



A Multilanguage Static Analysis of Python Programs with Native C Extensions

Raphaël Monat, Abdelraouf Ouadjaout, Antoine Miné

► To cite this version:

Raphaël Monat, Abdelraouf Ouadjaout, Antoine Miné. A Multilanguage Static Analysis of Python Programs with Native C Extensions. Static Analysis Symposium (SAS), Oct 2021, Chicago, Illinois, United States. hal-03313409

HAL Id: hal-03313409




<https://hal.science/hal-03313409>

Submitted on 3 Aug 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Multilanguage Static Analysis of Python Programs with Native C Extensions*

Raphaël Monat¹ , Abdelraouf Ouadjaout¹ , and Antoine Miné^{1,2} 
firstname.lastname@lip6.fr

¹ Sorbonne Université, CNRS, LIP6, F-75005 Paris, France

² Institut Universitaire de France, F-75005, Paris, France

Abstract. Modern programs are increasingly multilanguage, to benefit from each programming language’s advantages and to reuse libraries. For example, developers may want to combine high-level Python code with low-level, performance-oriented C code. In fact, one in five of the 200 most downloaded Python libraries available on GitHub contains C code. Static analyzers tend to focus on a single language and may use stubs to model the behavior of foreign function calls. However, stubs are costly to implement and undermine the soundness of analyzers. In this work, we design a static analyzer by abstract interpretation that can handle Python programs calling C extensions. It analyses directly and fully automatically both the Python and the C source codes. It reports runtime errors that may happen in Python, in C, and at the interface. We implemented our analysis in a modular fashion: it reuses off-the-shelf C and Python analyses written in the same analyzer. This approach allows sharing between abstract domains of different languages. Our analyzer can tackle tests of real-world libraries a few thousand lines of C and Python long in a few minutes.

Keywords: Formal Methods · Static Analysis · Abstract Interpretation · Dynamic Programming Language · Multilanguage Analysis



1 Introduction

Modern programs are increasingly multilanguage. This allows developers to combine the strengths of different languages and reuse libraries written in other languages. A *host* language may call a *guest* language through an interface; this

*This work is partially supported by the European Research Council under Consolidator Grant Agreement 681393 — MOPSA.

interface is also called a *boundary*. The guest language is frequently C and is usually referred to as native code or native C. In this paper, the host language is Python, and the guest language is C. This work supports the Python API [40] as the interoperability mechanism between Python and C.¹ Using native C modules in Python is frequent as it allows writing high-level Python code, itself calling efficient C code. As a matter of fact, one in five of the 200 most downloaded Python libraries available on GitHub contains C code. Although useful, multi-language programs generate additional sources of bugs. Indeed, developers need to take into account different safety mechanisms and memory representations. Python is safe to the extent that runtime errors in pure Python programs are encapsulated into exceptions, which can be caught later on. This safety property breaks when C modules are used since a runtime error in C may irremediably terminate the program or create an inconsistent state. Python and C also have different representations. For example, Python integer objects use at least 24 bytes of memory and have unlimited precision, while C integers have fixed lengths (generally ranging from 8 to 64 bits) and can suffer from overflows.

Static analysis aims at inferring program properties (e.g., the absence of runtime errors) by analyzing programs without executing them. Static analyzers tend to focus on analyzing one language at a time. They may use stubs to model the behavior of calls to other languages. These stubs may be time-consuming to implement if written by hand. They can undermine the soundness of the analyses since the actual code is not analyzed, and the stubs may be imprecise or wrong. For example, our previous work developing a type analysis for Python [33] uses official Python type annotations (defined by PEP 484 [39]) as stubs. While this analysis tracks uncaught Python exceptions, these type annotations do not declare which exceptions may be raised, thus adding an unchecked assumption to the soundness property [33, Section 6.2].

We aim at analyzing both the native C code and the Python code (including callbacks to Python code from the native side) within the same analyzer, called Mopsa [18]. We perform a precise, flow, and context-sensitive value analysis by abstract interpretation. Our analyzer works by induction on the syntax and switches from one language to the other just as the concrete program execution does. We present a multilanguage static analysis built upon pre-existing value static analyses by abstract interpretation [9] of C [37] and Python [33, 34]. It detects runtime errors in the native C code (invalid pointer operations, invalid memory accesses, integer overflows), in the Python code (raised exceptions), and at the boundary between the languages. The underlying address allocation and numeric abstractions are shared. A few multilanguage static analyses exist, and focus mostly on analyzing Java and C ([43, 12, 13, 14, 45, 22], cf. Section 6). They compute summaries of the effects of native code on the chosen abstract property in a bottom-up fashion. Those effects are then translated into the host language, where a standard analyzer for the host language can then be used. The use of

¹Other interoperability mechanisms such as the `ctypes` from the standard library, the `ffi` library, the `cython` project, or the `swig` project all use code using this API or generate code targeting this API. We could thus analyze this generated code.

summaries to convey the abstract meaning of functions makes it easier to rely on independent analyzers for each language. However, the language and properties we target require precise context-sensitive value analyses that are difficult to perform bottom-up. Since Python is a dynamic programming language with a flexible semantics, it is not possible to analyze programs precisely in a context-insensitive fashion. Additionally, a precise description of the Python heap at a native call is mandatory to analyze the called C code, check for pointer errors, and infer effects. We believe the approach described in this paper is general and could be extended to other multilanguage settings, such as the analysis of Java and C through the JNI.

Contributions.

- We define a multilanguage semantics for Python programs with native C modules using the Python C API.
- We show how to lift analyses of Python and C into the multilanguage setting. The underlying address allocation and numeric abstractions are shared, paving the way for relational invariants between Python and C variables.
- We built an implementation on top of an existing static analysis platform called Mopsa that was previously used to design independent analyses for C and for Python. We added support for multilanguage C/Python programs by only adding domains modeling the boundary. We reuse the previous domains analyzing C and Python off-the-shelf. Thanks to this construction, we can detect runtime errors at the boundary, but also in the Python code and in the native C code.
- We evaluate our approach on six real-world libraries found on GitHub. We show that we can scale to libraries of 5,800 combined lines of Python and C code within five minutes.

Artifact. An artifact [35] is available alongside this article. The artifact makes Table 2 and the claim about the percentage of Python packages containing C code on Github (made in Section 1) reproducible. The example codes displayed in Figure 1 and Figure 12 are provided in the artifact, along with instructions to run our analyzer on them. The source code of the modified version of Mopsa is also included. We plan on merging our changes into the public version of Mopsa [19].

Limitations. Our concrete semantics is high-level and makes the assumption that builtin Python objects are only manipulated through the API in C (this assumption is verified by our analyzer). The garbage collection based on reference counting is not supported by our semantics. Thus, we cannot detect deallocations that are performed too soon or that are not performed at all. There is no formal soundness proof that our concrete semantics effectively models the

behavior of the Python interpreter. Potential runtime errors in the API implementation modeled by our concrete semantics as builtins cannot be detected.² Our implementation supports relational analyses, but those do not scale to real-world examples (thus, we use intervals for the numeric abstraction by default). These limitations could be removed in future work.

Outline. We start by showing a self-contained motivating example in Section 2, giving insights on how native C modules are defined and how they work with Python. We define the concrete semantics of these multilanguage programs in Section 3, and explain the abstractions performed in Section 4. Section 5 presents our implementation and the analysis results. We discuss related work in Section 6 and conclude in Section 7.

2 An Extension Module Example

This section provides an in-depth motivating example. We show how to define a native C extension module, and how it can be used by a Python client code. We end the section by discussing which errors may happen.

