple conditional backward jump at address `0x90`.

DBA BINSEC translates machine code into an intermediate representation called **DBA** (for Dynamic Bitvector Automata), a simple imperative language with variables, memory accesses and arbitrary jumps. The syntax of **DBA** is given in Figure 3. The semantics of **DBA** is completely standard,

```

40: a1 1c c0 04 08      mov     eax,ds:0x804c01c
45: 8b 10                mov     edx,DWORD PTR [eax]
47: 85 d2                test    edx,edx
49: 74 54                je     9f <main+0x5f>
4b: 8d 4c 24 04          lea    ecx,[esp+0x4]
4f: 83 e4 f0             and    esp,0xffffffff0
52: ff 71 fc             push   DWORD PTR [ecx-0x4]
55: 55                   push   ebp
56: 89 e5                mov     ebp,esp
58: 56                   push   esi
59: 53                   push   ebx
5a: 31 db                xor     ebx,ebx
5c: 51                   push   ecx
5d: 83 ec 0c             sub    esp,0xc
60: 8b 50 04             mov     edx,DWORD PTR [eax+0x4]
63: 8d 0c dd 00 00 00 00 lea    ecx,[ebx*8+0x0]
6a: 8b 70 08             mov     esi,DWORD PTR [eax+0x8]
6d: 83 ec 08             sub    esp,0x8
70: 83 c3 01             add    ebx,0x1
73: 01 ca                add    edx,ecx
75: 01 ce                add    esi,ecx
77: ff 72 04             push   DWORD PTR [edx+0x4]
7a: ff 32                push   DWORD PTR [edx]
7c: e8 3f 01 00 00      call   80491c0 <create_descriptor>
81: 83 c4 10             add    esp,0x10
84: 89 06                mov     DWORD PTR [esi],eax
86: a1 1c c0 04 08      mov     eax,ds:0x804c01c
8b: 89 56 04             mov     DWORD PTR [esi+0x4],edx
8e: 39 18                cmp    DWORD PTR [eax],ebx
90: 77 ce                ja     60 <main+0x20>
92: 8d 65 f4             lea    esp,[ebp-0xc]
95: 31 c0                xor    eax,eax
97: 59                   pop     ecx
98: 5b                   pop     ebx
99: 5e                   pop     esi
9a: 5d                   pop     ebp
9b: 8d 61 fc             lea    esp,[ecx-0x4]
9e: c3                   ret
9f: 31 c0                xor    eax,eax
a1: c3                   ret

```

Figure 2: Running example compiled by GCC with -O3.

Expressions

$\mathcal{E} \ni e, e_{\text{loc}} ::=$	c	constant ($c \in \mathbb{N}$)
	$ r$	register ($r \in \mathbb{R}$)
	$ [e_{\text{loc}}]$	memory read
	$ \Theta e$	unary operation ($\Theta \in \{-, \neg\}$)
	$ e \oplus e$	binary operation ($\oplus \in \{+, -, \div, ==, \leq, \dots\}$)

Instructions

$\mathcal{I} \ni i ::=$	$r := e$	register assignment ($r \in \mathbb{R}$)
	$ [e_{\text{loc}}] := e$	memory write
	$ \text{jump } e_{\text{loc}}$	jump to address
	$ \text{if } e \text{ then } i \text{ else } i \text{ end}$	conditional

Figure 3: Syntax of DBA.

so we will not detail it.

Multicore systems In a multicore system, the state is an element of $M \times (\mathbb{R} \rightarrow V_{32})^n$, where n is the number of cores. Each core has a distinct set of registers —except for system registers which are common, but for simplicity’s sake we will not consider them here. The memory is shared. Each CPU performs an independent execution governed by the same semantics.

3 Abstract domains used in the analysis

Abstract interpretation revolves around the concept of abstract domains. An abstract domain is a set, the elements of which abstract elements. Each abstract element represents a set of concrete elements (e.g. a set of program states) through a concretization function, generally denoted γ . Following tradition, we will write abstract domains with a superscript sharp symbol (#). For instance, if the abstract domain $\mathbb{D}^\#$ represents sets of program states, its concretization has type:

$$\gamma_{\mathbb{D}} : \mathbb{D}^\# \rightarrow \mathcal{P}(\mathcal{S})$$

To be usable in static analysis, an abstract domain must be equipped with a partial order, representing precision. An abstract element is smaller than

```

40 :  eax := [ds + 0x804c01c]
      jump 0x45
45 :  edx := [ds + eax]
      jump 0x47
47 :  res32 := edx - edx
      ZF := (res32 == 0)
      CF := 0
      OF := 0
      jump 0x49
49 :  if ZF then jump 0x9f else jump 0x4b end