When developers want to run native C code in Python, they define native C extension modules using the Python API. These modules may contain attributes, methods, and classes, just as any other Python module. However, these methods and classes are now written in C. API functions are denoted by the `Py` prefix (and written in **magenta** in the listings). The semantics of some of these functions are described formally in Section 3.

Counter module, viewed from Python. Our example is a C module defining a `Counter` class, alongside some client code in Python. This example is self-contained and shown in Figure 1. From a high-level point of view, the `counter` module defines a `Counter` class. Instances of `Counter` can be created (`count.py`, line 4); their internal counter can be incremented using the `incr` method, which takes an optional integer argument being the increment (lines 6-7); they also have a read-only attribute `counter` returning their current value (line 8).

Counter, viewed from C. In C, instances of `Counter` will be stored using the `Counter0 struct`. This `struct` starts with a `PyObject ob_base` field. All Python objects are represented as `PyObject`s in C. Putting the `PyObject` as the first field in the `Counter` structure allows casting to and from Python objects.³ The `PyObject` definition is part of the API and shown in Figure 2. Its fields are a reference counter for the garbage collector and a pointer to the class to

²However, half of the API supported by our implementation uses the original C implementation, which is analyzed and where runtime errors would be detected.

³According to the ISO C reference, “a pointer to a structure object, suitably converted, points to its initial member, and vice versa”.

<pre> 1 import counter 2 import random 3 4 c = counter.Counter() 5 p = random.randrange(128) 6 c.incr(2**p-1) 7 c.incr() 8 r = c.counter </pre>	<pre> 40 static PyMemberDef CounterMembers[] = { 41 {"counter", T_INT, offsetof(Counter0, 42 count), READONLY, ""}, {NULL} 43 }; 44 45 static PyType_Slot CounterTSlots[] = { 46 {Py_tp_new, PyType_GenericNew}, 47 {Py_tp_init, CounterInit}, 48 {Py_tp_methods, CounterMethods}, 49 {Py_tp_members, CounterMembers}, {0, 0} 50 }; 51 52 static PyType_Spec CounterTSpec = { 53 .name = "counter.Counter", 54 .basicsize = sizeof(Counter0), 55 .itemsize = 0, 56 .flags = Py_TPFLAGS_DEFAULT 57 Py_TPFLAGS_BASETYPE, 58 .slots = CounterTSlots 59 }; 60 61 static struct PyModuleDef countermod = { 62 PyModuleDef_HEAD_INIT, .m_name = "counter", 63 .m_methods = NULL, .m_size = -1 64 }; 65 66 PyMODINIT_FUNC 67 PyInit_counter(void) 68 { 69 PyObject *m = PyModule_Create(&countermod); 70 if(m == NULL) return NULL; 71 PyObject* CounterT = 72 PyType_FromSpec(&CounterTSpec); 73 if(CounterT == NULL PyModule_AddObject(74 m, "Counter", CounterT) < 0) { 75 Py_DECREF(m); 76 return NULL; 77 } 78 return m; 79 } </pre>
<pre> 9 #include <Python.h> 10 #include "structmember.h" 11 12 typedef struct { 13 PyObject ob_base; 14 int count; 15 } Counter0; 16 17 static PyObject* 18 CounterIncr(Counter0 *self, PyObject *args) 19 { 20 int i = 1; 21 if(!PyArg_ParseTuple(args, " i", &i)) 22 return NULL; 23 self->count += i; 24 Py_RETURN_NONE; 25 } 26 27 static int 28 CounterInit(Counter0 *self, PyObject *args, 29 PyObject *kwargs) 30 { 31 self->count = 0; 32 return 0; 33 } 34 35 static PyMethodDef CounterMethods[] = { 36 {"incr", (PyCFunction) CounterIncr, 37 METH_VARARGS, ""}, {NULL} 38 }; </pre>	

Fig. 1. Example of a Python client program alongside a C counter module

<pre> 1 typedef struct PyObject { 2 Py_ssize_t ob_refcnt; 3 struct PyTypeObject *ob_type; 4 } PyObject; 5 6 typedef PyObject *(*PyCFunction) 7 (PyObject *, PyObject *); 8 typedef int (*initproc) 9 (PyObject *, PyObject *, PyObject *); 10 11 typedef struct PyMethodDef { 12 const char *ml_name; PyCFunction ml_meth; 13 int ml_flags; const char *ml_doc; 14 } PyMethodDef; 15 16 typedef struct PyMemberDef { 17 const char *name; int type; </pre>	<pre> 18 Py_ssize_t offset; int flags; 19 const char *doc; 20 } PyMemberDef; 21 22 typedef struct PyTypeObject { 23 PyObject ob_base; 24 const char *tp_name; 25 Py_ssize_t tp_basicsize; 26 Py_ssize_t tp_itemsize; 27 unsigned long tp_flags; 28 struct PyMethodDef *tp_methods; 29 struct PyMemberDef *tp_members; 30 struct PyTypeObject *tp_base; 31 PyObject *tp_dict; 32 initproc tp_init; 33 newfunc tp_new; 34 } PyTypeObject; </pre>
--	--

Fig. 2. Extract of Python's API header files

which it belongs. `PyObject` is Python’s `type` object, from which all classes derive. The second field of `Counter0` is the instance’s data: an integer `count`, not directly exposed to Python.

The `Counter` class’ specification is defined lines 52-59. It has three methods stored in the `Py_tp_new`, `Py_tp_init`, and `Py_tp_methods` fields. It also defines a special attribute member `counter` lines 41-42. The `PyObject` structure is synthesized from the specification by `PyType_FromSpec` (line 72). These methods and members are lifted to become Python attributes and methods when the class is initialized (in `PyType_FromSpec`). Other fields are defined in the `PyObject` structure. `tp_basicsize`, `tp_itemsize` define the size (in bytes) of instances. `tp_flags` is used to perform fast instance checks for builtin classes. `tp_base` points to the parent of the class. `tp_dict` is the class’ dictionary used by Python to resolve attribute accesses (created during class initialization).

Module import. When executing the `import counter` statement, the C function `PyInit_counter` is called. This function starts by creating a module whose name (line 62) is `counter` with no methods attached to it (line 63). Then, the `CounterT` class is created (lines 71-72), and the class is bound to the module (lines 73-74). Python uses a reference-counting-based garbage collector, which has to be manually handled using the `Py_INCREF` and `Py_DECREF` macros in C. If no errors have been encountered, the module object is returned at the end of the function, and `counter` will now point to this module in Python.

Class initialization. The call to `PyType_FromSpec` creates the `Counter` class from its specification. The function `PyType_FromSpec` starts by creating the `PyObject` and fills its fields according to the specification. This structure is then passed to `PyType_Ready` which populates the `tp_dict` field of the class. This field is the class’ dictionary used by Python to resolve attribute accesses. Before this call, the attribute `counter` and the methods `__new__`, `__init__`, and `incr` do not exist on the Python side. We explain how these C functions are encapsulated into Python objects by `PyType_Ready`. The prototype of C functions for Python is `PyCFunction` (Figure 2, line 6). Some signature adaptations may be needed for specific kinds of functions. For example, initialization methods (such as `CounterInit`) return a C `int` by convention. Thus, `CounterInit` will be wrapped into a function called `wrap_init`, which behaves as a `PyCFunction`. It is then encapsulated into a builtin Python descriptor object. Upon a call to this object, this descriptor performs pre- and post-call checks (described in Section 3). Continuing our example, `wrap_init` will be stored into an instance of the builtin `wrapper_descriptor` object. These descriptors are then added to the class dictionary. Table 1 describes the three other fields.

Counter creation. When a new instance of `Counter` is created (line 4), Python starts by allocating it by calling `Counter.__new__`. This call will eventually be resolved into `PyType_GenericNew` (from `tp_new`), allocating the object and

Table 1. Python `Counter` structure summary

Attribute	Encapsulating Object	Underlying Wrapper	Underlying C definition
<code>__new__</code>	<code>builtin_function</code>	<code>tp_new_wrapper</code>	<code>PyType_GenericNew</code>
<code>__init__</code>	<code>wrapper_descriptor</code>	<code>wrap_init</code>	<code>CounterInit</code>
<code>incr</code>	<code>method_descriptor</code>	<code>∅</code>	<code>CounterIncr</code>
<code>counter</code>	<code>member_descriptor</code>	<code>∅</code>	<code>CounterMembers[0]</code>

initializing the necessary fields (such as `ob_refcnt` and `ob_type` of `PyObject`). Then, `Counter.__init__` is called and the C function `CounterInit` ends up being called. It initializes the `count` field of the `Counter0 struct` to 0.

Counter increment. When the `incr` function is called, it is resolved through Python’s attribute accesses into `CounterIncr`. `CounterIncr` uses the standard Python function prototype, corresponding to `PyCFunction` (Figure 2). Its first argument is the object instance (here, the instance stored in variable `c`), and the second argument is a tuple of the arguments passed to the function (for the call at line 6 it is a tuple of length one containing `2**p-1`, and an empty one for the second call at line 7). `PyArg_ParseTuple` is a helper C function from the Python API converting the Python arguments wrapped in the tuple into C values.⁴ It uses a format string to describe the conversion. The `|` character separates mandatory arguments from the optional ones, while `i` signals a conversion from a Python integer to a C `int`. Internally, the conversion is done from a Python integer to a `long` (which may fail with an exception since Python integers are unbounded), which is then cast to an `int` if size permits (otherwise, another exception is set). In the first call to `CounterIncr`, `i` will be assigned `2**p-1` if the conversion is successful. In the second call, `i` will keep its default value, `1`. The internal value of the counter is then incremented by `i`, and then Python’s `None` value is returned.

Counter access. Thanks to the complex semantics of Python, attribute accesses can actually hide calls to custom getter and setter functions through the descriptor protocol [33, Figure 6]. In our case, `PyType_Ready` takes the member declaration lines 40-43, and creates those custom getters and setters through a `member_descriptor` builtin object. The access to attribute `counter` at line 8 calls the getter of this `member_descriptor` object. This getter accesses the `count` field of the `Counter0 struct` and converts the C integer into a Python integer. The `READONLY` flag in the declaration ensures that any call to the setter function raises an exception. These member descriptors are supported by our analysis.

⁴`Py_BuildValue` is the converse function translating C values to Python ones.

What can go wrong? Depending on the chosen value of p the result r will range from (i) the expected value ($r = 2^p$ when $0 \leq p \leq 30$), (ii) conversion errors raised as `OverflowError` exceptions (with messages depending on whether $p \leq 62$), (iii) a silent integer overflow on the C side, causing a wrap-around which is an unexpected behavior for Python developers ($r = -2^{31}$ for $p = 31$). All these errors are due to different representations between the builtin values of the language. The C integer overflow does not interrupt the execution. This motivates the creation of our analysis targeting all kinds of runtime errors in both host and guest languages as well as at the boundary. Our analysis infers all reachable C and Python values, as well as raised exceptions in order to detect these runtime errors. In this example, our analyzer is able to detect all these cases. Our analyzer is also able to prove that the program is safe when p ranges between 0 and 30.

Common bugs at the boundary. We refer the reader to the work of Hu and Zhang [16] for an empirical evaluation of bugs at the boundary between Python and C. The most frequent bugs happening at the boundary are:

- mismatches between a returned NULL and the exception being set in C (NULL should be returned by Python C functions if and only if an exception has been set – cf. Figure 8 in the next section),
- mismatches between the C and Python datatypes during conversion (in calls to `PyArg_ParseTuple`, `PyLong_FromLong`),
- integer overflows during conversions from arbitrary-precision Python integers to C,
- reference-counting errors (not supported by our analyzer).

3 Concrete Semantics

This section defines the semantics of the interface between Python and C. It is built upon the semantics of each language. Our goal is to delegate most of the work to the semantics of each language, each used in a black-box fashion. This delegation will also simplify our implementation as we will reuse the analyses of Python and C in a similar black-box fashion.

A key assumption of our semantics is that builtin Python objects (such as integers and dictionaries) can only be accessed from C through the API provided by the interpreter. As such, any access to the builtins through the C API can be encoded as a call back to the Python language. Thus, each language will have complementary representations of Python objects. Each language has a view of any Python object. Accesses to the other view are done by switching language. To illustrate this representation on our running example, the `int count` field from a `Counter` instance is only exposed from the C. It is possible to read the counter’s value from Python. This can only be done by calling a C function performing the conversion of the C integer into a Python integer. This object is new and independent from the C integer. Hence, only C code directly dereferences the field value from the object memory. Conversely, an attribute that

is dynamically added to a Python object is stored in the instance’s dictionary. This dictionary is opaque from C. Accessing it through the C language using the API will ultimately be evaluated as a Python attribute access in our semantics. As illustrated in these examples, mutable parts of the objects are not *directly* available through both languages. There is thus no need to perform systematic state reductions when switching from one language to the other.

Our concrete semantics defines the operators working at the boundary. These operators allow switching from one language to another or converting the value of one language to the other. We define how Python may call C functions, and how C may perform callbacks to Python objects. We also define conversions between Python integers and C **longs**. API functions working on other builtin datatypes (such as floats, strings, lists, ...) exist and are supported by our analysis. They are similar to the integer case and are not described in this article. The definitions of these operators rely on calls to boundary functions between the two languages. These boundaries ensure that objects have the correct shape in the other language state. The boundary from C to Python also checks that only valid Python objects are passed back.

We define the state on which each semantics operates. In the following, Python-related states and expressions will be written in **green**. C-related states and expressions will be in **orange**. We reuse the states defined in the work of Monat et al. [33] for Python and Ouadjaout and Miné [37] for C. A set of heap addresses **Addr** (potentially infinite) is common to the states. Previous works [11, 33, 37] define the semantics of Python and C.