```

Figure 4: First few instructions from Figure 2 translated to DBA.

numeric domain	$\mathbb{N}^\# = \text{any conjunction of numeric constraints over the values bound to } A_F \text{ and } \mathcal{R}$	$\gamma_N : \mathbb{N}^\# \rightarrow \mathcal{P}(R \uplus A_F \rightarrow V_{32})$
weak shape domain	$\mathbb{W}^\# = \mathbb{N}^\# \times (R \uplus A_F \rightarrow \mathbb{T}^\#)$	$\gamma_W : \mathbb{W}^\# \rightarrow \mathcal{P}(\mathcal{S})$
control flow domain	$\mathbb{C}^\# = \mathcal{P}(\mathcal{G}) \times (\mathbb{L} \rightarrow \mathbb{W}^\#)$	$\gamma_C : \mathbb{C}^\# \rightarrow \mathcal{P}(\mathcal{S})$

$\mathcal{G} = \mathcal{P}(\mathbb{L} \times \mathbb{L})$ is the set of control flow graphs,
 $\mathbb{T}^\#$ is the single type abstract domain,
 A_F the set of fixed addresses in the kernel,
 R the set of register names, \mathbb{L} the set of program locations.

Figure 5: Implementation of the abstract domains.

another ($a \leq^\# b$) only if the concrete set that a represents is a subset of the one b represents:

$$\forall a, b \in \mathbb{D}^\#, a \leq^\# b \implies \gamma_D(a) \subseteq \gamma_D(b)$$

In this section, the abstract domains that we use will be detailed. They are summarized in Figure 5.

3.1 Control flow domain

Central to our abstractions is the notion of program locations. What is a program location is an implementation choice of the analyzer: a natural choice is to consider that a program locations is a valid code address in the program's executable. For our needs, we chose a more precise abstraction: a program location consists in a kernel address and a call stack. Denoting

the set of valid code addresses by $A_C \subseteq A$ and E^* the set of finites sequences of elements of a set E , we get:

$$\mathbb{L} = A_C \times A_C^*$$

We denote by $\mathcal{L}(s)$ the program location of a state s .

Then, our control flow abstraction is a graph between program locations (we denote $\mathcal{G} = \mathcal{P}(\mathbb{L} \times \mathbb{L})$ the set of such graphs) and a mapping from every location to an abstract state $M^\#$ (described further):

$$\mathbb{C}^\# = \mathcal{P}(\mathcal{G}) \times (\mathbb{L} \rightarrow M^\#)$$

The graph is an over-approximated control flow graph (CFG), meaning that (1) the only reachable program locations are the nodes in the graph, and (2) all possible control flow transfer are represented by an edge in the graph. Formally:

$$\begin{aligned} \gamma : \mathbb{C}^\# &\rightarrow \mathcal{P}(\mathcal{S}) \\ \gamma(\mathcal{G}, \text{states}) &= \{s \in \mathcal{S} \mid \exists \ell \in \mathbb{L}, \text{states}(\ell) = s \wedge \ell \text{ is a node of } \mathcal{G}\} \end{aligned}$$

The analysis works by performing multiple rounds of the following steps in sequence:

1. Perform a standard data-flow analysis using the current abstract CFG, to compute a new state for every location of the graph.
2. Iterate over all locations ℓ to compute all possible outgoing edges, given the possible states at ℓ (this uses the same `resolve` function than [7]). Newly-discovered edges are added to the over-approximated CFG.

The iteration sequence starts with an abstraction (\mathcal{G}_0, m_0) where \mathcal{G}_0 is a single node (the start instruction address), and m_0 maps this address to the initial abstract state. The analysis terminates when the fixpoint is reached, i.e., no new edge is discovered in the CFG. In practice, several small optimisations are used to reuse results between rounds (e.g., caching the results), and to have fewer rounds (by early exploration of the newly-discovered CFG nodes).