Python state (Figure 3). Python objects are split into a nominal part and a structural part. The nominal part **ObjN** can be a builtin object such as an integer (we omit other builtins for the sake of concision), a function, a class, or an instance (defined by the address of the class from which it is instantiated). The structural part **ObjS** maps attribute names to their contents’ addresses. A state Σ_p consists of an environment and a heap. The environment \mathcal{E}_p maps variable identifiers **Id_p** to addresses (or **LocalUndef** for local variables with an undefined value). The heap \mathcal{H}_p maps addresses to objects. Given a state $\sigma_p \in \Sigma_p$, we write as $\sigma.\epsilon_p$ its environment and $\sigma.\eta_p$ its heap. Following Fromherz et al. [11], the state is additionally tagged by a flow token to handle non-local control-flow: *cur* represents the current flow on which all instructions that do not disrupt the control flow operate (e.g., assignments, but not **raise**); *exn* collects the states given by a raised exception. *exn* is indexed by the address of the exception object, so each exception will be kept separate. These continuation-tagged states are standard in abstract interpreters working by induction on the syntax to model non-local control-flow. Evaluating an expression **e** through $\mathbb{E}_p[e]$ in a given state yields a Python value in a new state. This value may be \perp if the evaluation fails. We use **letb** $v, \sigma = f(e)$ **in** *body* as syntactic sugar for **let** $r, \sigma = f(e)$ **in** *if* $r \neq \perp$ **then** *body* **else** \perp, σ .

$$\begin{aligned}
\mathbf{ObjN} &\stackrel{\text{def}}{=} \mathbf{int}(i \in \mathbb{Z}) \cup \mathbf{Fun}(f) \\
&\cup \mathbf{Class}(c) \cup \mathbf{Inst}(a \in \mathbf{Addr}) \\
\mathbf{ObjS} &\stackrel{\text{def}}{=} \mathit{string} \rightarrow \mathbf{Addr} \\
\mathbf{Value}_p &\stackrel{\text{def}}{=} \mathbf{Addr} \\
\mathcal{F}_p &\stackrel{\text{def}}{=} \{ \mathit{cur}, \mathit{exn} \mid a \in \mathbf{Addr} \} \\
\mathcal{E}_p &\stackrel{\text{def}}{=} \mathbf{Id}_p \rightarrow \mathbf{Addr} \cup \mathbf{LocalUndef} \\
\mathcal{H}_p &\stackrel{\text{def}}{=} \mathbf{Addr} \rightarrow \mathbf{ObjN} \times \mathbf{ObjS} \\
\Sigma_p &\stackrel{\text{def}}{=} \mathcal{F}_p \times \mathcal{E}_p \times \mathcal{H}_p \\
\mathbb{E}_p[\mathit{expr}] &: \Sigma_p \rightarrow \mathbf{Value}_p^\perp \times \Sigma_p \\
\mathbb{S}_p[\mathit{stmt}] &: \Sigma_p \rightarrow \Sigma_p
\end{aligned}$$

Fig. 3. Concrete Python State

$$\begin{aligned}
\mathbf{Cells} &\stackrel{\text{def}}{=} \{ \{b, o, t\} \mid b \in \mathbf{Base}, t: \text{scalar} \\
&\quad \text{type}, 0 \leq o \leq \text{sizeof}(b) - \text{sizeof}(t) \} \\
\mathbf{Ptr} &\stackrel{\text{def}}{=} (\mathbf{Base} \times \mathbb{Z}) \cup \{ \mathbf{NULL}, \mathbf{invalid} \} \\
\mathbf{Base} &\stackrel{\text{def}}{=} \mathbf{Id}_c \cup \mathbf{Addr} \\
\mathbf{Value}_c &\stackrel{\text{def}}{=} \mathbf{MNum} \cup \mathbf{Ptr} \\
\mathcal{E}_c &\stackrel{\text{def}}{=} \mathbf{Cells} \rightarrow \mathbf{Value}_c \\
\mathcal{H}_c &\stackrel{\text{def}}{=} \mathbf{Addr} \rightarrow \mathit{ident} \times \mathbb{N} \\
\Sigma_c &\stackrel{\text{def}}{=} \mathcal{E}_c \times \mathcal{H}_c \\
\mathbb{E}_c[\mathit{expr}] &: \Sigma_c \rightarrow \mathbf{Value}_c^\perp \times \Sigma_c \\
\mathbb{S}_c[\mathit{stmt}] &: \Sigma_c \rightarrow \Sigma_c
\end{aligned}$$

Fig. 4. Concrete C State

$$\begin{aligned}
\Sigma &= \Sigma_p \times \Sigma_c \\
p \hookrightarrow_c &: \mathbf{Value}_p \times \Sigma \rightarrow \mathbf{Value}_c \times \Sigma \\
c \hookrightarrow_p &: \mathbf{Value}_c \times \Sigma \rightarrow \mathbf{Value}_p^\perp \times \Sigma \\
\mathbb{E}_{p \times c}[\mathit{expr}_p] &: \Sigma \rightarrow \mathbf{Value}_p^\perp \times \Sigma \\
\mathbb{E}_{p \times c}[\mathit{expr}_c] &: \Sigma \rightarrow \mathbf{Value}_c^\perp \times \Sigma
\end{aligned}$$

Fig. 5. Combined State

C state (Figure 4). The memory is decomposed into blocks **Base** which are either variables **Id_c** or heap addresses **Addr**. Each block is decomposed into scalar elements (machine integers, floats, pointers). $\{b, o, \tau\}$ denotes the memory region in block b , starting at offset o and having type τ . C values **Value_c** are either machine numbers **MNum**, or pointers **Ptr** defined by their block and offset. Additionally, pointers can be **NULL** or **invalid**. The state Σ_c consists of an environment and a heap. The environment \mathcal{E}_c maps scalar elements to values. The heap \mathcal{H}_c maps addresses to the type of allocated resource and their size. The type of allocated resource is **Malloc** when the standard C library **malloc** is used⁵. The Python allocator (called by **PyType_GenericNew**) will create resources of type **PyAlloc**, ensuring that: (i) Python objects are well constructed by the correct allocator (ii) the C code cannot access these “opaque” objects and needs to use the API.

⁵Other resources (such as file descriptors) can also be defined [37].

Combined state (Figure 5). Two new kinds of nominal objects are added to Python: **CFun** f for Python functions defined in C, **CClass** c for Python classes defined in C (where f and c denote the name of the underlying C function or class declaration). The combined state used for the multilanguage semantics is the product of the Python and C states, written Σ . Note that each state may reference addresses originally allocated by the other language (in the running example, the Python variable c points to the address of the **Counter** instance, which has been allocated on the C side by **PyType_GenericNew**). In the following, we define two boundary functions converting a Python value into a C value and conversely (written $p \hookrightarrow_c$ and $c \hookrightarrow_p$). The multilanguage semantics $\mathbb{E}_{p \times c}[\cdot]$ is defined over Python and C expressions. It operates over the whole state Σ and its return value matches the language of the input expression. We define the semantics of some builtins working at the boundary between Python and C, which require the whole state. For expressions not working at the boundary, the multilanguage semantics defaults to the usual Python or C semantics:

$$\begin{aligned}\mathbb{E}_{p \times c}[\text{expr}_p](\sigma_p, \sigma_c) &= \mathbb{E}_p[\text{expr}_p](\sigma_p), \sigma_c \\ \mathbb{E}_{p \times c}[\text{expr}_c](\sigma_p, \sigma_c) &= \sigma_p, \mathbb{E}_c[\text{expr}_c](\sigma_c)\end{aligned}$$

Boundary functions. Boundary functions ensure that Python objects are correctly represented in the heap of each language. The C to Python boundary also ensures that only Python objects are passed back to Python. These functions do not convert values from one language to another. This kind of conversion is handled by builtin conversion functions such as **PyLong_AsLong**, **PyLong_FromLong** for Python integer to C long conversion. These boundary functions are lazy and shallow: (i) only objects switching languages are passed through those boundaries, (ii) an object that has already been converted does not need to be converted again (i.e., when its address is already in the other language’s heap).

The boundary from Python to C is described in Figure 6. The boundary is first applied recursively to the class of the object (using the **type** operator of Python). Then, the **ob_type** field of the object is initialized to point to its class. The last step performed is to update the heap: the object has been allocated by Python, and has the size of **PyObject** (if the object is a class, it has the size of **PyTypeObject**, and we call the class initializer afterward).

The converse boundary (Figure 7) starts by checking that the value is a heap allocated Python object, allocated with resource type **PyAlloc**. It calls itself recursively on the class of the object (using the **ob_type** field in C). The Python heap is updated with the converted object.

C call from Python (Figure 8). The semantics of C function calls from Python is shown in Figure 8. The function **in_check** enforces that e_1 should be an instance of the class to which f is bound. Otherwise, a **TypeError** exception is raised. C functions callable from Python can only have two arguments (cf. the type of **PyCFunction**, Figure 2, line 6). Thus, the Python arguments are split

into the first one e_1 and the other ones, bundled in a tuple. The boundary from Python to C is applied to e_1 , and to the tuple containing the other arguments. Then, the C function is evaluated using the standard C semantics. Afterward, `out_check` ensures that the function returned NULL if and only if an exception has been set in the interpreter state. Otherwise, a `SystemError` exception is raised. Finally, the C value is passed through the boundary function.

Python call from C (Figure 9). Calls back to Python from the C code are possible using the `PyObject_CallObject` function. The first argument is the object being called. The second argument is a tuple containing all the parameters. These two arguments are first passed through the C to Python boundary. Then, we use the Python semantics to evaluate the call (the `*` operator in Python unpacks the tuple into the arguments of the variadic function). If the call is successful (i.e., the execution is normal, shown by flow token *cur*), the converse boundary function is applied. If an exception has been raised during the evalu-

$$\begin{aligned}
&{}^p\hookrightarrow_c(v_p, \sigma_p, \sigma_c) = \\
&\text{if } v_p \in \sigma.\eta_c \text{ then } (v_p, 0), \sigma_p, \sigma_c \text{ else} \\
&\text{letb } ty_p, \sigma_p = \mathbb{E}_p[\![type(v_p)]\!] \sigma_p \text{ in} \\
&\text{letb } (ty_c, 0), \sigma_p, \sigma_c = {}^p\hookrightarrow_c(ty_p, \sigma_p, \sigma_c) \text{ in} \\
&\text{let } \sigma_c = \mathbb{S}_c[\![v_p \rightarrow \text{ob_type} = ty_c]\!] \sigma_c \text{ in} \\
&\text{let } \sigma_p, \sigma_c = \\
&\quad \text{if } \sigma_p(v_p) = \text{Class}(c) \text{ then} \\
&\quad \quad \text{let } \sigma_c = \sigma.\epsilon_c, \sigma.\eta_c[v_p \mapsto \text{PyAlloc, sizeof(PyTypeObject)}] \text{ in} \\
&\quad \quad \mathbb{E}_{p \times c}[\![\text{PyType_Ready}(v_p)]\!] (\sigma_p, \sigma_c) \\
&\quad \text{else } \sigma_p, (\sigma.\epsilon_c, \sigma.\eta_c[v_p \mapsto \text{PyAlloc, sizeof(PyObject)}]) \\
&\text{in } (v_p, 0), \sigma_p, \sigma_c
\end{aligned}$$

Fig. 6. Python to C value boundary

$$\begin{aligned}
&{}^c\hookrightarrow_p(v_c, \sigma_p, \sigma_c) = \\
&\text{if } v_c \notin \text{Addr} \times \{0\} \parallel \sigma.\eta_c(\text{fst } v_c) \neq (\text{PyAlloc}, _) \text{ then } \perp, \sigma_p, \sigma_c \text{ else} \\
&\text{let } v = \text{fst } v_c \text{ in} \\
&\text{if } v \in \sigma.\eta_p \text{ then } v, \sigma_p, \sigma_c \text{ else} \\
&\text{letb } ty_c, \sigma_c = \mathbb{E}_c[\![(\text{PyObject}*)v) \rightarrow \text{ob_type}]\!] \sigma_c \text{ in} \\
&\text{letb } ty_p, \sigma_p, \sigma_c = {}^c\hookrightarrow_p(ty_c, \sigma_p, \sigma_c) \text{ in} \\
&\text{let } \sigma_p = \sigma.\epsilon_p, \sigma.\eta_p[v \mapsto \text{Inst}(ty_p), \emptyset] \text{ in } v, \sigma_p, \sigma_c
\end{aligned}$$

Fig. 7. C to Python value boundary

$$\mathbb{E}_{p \times c} \llbracket (\mathbf{CFun} \ f)(e_1, e_2, \dots, e_n) \rrbracket (\sigma_p, \sigma_c) =$$

```

letb  $\sigma_p = \text{in\_check}(\mathbf{CFun} \ f, e_1, \sigma_p)$  in
letb  $c_1, \sigma_p, \sigma_c = {}^p\hookrightarrow_c(e_1, \sigma_p, \sigma_c)$  in
letb  $p_2, \sigma_p = \mathbb{E}_p \llbracket \text{tuple}(e_2, \dots, e_n) \rrbracket \sigma_p$  in
letb  $c_2, \sigma_p, \sigma_c = {}^p\hookrightarrow_c(p_2, \sigma_p, \sigma_c)$  in
letb  $c_f, \sigma_c = \mathbb{E}_c \llbracket f(c_1, c_2) \rrbracket \sigma_c$  in
letb  $c_f, \sigma_c = \text{out\_check}(c_f, \sigma_c)$  in
 ${}^c\hookrightarrow_p(c_f, \sigma_p, \sigma_c)$ 

```

Fig. 8. C call from Python

$$\mathbb{E}_{p \times c} \llbracket \text{PyObject_CallObject}(f, a) \rrbracket (\sigma_p, \sigma_c) =$$

```

letb  $f_p, \sigma_p, \sigma_c = {}^c\hookrightarrow_p(f, \sigma_p, \sigma_c)$  in
letb  $a_p, \sigma_p, \sigma_c = {}^c\hookrightarrow_p(a, \sigma_p, \sigma_c)$  in
let  $r_p, \sigma_p = \mathbb{E}_p \llbracket f_p(*a_p) \rrbracket \sigma_p$  in
if  $\sigma_p = (cur, \_, \_)$  then  ${}^p\hookrightarrow_c(r_p, \sigma_p, \sigma_c)$ 
else let  $exn \ e, \epsilon_p, \eta_p = \sigma_p$  in
letb  $e_c, \sigma_p, \sigma_c = {}^p\hookrightarrow_c(e, (cur, \epsilon_p, \eta_p), \sigma_c)$  in
 $\text{NULL}, \sigma_p, \mathbb{S}_c \llbracket \text{PyErr\_SetNone}(e_c) \rrbracket \sigma_c$ 

```

Fig. 9. Python call from C

ation of the Python call, we revert to the *cur* flow token and pass the exception object *e* through the boundary. The result of the call will be `NULL`, and the exception will be set on the C side by calling `PyErr_SetNone`.

Python exceptions in C. Python exceptions may be raised from the C code using the `PyErr_SetNone` builtin. In the Python interpreter, this sets a flag in a structure describing the interpreter’s state. We model this by setting a global variable `exc` with the Python object passed as an argument. Additional functions such as `PyErr_Occurred` checking if an exception has been raised and `PyErr_Clear` removing raised exceptions are modeled by accessing and setting this same global variable.

Conversion builtins of the API. We show conversion functions from C `long` to Python integers and back in Figure 10. Converting a C `long` of value v_c to a Python integer is done by calling the integer constructor in the Python semantics, and applying the boundary afterwards. To perform the other conversion, we apply the boundary function to the C value. Then, we check if the corresponding Python value v_p is an integer by looking into the Python heap. If that is the case, we check that this integer fits in a C `long` (Python integers are unbounded). Otherwise we raise an `OverflowError` and the function returns -1. A `TypeError` exception is raised and the function returns -1 if the object is not an integer.