Theorem 1. *If the transfer functions for $M^\#$ are sound, the result $s_{final}^\#$ of the analysis is a sound abstraction of all the reachable states in the system (and thus a state invariant).*

In the running example, the analysis proceeds as follows:

1. While instruction 0x90 is not reached, every instruction has the next instruction as its only successor, so the graph is linear.
2. After analyzing instruction 0x90 for the first time, the conditional is false (unless `params->nb_tasks` is equal to 1, but during the analysis we do not have this information), so there is only one successor, 0x60. This constitutes a back edge in the CFG (i.e. an edge to an ancestor in a spanning tree of the CFG). This characterizes a control flow loop, and therefore the control flow domain chooses a loop head and inserts a widening point.
3. After a fixpoint is reached for every location in the loop, there should be a new successor to 0x90 (since the loop is not an infinite one), namely 0x92, from which new paths can be discovered.

It should be noted that the choice of including the call stack in the abstract locations has the consequence that all function calls are inlined. On the one hand, it allows for a precise interprocedural analysis, since a function is always analyzed in the most precise possible calling context. On the other hand, it prevents the analyzer to handle some recursive functions (except when all recursive calls are tail calls and are optimized into jumps), because the analyzer may compute abstract states with ever-growing call stacks, without reaching a stable invariant. Since low-level code rarely contains recursive code, this is usually not a problem.

3.2 Numeric domain

What we call numeric domain, in this context, is an abstraction of both numeric values and memory. We call this domain $\mathbb{N}^\#$ and it represents the contents of all memory cells which we want to fully enumerate, and registers. The designation “all memory cells which we want to fully enumerate” matters because part of memory will be represented only by type information in the weak shape domain, and not explicitly represented at the byte level.

The concretization of the numeric domain thus has type:

$$\gamma_{\mathbb{N}} : \mathbb{N}^\# \rightarrow \mathcal{P}(\mathbb{R} \uplus A_{\mathbb{F}} \rightarrow V_{32})$$

where $A_{\mathbb{F}} \subseteq A$. (The F stands for “flat memory model”.)

In principle, any domain that has such a concretization could be used. In practice, of course, it must be sufficiently precise to prove the properties we want. We use a combination of standard abstractions, the full description of which is not the intent of this report.

- For values, we mainly use efficient non-relational abstract domains: intervals, based on the reduced product between the signed and unsigned meaning of bitvectors [3], with congruence information [5]. They are complemented with symbolic relational information [3, 9, 4] for local simplifications of sequences of machine code.
- Regarding the memory abstraction: our memory model is ultimately byte-level in order to deal with very low-level coding aspects of kernels. Yet, as representing each memory byte separately is inefficient and imprecise, we use a stratified representation of memory caching multi-byte loads and stores, like Miné [8]. Moreover, we do not track memory addresses whose contents is unknown.

3.3 Weak shape domain

3.3.1 Notations

We define a bitvector concatenation operation. It corresponds to the idea of joining two binary representations next to each other and interpreting them as a single integer. It depends the bit length of each operand.

$$\forall v_1, v_2, s_1, s_2 \in \mathbb{N}, v_1 \mathbin{s_1} \mathbin{s_2} v_2 = v_1 + 2^{s_1} v_2$$

Intuitively, s_1 and s_2 are the lengths of the two bit vectors v_1 and v_2 . We may omit the lengths when they are clear from the context. We note $\mathbb{H} = A \rightarrow V_8$ the set of heaps. For $h \in \mathbb{H}$, we will write indifferently $h(a)$ or $h[a]$. We will write $h[a, a + n]$ to mean $h[a] :: h[a + 1] :: \dots :: h[a + n]$.

3.3.2 Types and their meaning

We want to abstract properties of memory structures using types. To that end, we are going to define a type system encompassing the low-level types of the C language: scalars, pointers, and structures (which, following the conventions of type system research, we will call product types). We do not have so-called sum types, and thus cannot express C unions; however, nothing in principle prevents sum types to be added. One of the nice features of this system is that it is easy to extend it by enriching the grammar of types, and make the necessary adjustments to the definitions and proofs. By enriching the type system, one directly enriches the abstract domain that derives from it.

Speaking of enriching the types, we also define a form of refinement types, i.e. types augmented with predicates. Such types do not exist in

C, but by using them we can specify –and verify– useful properties on values. The predicates can be any unary predicates from a given logic. In our analyzer, predicates are from a simple, first-order logic containing equality and inequality operators, as well as usual operators from integer arithmetic (addition, multiplication, ...) and bitwise operators (and, xor, ...).