Thanks to the definition of builtins such as `PyLong_AsLong`, other builtins calling this function can be analyzed directly using their source code from the Python interpreter’s implementation. For example, when `PyArg_ParseTuple` encounters an `'i'` `char` in its conversion string, it executes the code shown in Figure 11.⁶ As explained in our example from Section 2, it first calls `PyLong_AsLong` and converts the `long` to `int` checking for additional overflows. Our analyzer is able to analyze this piece of code directly. In our implementation, about half of

⁶`PyArg_ParseTuple` is defined as stub (just as `PyLong_AsLong` is), but the case of integers is delegated to the interpreter’s implementation shown in Figure 11.

$$\begin{aligned}
\mathbb{E}_{p \times c} \llbracket \text{PyLong_FromLong}(v_c) \rrbracket (\sigma_p, \sigma_c) &= \overset{p}{\hookrightarrow}_c (\mathbb{E}_p \llbracket \text{int}(v_c) \rrbracket (\sigma_p), \sigma_c) \\
\mathbb{E}_{p \times c} \llbracket \text{PyLong_AsLong}(v_c) \rrbracket (\sigma_p, \sigma_c) &= \\
&\text{let } v_p, \sigma_p, \sigma_c = \overset{c}{\hookrightarrow}_p (v_c, \sigma_p, \sigma_c) \text{ in} \\
&\text{if } \sigma.\eta_p = \text{int}(i) \text{ then} \\
&\quad \text{if } i \in [-2^{63}, 2^{63} - 1] \text{ then } i, \sigma_p, \sigma_c \\
&\quad \text{else } -1, \sigma_p, \mathbb{S}_c \llbracket \text{PyErr_SetNone(PyExc_OverflowError)} \rrbracket \sigma_c \\
&\text{else } -1, \sigma_p, \mathbb{S}_c \llbracket \text{PyErr_SetNone(PyExc_TypeError)} \rrbracket \sigma_c
\end{aligned}$$

Fig. 10. Conversion from Python builtin integers to C long

<pre> 1 long ival = PyLong_AsLong(obj); 2 if(ival == -1 && PyErr_Occurred()) { 3 return 0; 4 } else if (ival > INT_MAX) { 5 PyErr_SetString(PyExc_OverflowError, 6 "signed integer is greater than maximum"); 7 return 0; 8 } </pre>	<pre> 10 else if (ival < INT_MIN) { 11 PyErr_SetString(PyExc_OverflowError, 12 "signed integer is less than minimum"); 13 return 0; 14 } else { 15 *result = ival; 16 return 1; 17 } </pre>
---	--

Fig. 11. Python integer to C int conversion as done by PyArg_ParseTuple

the builtins use the original C implementation; their code is around 650 lines long.

Threats to validity. Our goal is to analyze Python programs with native C modules and detect all runtime errors that may happen. Assuming that those C modules use Python’s API rather than directly modify the internal representation of builtins seems reasonable when analyzing third-party modules. This is the recommended approach for developers, as it eases maintenance of the codebase since API changes are not frequent and documented. Our analysis is able to detect if a program does not use the Python API and tries to modify a builtin Python object directly.

This concrete semantics is already high-level. A lower-level semantics where implementation details of builtins are exposed would be much more complex. It would not benefit our analysis, which aims at detecting runtime errors in the user’s codebase. We have established this semantics by reading the code of the reference Python interpreter. Proving that our semantics is a sound abstraction of such lower-level semantics is left as future work.

4 Abstract Semantics

Concrete states use numeric values in different places (e.g, the C state has machine numbers, pointer offsets and resource sizes). All these values will be

```

rel_count.py
1 c = counter.Counter()
2 for i in range(randint(1, 100)):
3     c.incr()
4 r = c.counter
5 assert(r == i+1)

```

Fig. 12. Example code where relationality between C and Python improves precision

centralized in a common numeric domain in the abstract state. This centralization allows expressing invariants between all those numeric variables, possibly improving the precision. We show a generic construction of the abstract multilanguage state. We assume the abstract semantics of Python and C are provided through $\mathbb{E}_p^\#[\cdot]$, $\mathbb{E}_c^\#[\cdot]$. These can be instantiated in practice using previous works [37, 34]. We assume that each language's abstract state relies on an address allocation abstraction (such as the callsite abstraction or the recency abstraction [2]) and a numeric abstraction (such as intervals, octagons, ...). We write $\Sigma_u^\#$ the cartesian product of these two abstractions. The abstract Python (resp. C) state can then be decomposed as a product $\Sigma_p^\# = \tilde{\Sigma}_p^\# \times \Sigma_u^\#$ (resp. $\Sigma_c^\# = \tilde{\Sigma}_c^\# \times \Sigma_u^\#$). The multilanguage abstract state consists in the cartesian product of the Python and C abstract states, where the address allocation and numeric states are shared: $\Sigma_{p \times c}^\# = \tilde{\Sigma}_p^\# \times \tilde{\Sigma}_c^\# \times \Sigma_u^\#$.

Just as the concrete semantics builds upon the underlying C and Python semantics, so does our abstract semantics. The abstract semantics of the boundary operators is structurally similar to the concrete ones (each concrete operator is replaced with its abstract counterpart). We show this transformation on the abstract semantics of `PyLong_FromLong` (to be compared with Figure 10).

$$\mathbb{E}_{p \times c}^\# [\text{PyLong_FromLong}(v_c)](\sigma_p^\#, \sigma_c^\#) = \overset{p \hookrightarrow c}{\hookrightarrow} (\mathbb{E}_p^\# [\text{int}(v_c)](\sigma_p^\#), \sigma_c^\#)$$

Sharing the address allocation and numeric abstractions allows expressing relational invariants between the languages. In the example in Figure 12, a non-relational analysis would be able to infer that $0 \leq i \leq 99$, but it cannot infer that the number of calls to `incr` is finite. It would thus infer that $-2^{31} \leq r < 2^{31}$, report an overflow error and be unable to prove the assertion at the end. Note that the value of `r` originates from the C value of the `count` field in the instance defined in `c`. With a relational analysis where C and Python variables are shared in the numeric domains, it is possible to infer that $\text{num}(@_{int}) + 1 = \text{num}(@_{Counter}, 16, s32)$. $\text{num}(@_{int})$ is the numeric value of the integer bound to `i`. $\text{num}(@_{Counter}, 16, s32)$ is the numeric value of the `Counter` instance (i.e., the value of `count` in the `Counter struct`, here represented as the cell [31] referenced by the `Counter` instance, at offset 16 being a 32-bit integer). Our analyzer is able to prove that the assertion holds using the octagon abstract domain [32].

Soundness. Assuming the underlying abstract semantics of Python and C are sound, the only cases left in the soundness proof are those of the operators

working at the boundary. Since the abstract semantics of those operators is in point-to-point correspondence with the concrete semantics, the soundness proof is straightforward.

5 Experimental Evaluation

Implementation. We have implemented our multilanguage analysis within the open-source, publicly-available static analysis framework called Mopsa [18, 19]. A specific feature of Mopsa is to provide loosely-coupled abstract domains that the user combines through a configuration file to define the analysis. We were able to reuse off-the-shelf value analyses of C programs [37] and Python programs [34] already implemented into Mopsa. The only modification needed was to add a multilanguage domain, implementing the semantics of the operators at the boundary, just as our semantics (both concrete and abstract) do.

The configuration for the multilanguage analysis is shown in Figure 13. A configuration is a directed acyclic graph, where nodes are either domains or domain combinators. Domains can have their own local abstract state. The global abstract state is the product of these local states. Domains can also be iterators over language constructions. Each analyzed expression (or statement) flows from the top domain of the configuration to the bottom until one domain handles it. The multilanguage domain is at the top. It dispatches statements not operating at the boundary to the underlying Python or C analysis. The Python and C analyses are in a cartesian product, ensuring that when a statement goes through them, it will be handled by only one of the two sub-configurations. An example of stateful domain is “C.pointers”, which keeps points-to information. “Py.exceptions” is an example of iterator domain. It handles the `raise` and `try` operators of Python. Both Python and C analyses share an underlying “universal” analysis, to which they can delegate some statements. This “universal” part provides iterators for intra- and inter-procedural statements, as well as the address allocation and numeric abstractions. The numeric abstraction displayed here only uses intervals, but it can be changed to a reduced product between a relational domain and intervals, for example.

This multilanguage domain consists in 2,500 lines of OCaml code (measured using the `cloc` tool), implementing 64 builtin functions such as the ones presented in the concrete semantics. This is small compared to the 11,700 lines of OCaml for the C analysis and 12,600 lines of OCaml for the Python analysis. These domains rely on “universal” domains representing 5,600 lines of OCaml and a common framework of 13,200 lines of OCaml. We also reuse the C implementation of 60 CPython functions as-is.

Corpus selection. In order to perform our experimental evaluation, we selected six popular Python libraries from GitHub (having in average 412 stars). These libraries are written in C and Python and do not have external dependencies. The `noise` library [10] aims at generating Perlin noise. Libraries

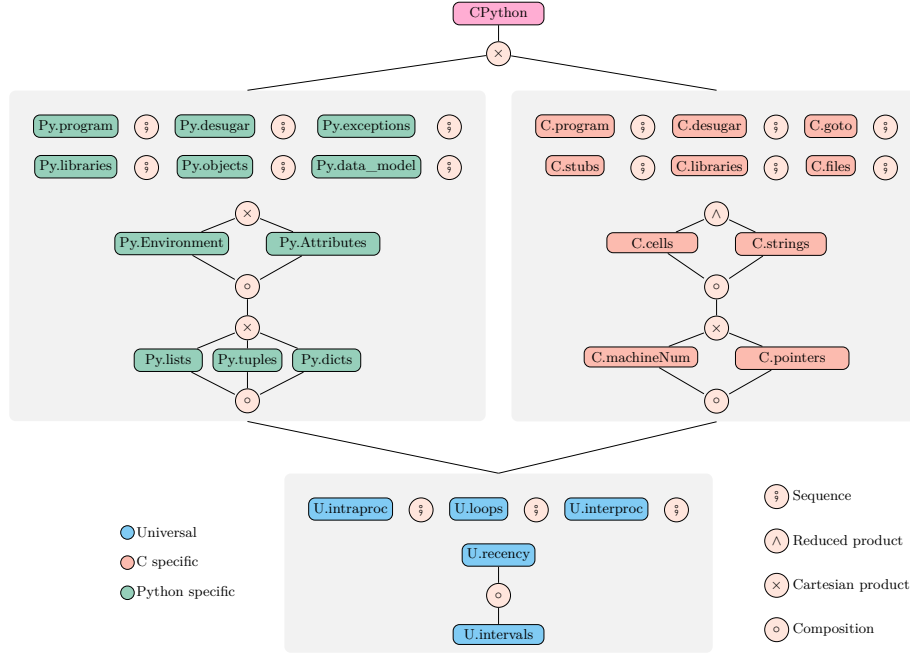


Fig. 13. Multilanguage configuration in Mopsa

`levenshtein`, `ahocorasick`, `cdistance` [15, 36, 30] implement various string-related algorithms. `l1ist` [17] defines linked-list objects (both single and double ones). `bitarray` [41] provides an implementation for efficient arrays of bits. Our analysis is context-sensitive in order to perform a precise value analysis. Thus, we needed some client code to analyze those libraries. We decided to analyze the tests defined by those libraries: they should cover most usecases of the library, and ensure that the transition between Python and C are frequent, which is ideal to stress-test our analysis. Some libraries (`noise`, `bitarray`, and `l1ist`) come with example programs with less than 50 lines of code that we analyze within 15 seconds. We have not been able to find applications with a well-defined entry point using those libraries (or they had big dependencies such as `numpy`). Our experimental evaluation thus focuses on the analysis of the libraries' tests.

Analysis results. We show the results of our analysis in Table 2. The column $|C|$ (resp. $|Py|$) shows the lines of code in C (resp. Python), measured using the `cloc` tool. The Python code corresponds mainly to the tests. It may also wrap C classes in custom classes. For example, `frozenbitarray` is defined in Python, on top of the `bitarray` class. The “tests” column shows the number of tests we are able to analyze, compared to the total number of tests defined by the library. The 🕒 column shows the time taken to analyze all the tests. Columns 🟢_c (resp. 🟢_p) show the selectivity of our analysis – the number of safe operations

Table 2. Analysis results of libraries using their unit tests

Library	C	Py	Tests	🕒	🔴 _c	🟢 _p	Assert. Py↔C	
noise	722	675	¹⁵ / ₁₅	19s	99.6% (4952)	100.0% (1738)	⁰ / ₂₁	6.6
ahocorasick	3541	1336	⁴⁶ / ₉₂	59s	93.1% (1785)	98.0% (4937)	³⁰ / ₈₈	5.4
levenshtein	5441	357	¹⁷ / ₁₇	1.6m	79.9% (3106)	93.2% (1719)	⁰ / ₃₈	2.7
cdistance	1433	912	²⁸ / ₂₈	1.9m	95.3% (1832)	98.3% (11884)	⁸⁸ / ₂₀₇	8.6
l1ist	2829	1686	¹⁶⁷ / ₁₉₄	4.3m	99.0% (5311)	98.8% (30944)	²³⁵ / ₆₉₁	51.7
bitarray	3244	2597	¹⁵⁹ / ₂₁₆	4.6m	96.3% (4496)	94.6% (21070)	⁹⁷ / ₃₇₄	14.9

compared to the total number of runtime error checks, the latter being also displayed in parentheses – performed by the analyzer for C (resp. Python). The selectivity is computed by Mopsa during the analysis. The C analysis checks for runtime errors including integer overflows, divisions by zero, invalid memory accesses and invalid pointer operations. The Python analysis checks also for runtime errors, which include the `AttributeError`, `TypeError`, `ValueError` exceptions. Runtime errors happening at the boundary are considered as Python errors since they will be raised as Python `SystemError` exceptions. The second to last column shows the number of functional properties (expressed as assertions) defined by the tests that our analyzer is able to prove correct automatically. The last column shows the number of transitions between the analyzed Python code and the C code, averaged per test.

We observe that Mopsa is able to analyze these libraries in a few minutes with high selectivity for the detection of Python and C runtime errors. Our analysis is able to detect some bugs that were previously known. For example, the **ahocorasick** module forgets to initialize some of its iterator classes, and some functions of the **bitarray** module do not set an exception when they return an erroneous flag, raising a `SystemError` exception. We have not manually checked if unknown bugs were detected by our analysis. We have instrumented Mopsa to display the number of crossings (from Python to C, or C to Python). The average number of crossings per test is shown in the last column of Table 2. The minimal number of crossings is one per test. Thus these tests seem correct to benchmark our approach since they all alternate calls to native C code and Python code.