The grammar of types is:

$\mathbb{T} \ni t$	$::= $	<code>byte</code> <code>word</code>	base types
		t_a^*	pointer
		$t \times t$	product type
		$\{x : t \mid p(x)\}$	refinement type with a predicate p
		$t[s]$	array type ($s \in \mathbb{N}$)
$\mathbb{T}_A \ni t_a$	$::= $	$t.o$	cell type ($o \in \mathbb{Z}$)

Each type has a size (in bytes) given by the function:

$$\begin{aligned}
 \text{size} &: \mathbb{T} \rightarrow \mathbb{N} \\
 \text{size}(\text{byte}) &= 1 \\
 \text{size}(\text{word}) &= 4 \\
 \text{size}(t^*) &= 4 \\
 \text{size}(t_1 \times t_2) &= \text{size}(t_1) + \text{size}(t_2) \\
 \text{size}(\{x : t \mid p(x)\}) &= \text{size}(t) \\
 \text{size}(t[s]) &= s \cdot \text{size}(t)
 \end{aligned}$$

The size of `word` being equal to 4 comes from the fact that the architecture is 32-bit. We will give the meaning of types in terms of sets of values shortly, but for that we need to introduce labellings.

Definition 2 (Labelling). *A labelling \mathcal{L} is a function of $A \rightarrow \mathbb{T}_A$ such that all instances of a type are whole and contiguous in memory, i.e. for all type $t \in \mathbb{T}$ and address $a \in A$, if we define $n = \text{size}(t)$:*

$$(\exists k \in [0, n - 1], \mathcal{L}(a + k) = t.k) \implies \begin{cases} \mathcal{L}(a) = t.0 \\ \mathcal{L}(a + 1) = t.1 \\ \vdots \\ \mathcal{L}(a + n - 1) = t.(n - 1) \end{cases}$$

The set of labellings is denoted `Lab`.

Two address types can express similar things, with one being more precise than the other. Consider a C structure like the following:

```

struct s {
    uint8_t a;
    uint32_t *b;
};

```

In the language of our types, this would be expressed: $s = \text{byte} \times \text{word}.0^*$. Now consider the type of the first memory cell in such a structure: the type $s.0$. In some sense, an $s.0$ is also a $\text{byte}.0$. However, $s.0$ is more precise than $\text{byte}.0$, in the sense that not all memory cells that contain a byte are necessarily part of a structure s . We say that $s.0$ subsumes $\text{byte}.0$.

Definition 3 (Subsumption relation). *A cell type (an element of \mathbb{T}_A) can subsume another cell type in the following conditions:*

$$(t_1 \times t_2).k \text{ subsumes } t_1.k \text{ if } 0 \leq k < \text{size}(t_1)$$

$$(t_1 \times t_2).(\text{size}(t_1) + k) \text{ subsumes } t_2.k \text{ if } 0 \leq k < \text{size}(t_2).$$

$$\{x : t \mid p(x)\}.k \text{ subsumes } t.k$$

$$t[s].k \text{ subsumes } t.o, \text{ if } \begin{cases} k = q \cdot \text{size}(t) + o \\ 0 \leq o < \text{size}(t) \end{cases}$$

We define the relation $\preceq \in \mathbb{T}_A \times \mathbb{T}_A$ as the transitive reflexive closure of the relation “subsumes”.

Note that defining \preceq as reflexive requires a notion of equality on \mathbb{T}_A . We use for that the equality that can be defined inductively on the grammar in a straightforward way.

Lemma 4. *For all types $u, a, b \in \mathbb{T}_A$, if u subsumes a and u subsumes b , then $a = b$.*

Proof. By the definition of “subsumes”, necessarily u is one of the following:

- $u = \{x : t \mid p(x)\}.k$. Then necessarily $a = b = t.k$.
- $u = t[s].k$. Then it is easy to show that $a = b$.
- $u = (t_1 \times t_2).k$. Also from the definition of subsumption, a is either the type $t_1.k$ or the type $t_2.(k - \text{size}(t_1))$. Same for b . Let us proceed *ab absurdum* and assume $a \neq b$. Then, there are two symmetrical possibilities. Let us devise the case where $a = t_1.k$ and $b = t_2.(k - \text{size}(t_1))$.

Again from Definition 3, we deduce:

$$\begin{cases} 0 \leq k < \text{size}(t_1) \\ 0 \leq k - \text{size}(t_1) < \text{size}(t_2) \end{cases}$$

from which we derive two contradictory statements: $k < \text{size}(t_1)$ and $k \geq \text{size}(t_1)$. Therefore $a = b$.

□

We now give the meaning of types in terms of an interpretation function.

Definition 5 (Interpretation of a type). *The interpretation operator with respect to a labelling \mathcal{L} , denoted $\langle \cdot \rangle_{\mathcal{L}} : \mathbb{T} \rightarrow \mathbb{N}$, is defined by:*

$$\begin{aligned} \langle \text{byte} \rangle_{\mathcal{L}} &= [0, 2^8 - 1] \\ \langle \text{word} \rangle_{\mathcal{L}} &= [0, 2^{32} - 1] \\ \langle t_1 \times t_2 \rangle_{\mathcal{L}} &= \{v_1 \text{ size}(t_1) : \text{size}(t_2)v_2 \mid v_1 \in \langle t_1 \rangle_{\mathcal{L}}, v_2 \in \langle t_2 \rangle_{\mathcal{L}}\} \\ \langle \{x : t \mid p(x)\} \rangle_{\mathcal{L}} &= \{v \in \langle t \rangle_{\mathcal{L}} \mid p(v)\} \\ \langle t[s] \rangle_{\mathcal{L}} &= \{v_1 :: \dots :: v_s \mid v_1, \dots, v_s \in \langle t \rangle_{\mathcal{L}}\} \\ \langle t.k^* \rangle_{\mathcal{L}} &= \{a \in A ; \mathcal{L}(a) \preceq t.k\} \cup \{0\} \end{aligned}$$

A labelling is a labelling for a heap h when values are laid out in memory in a manner consistent with their types.

Definition 6. *We say that a labelling $\mathcal{L} \in \text{Lab}$ is a labelling for $h \in \mathbb{H}$ if:*

$$\forall a \in A, \mathcal{L}(a) = t.0 \implies h[a, a + \text{size}(t) - 1] \in \langle t \rangle_{\mathcal{L}}$$

Note that by Definition 5, we have

$$t.n \preceq u.m \implies \langle t.n^* \rangle_{\mathcal{L}} \subseteq \langle u.m^* \rangle_{\mathcal{L}}$$

So \preceq can be seen as a subtyping relation on pointers.

Definition 7 (Set of addresses of a given type).

$$\begin{aligned} \text{addr} &: \text{Lab} \times \mathbb{T} \rightarrow \mathcal{P}(A) \\ \text{addr}_{\mathcal{L}}(t) &= \bigcup_{i=0}^{\text{size}(t)-1} \{a \in A; \mathcal{L}(a) \preceq t.i\} \end{aligned}$$

In any reasonable program of any language, values of different types should reside in different zones of memory. More precisely, *incompatible* types should reside in different zones of memory. What makes one type “compatible” with another is a subsumption relation between them. If some struct contains an `uint32` at offset 4, then it is expected that, wherever that struct exists in memory, the object existing at bytes 4 to 8 of the struct should be interpretable as an `uint32`. Therefore, we define a labelling as separated when only the types that are in a subsumption relation can be attributed to the same cells.

Definition 8 (Separated labellings). A labelling $\mathcal{L} \in \mathbb{A} \rightarrow \mathcal{P}(\mathbb{T}_{\mathbb{A}})$ is separated if incomparable cell types cannot be labels of the same cell:

$$\text{addr}_{\mathcal{L}}(t) \cap \text{addr}_{\mathcal{L}}(u) \neq \emptyset \implies \exists n, m, t.n \leq u.m \vee u.m \leq t.n$$

We will prove that this desirable property holds in fact for all labellings. But first, a few properties of \leq need to be shown.

Proposition 9. $(\mathbb{T}_{\mathbb{A}}, \leq)$ is a partial order.

Proof. \leq is reflexive and transitive by definition; it is also antisymmetric, because the elements of $\mathbb{T}_{\mathbb{A}}$, by their definition, cannot contain themselves. \square

Lemma 10. For any cell types $\tau, \upsilon, \phi \in \mathbb{T}_{\mathbb{A}}$:

$$\phi \leq \tau \wedge \phi \leq \upsilon \implies \tau \leq \upsilon \vee \upsilon \leq \tau$$

Proof. Let us assume that $\phi \leq \tau$ and $\phi \leq \upsilon$. Then, if we write $x \xrightarrow{s} y$ for “ x subsumes y ”, there exists two finite chains $\phi \xrightarrow{s} \tau_1 \xrightarrow{s} \tau_2 \xrightarrow{s} \dots \xrightarrow{s} \tau_n \xrightarrow{s} \tau$, and $\phi \xrightarrow{s} \upsilon_1 \xrightarrow{s} \upsilon_2 \xrightarrow{s} \dots \xrightarrow{s} \upsilon_m \xrightarrow{s} \upsilon$. Either one chain is included in the other, or not. If yes, then either $\tau \leq \upsilon$ or $\upsilon \leq \tau$, which ends the proof. If not, then let p be the minimal index such that $\tau_p \neq \upsilon_p$. But τ_{p-1} subsumes both τ_p and υ_p , so using Lemma 4, $\tau_{p-1} = \upsilon_{p-1}$, which contradicts the minimality of p . \square

In other words, the graph of the order relation $(\mathbb{T}_{\mathbb{A}}, \leq)$ is a forest.

Theorem 11. All labellings are separated.

Proof. Let us assume that the intersection of $\text{addr}_{\mathcal{L}}(t)$ and $\text{addr}_{\mathcal{L}}(u)$ contains some address a . Then there exists i and j such that $\mathcal{L}(a) \leq t.i$ and $\mathcal{L}(a) \leq u.j$. Thus by Lemma 10, either $t.i \leq u.j$, or $u.j \leq t.i$. \square

3.3.3 The single type abstract domain

Here we show how \mathbb{T} , supplemented with “top” and “bottom” elements, can be seen as an abstraction for a set of values. Its concretization is only defined relatively to a labelling \mathcal{L} .

Definition 12. Given a labelling $\mathcal{L} \in \text{Lab}$, the single type abstract domain is:

$$\mathbb{T}^{\#} = \mathbb{T} \uplus \{\perp, \top\}$$

Its concretization is:

$$\begin{aligned} \gamma_{\mathcal{L},\mathbb{T}} &: \mathbb{T}^\# \rightarrow \mathcal{P}(\mathbb{N}) \\ \gamma_{\mathcal{L},\mathbb{T}}(\perp) &= \emptyset \\ \gamma_{\mathcal{L},\mathbb{T}}(t) &= \langle t \rangle_{\mathcal{L}} \\ \gamma_{\mathcal{L},\mathbb{T}}(\top) &= \mathbb{N} \end{aligned}$$

Its abstract inclusion $\sqsubseteq_{\mathcal{L},\mathbb{T}}$ is defined by:

- $\perp \sqsubseteq_{\mathcal{L},\mathbb{T}} t$
- $t \sqsubseteq_{\mathcal{L},\mathbb{T}} \top$
- $t.n \leq u.n \implies t.n^* \sqsubseteq_{\mathcal{L},\mathbb{T}} u.m^*$

Its lub $\sqcup_{\mathcal{L},\mathbb{T}}$ is the one induced by $\sqsubseteq_{\mathcal{L},\mathbb{T}}$.

Proposition 13. $\sqsubseteq_{\mathcal{L},\mathbb{T}}$ and $\sqcup_{\mathcal{L},\mathbb{T}}$ are sound with respect to $\gamma_{\mathbb{T}}$.

3.3.4 The weak shape abstract domain

The weak shape abstract domain $\mathbb{W}^\#$ is a memory domain which abstracts part of the memory using types. The weak shape domain is essentially an augmentation upon another domain $\mathbb{M}^\#$, which concretizes to $\mathcal{P}(\mathbb{R} \uplus A_F \rightarrow V_{32})$, where $A_F \subseteq A$. (In the context of our analyzer, we instantiate $\mathbb{M}^\#$ to the domain $\mathbb{N}^\#$ described above.)

$$\mathbb{W}^\# = \mathbb{M}^\# \times (\mathbb{R} \uplus A_F \rightarrow \mathbb{T}^\#)$$

Its concretization has type:

$$\gamma_{\mathbb{W}} : \mathbb{W}^\# \rightarrow \mathcal{P}(\mathbb{R} \uplus A_F \uplus A_T \rightarrow V_{32})$$

where A_F and A_T are disjoint. If one assumes that $A_F \uplus A_T = A$, then the weak shape domain is, in fact, an abstraction of a set of states: the concretization has type $\mathbb{W}^\# \rightarrow \mathcal{P}(\mathcal{S})$.

We see here that $\mathbb{W}^\#$ augments $\mathbb{M}^\#$ with the abstraction of a new memory zone A_T , disjoint from A_F . This disjoint heap is represented only by types. A type is associated to every register, and memory cell indexed by A_F . But for this representation to have meaning, it is necessary that the values the fixed memory be consistent with their types, relatively to a certain labelling: we say that $h_F \in \mathbb{R} \uplus A_F \rightarrow V_8$ and $\ell^\# \in A_F \rightarrow \mathbb{T}^\#$ are \mathcal{L} -consistent if $\forall a \in A_F, h(a) \in \gamma_{\mathcal{L},\mathbb{T}}(\ell^\#(a))$.

Let $(h^\#, \ell^\#) \in \mathbb{W}^\#$. We define

$$\gamma_{\mathbb{W}}(h^\#, \ell^\#) = \emptyset$$

if $\gamma_M(h^\#)$ and $\ell^\#$ are not \mathcal{L} -consistent for any labelling \mathcal{L} . Otherwise:

$$\gamma_W(h^\#, \ell^\#) = \{h_F \uplus h_T \mid \begin{array}{l} h_F \in \gamma_M(h^\#) \\ \wedge \text{dom}(h_F) \cap \text{dom}(h_T) = \emptyset \\ \wedge \exists \mathcal{L} \in \text{Lab}, \mathcal{L} \text{ is a labelling for } h_T \\ \wedge h_F \text{ and } \ell \text{ are } \mathcal{L}\text{-consistent} \\ \wedge \forall a \in A_F, \mathcal{L}(a) \sqsubseteq_{\mathcal{L}, \mathbb{T}} \ell^\#(a) \end{array}\}$$

Transfer functions The store and load transfer functions need to preserve 1. the separation between the flat heap and the typed heap and 2. the types of the typed heap. The load function is easier. Its type is:

$$\text{load} : W^\# \times (\mathbb{N}^\# \times \mathbb{T}^\#) \times \mathbb{N} \rightarrow (\mathbb{N}^\# \times \mathbb{T}^\#)$$

where the parameters are the weak shape abstract heap, a typed address with a numeric and a type component, and the size of the region to read, respectively. It returns a typed value, i.e. with a numeric and a type component. $\text{load}((h^\#, \ell^\#), (n^\#, t^\#), s)$ proceeds as follows:

- If $t^\#$ is a pointer type, or a subtype of a pointer type, i.e. if there exists $u \in \mathbb{T}$ such that $t^\# \sqsubseteq_{\mathcal{L}, \mathbb{T}} u.0^*$, then the abstract address concretizes to a subset of A_T . The numeric component $n^\#$ is ignored, and the load returns $(\llbracket u \rrbracket_{\mathcal{L}}, u)$.
- Otherwise, if $\gamma_N(n^\#) \subseteq A_F$, then the type component of the address is ignored and the load is forwarded to $h^\#$. We return $(\text{load}_M(h^\#, n^\#, s), \top)$.
- Otherwise, \top is returned (or an alarm is emitted).

store has type:

$$\text{store} : W^\# \times (\mathbb{N}^\# \times \mathbb{T}^\#) \times \mathbb{N} \times (\mathbb{N}^\# \times \mathbb{T}^\#) \rightarrow W^\#$$

where the arguments are the weak shape abstract heap, a typed address with a numeric and type component, the size of the region to read, and a typed value to store, respectively. $\text{store}((h^\#, \ell^\#), (n_a^\#, t_a^\#), s, (n_v^\#, t_v^\#))$ proceeds as follows:

- If $t_a^\#$ is a pointer type, or a subtype of a pointer type, i.e. if there exists $u \in \mathbb{T}$ such that $t_a^\# \sqsubseteq_{\mathcal{L}, \mathbb{T}} u.0^*$, then the abstract address points into the typed heap. We check that the types match between value and address, i.e. that $t_v^\# \sqsubseteq_{\mathcal{L}, \mathbb{T}} u$, and the abstract heap $(h^\#, \ell^\#)$ is returned unchanged. Otherwise, if the types don't match, \top is returned (or an alarm is emitted).

- Otherwise, if $\gamma_N(n_a^\#) \subseteq A_F$, then the store is performed in the memory subdomain: we return $(\text{store}_M(h^\#, n_a^\#, s, n_v^\#), \ell_{\text{new}}^\#)$ where $\ell_{\text{new}}^\#$ is the result of updating $\ell^\#$: $\forall a \in \gamma_N(n_a^\#), \ell_{\text{new}}^\#(a) = t_v^\#$.
- Otherwise, we return \top (or an alarm is emitted).

Let us see how this domain enables us to verify our motivating example (Figure 1): all parameters (number of tasks, allocated regions, address of the task array) can be numerically unknown, we still can perform the verification thanks to the weak shape domain. We use the following set of type definitions:

```

type segment_base = {x : uint32 | x > kernel_last_addr}
type context = segment_base × uint32
type desc = {x : uint64 | base(x) > kernel_last_addr}
type params = uint32 × context[N_tasks].0 * × desc[N_tasks].0*

```

Here, `kernel_last_addr` and `N_tasks` are what we call symbolic constants: they represent values which are not known precisely at the time of the analysis. It is simple to extend the type syntax to use such symbolic constants in type predicates and as lengths of arrays; we did not detail this point in the formalization above for clarity reasons.

- `kernel_last_addr` represents the highest address of the kernel memory. By constraining all segment bases to be greater than this constant, we make sure that the memory regions accessible to tasks never intersect kernel memory.
- `N_tasks` represents the number of tasks that run on the system. Following a common trend in embedded systems design, that number is a constant, i.e. to change the number of tasks running on the kernel, one must compile a new application image and link it with the kernel.

With this set of type definitions, the weak shape domain is able to analyze the machine code resulting from the compilation of Figure 1 and verify:

1. The absence of out-of-bounds array accesses, and other run-time errors.
2. That the two predicates constraining memory protection segments hold.

Combined, these two properties suffice to prove the absence of privilege escalation on this example kernel.

4 Conclusion

By combining standard abstract interpretation techniques with a novel weak shape abstract domain, we were able to build an analyzer to verify safety and security properties directly on machine code efficiently, with a low annotation burden.

References

- [1] Patrick Cousot and Radhia Cousot. “Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints.” In: *POPL*. 1977.
- [2] Robin David et al. “BINSEC/SE: A Dynamic Symbolic Execution Toolkit for Binary-Level Analysis.” In: *SANER*. 2016 IEEE 23rd International Conference on Software Analysis, Evolution and Reengineering (SANER). Suita: IEEE, Mar. 2016, pp. 653–656.
- [3] Adel Djoudi, Sébastien Bardin, and Éric Goubault. “Recovering High-Level Conditions from Binary Programs.” In: *FM 2016: Formal Methods*. Ed. by John Fitzgerald et al. Lecture Notes in Computer Science. Springer International Publishing, 2016, pp. 235–253.
- [4] Graeme Gange et al. “An Abstract Domain of Uninterpreted Functions.” In: *VMCAI*. Verification, Model Checking, and Abstract Interpretation. Ed. by Barbara Jobstmann and K. Rustan M. Leino. Vol. 9583. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 85–103.
- [5] Philippe Granger. “Static Analysis of Arithmetical Congruences.” In: *Int. J. Comput. Math.* 30.3-4 (1989), pp. 165–190.
- [6] Gary A. Kildall. “A Unified Approach to Global Program Optimization.” In: *Proceedings of the 1st Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages - POPL '73*. The 1st Annual ACM SIGACT-SIGPLAN Symposium. Boston, Massachusetts: ACM Press, 1973, pp. 194–206.

- [7] Johannes Kinder, Florian Zuleger, and Helmut Veith. “An Abstract Interpretation-Based Framework for Control Flow Reconstruction from Binaries.” In: *VMCAI. Verification, Model Checking, and Abstract Interpretation*. Ed. by Neil D. Jones and Markus Müller-Olm. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2009, pp. 214–228.
- [8] Antoine Miné. “Field-Sensitive Value Analysis of Embedded C Programs with Union Types and Pointer Arithmetics.” In: *LCTES. Proceedings of the 2006 ACM SIGPLAN/SIGBED Conference on Language, Compilers, and Tool Support for Embedded Systems. LCTES '06*. New York, NY, USA: ACM, 2006, pp. 54–63.
- [9] Antoine Miné. “Symbolic Methods to Enhance the Precision of Numerical Abstract Domains.” In: *VMCAI. Verification, Model Checking, and Abstract Interpretation*. Ed. by E. Allen Emerson and Kedar S. Namjoshi. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2006, pp. 348–363.
- [10] Olivier Nicole et al. “No Crash, No Exploit: Automated Verification of Embedded Kernels.” In: *RTAS. IEEE Real-Time and Embedded Technology and Applications Symposium*. Online, 2021.
- [11] H. G. Rice. “Classes of Recursively Enumerable Sets and Their Decision Problems.” In: *Trans. Amer. Math. Soc.* 74.2 (1953), pp. 358–366. ISSN: 0002-9947, 1088-6850.