The multilanguage analysis is limited by the current precision level of the underlying C and Python analyses but would naturally benefit immediately from any improvements in these. However, we focused on the multilanguage domains only in this study. We leave as future work the improvements required independently on the C and Python analyses for our benchmarks. We now describe a few areas where the analysis could benefit from improvements. Mopsa is unable to support some tests for now, either because they use unsupported Python libraries or because the C analysis is too imprecise to resolve some pointers. The unsupported tests of the **ahocorasick** analysis are due to imprecisions in the C analysis, which is not able to handle a complex trie data structure being stored in a dynamic array and reallocated over and over again. In **l1ist**, some

tests use the `getrefcount` method of the `sys` which is unsupported (and related to CPython’s reference-based garbage collector, which we do not support). In addition, some tests make pure-Python classes inherit from C classes: this is currently not supported in our implementation, but it is an implementation detail that will be fixed. For the `bitarray` tests, tests are unsupported because they use the unsupported `pickle` module performing object serialization, or they use the unsupported `sys.getsizeof` method, or they perform some unsupported input-output operations in Python. In addition, the C analysis is too imprecise to resolve some pointers in 18 tests.

The selectivity is lower in the C analysis of `levenshtein`, where dynamic arrays of structures are accessed in loops: the first access at `tab[i].x` may raise an alarm and continue the analysis assuming that `i` is now a valid index access. However, subsequent accesses to `tab[i].y`, `tab[i].z` will also raise alarms as the non-relational numeric domain is unable to express that `i` is a valid index access. Proving the functional properties is more challenging, and not the main goal of our analysis, which aims detecting runtime errors. For example, the assertions of the `noise` library check that the result of complex, iterative non-linear arithmetic lies in the interval $[-1, 1]$. Some assertions in the `l1list` or `bitarray` library aim at checking that converting their custom container class to a list preserves the elements. Due to the smashing abstraction [3] of the Python lists, we cannot prove these assertions.

6 Related Work

Native code analysis. Some works focus on analyzing native C code in the context of language interoperability, without analyzing the host language. Tan and Croft [42] perform an empirical study of native code use in Java and provide a classification by bug patterns; a similar study has been performed by Hu and Zhang [16] for the Python/C API. Kondoh and Onodera [20] check that native calls to Java methods should handle raised exceptions using a typestate analysis. Li and Tan [23] ensure that the native control-flow is correctly interrupted when a Java exception is raised. The work of Li and Tan [24, 25] infers which Java exceptions may be raised by JNI code, allowing the exception type-safety property of Java programs to be extended to the JNI. CpyChecker [27] is a GCC plugin searching for common erroneous patterns in C code using the CPython API. Two works [26, 28] aim at detecting reference counting errors in C code using the CPython API. Brown et al. [4] define specialized analyses for specific patterns of C++ interoperability that may jeopardize type or memory safety of JavaScript. Contrary to these works, we analyze both host and guest languages.

Multilanguage semantics. The seminal work of Matthews and Findler [29] defines the first semantics of multilanguage systems, using the notion of boundaries to model conversion between languages. Buro and Mastroeni [7] generalize this construction using an algebraic framework. We use a similar notion of boundary in our concrete semantics.

Multilanguage analyses. Buro et al. [6] define a theory based on abstract interpretation to combine analyses of different languages, and show how to lift the soundness property to the multilanguage setting. They provide an example of multilanguage setting where they combine a toy imperative language with a bit-level arithmetic language. The notion of boundary functions used in their work performs a full translation from the state of one language to the other. Our semantics works on the product of the states, although it can be seen as an abstraction of the semantics of C and Python, where the boundary performs a full state conversion (but the boundary from Python to C would be a concretization). From an implementation standpoint, our approach avoids costly state conversions at the boundary and allows sharing some abstract domains.

Chipounov et al. [8] perform symbolic execution of binaries, thus avoiding language considerations. Their approach is extended by the work of Bucur et al. [5], which supports any interpreted language by performing symbolic execution over the interpreter. Our approach is more costly to implement since we do not automatically lift the interpreter’s code to obtain our analyzer. Thanks to its higher-level, we think our approach should be more precise and efficient.

The next works perform multilanguage analyses by translating specific effects of what native functions do (they usually generate a summary using a bottom-up analysis) to the host language. This allows removing the native code and use existing analyses of the host language. Tan and Morrisett [43] compile C code into an extended JVMML syntax form, allowing the use of the bug-finding tool Jlint afterwards. Furr and Foster [12, 13, 14] perform inference of OCaml and Java types in C FFI code, which they crosscheck with the types used in the client OCaml/Java code. They assume there are no out-of-bounds accesses and no type casting in the C code. An inter-language, bottom-up taint analysis for Java and native binary code in the setting of Android applications is proposed by Wei et al. [45]. Lee et al. [22] aim at detecting wrong foreign function calls and mishandling of Java exceptions in Java/JNI code. They extract summaries of the Java callbacks and field accesses from the JNI code using Infer, transform these summaries into Java code, and call the FlowDroid analyzer on the whole. Contrary to these works, our analyzer supports both languages, and it switches between languages just as the real code does. The properties we target require precise context-sensitive value analyses that are difficult to perform bottom-up.

Library analyses. Previous work aim at analyzing libraries with no access to their client code [1, 38] using a “most-general client”. The work of Kristensen and Møller [21] refines the notion of most-general client in the setting of dynamic programming languages. However, it focuses on libraries where functions are typed. Python libraries are not explicitly typed. Extending their work to our untyped, multilanguage setting is left as future work.

7 Conclusion

This article presents a multilanguage analysis able to infer runtime errors in Python code using native C extensions. Our analyzer is able to reuse value analyses of Python and C off-the-shelf. It shares the address allocation and numeric abstractions between the Python and C abstract domains. We are able to analyze within a few minutes real-world Python libraries written in C and having up to 5,800 lines of code. To the best of our knowledge, we have implemented the first static analyzer able to soundly detect runtime errors in multilanguage programs.

Future work includes extending our implementation to analyze the standard Python library and large Python applications. This will require handling more dependencies, having a relational analysis that scales, and addressing the precision limitations of the underlying C and Python analyses. We plan to instrument our implementation to verify (or infer) type annotations of the standard library provided in the `typeshed` [44] project. It would also be interesting to target multilanguage-specific safety properties (such as correct garbage collection counts). Another future work is to try our approach in other multilanguage settings (such as Java/C).

Bibliography

- [1] Allen, N., Krishnan, P., Scholz, B.: Combining type-analysis with points-to analysis for analyzing Java library source-code. In: SOAP@PLDI, ACM (2015)
- [2] Balakrishnan, G., Reps, T.W.: Recency-abstraction for heap-allocated storage. In: SAS, Springer (2006)
- [3] Blanchet, B., Cousot, P., Cousot, R., Feret, J., Mauborgne, L., Miné, A., Monniaux, D., Rival, X.: Design and implementation of a special-purpose static program analyzer for safety-critical real-time embedded software. In: The Essence of Computation, Lecture Notes in Computer Science, vol. 2566, pp. 85–108, Springer (2002)
- [4] Brown, F., Narayan, S., Wahby, R.S., Engler, D.R., Jhala, R., Stefan, D.: Finding and preventing bugs in JavaScript bindings. In: SP, IEEE Computer Society (2017)
- [5] Bucur, S., Kinder, J., Candea, G.: Prototyping symbolic execution engines for interpreted languages. In: ASPLOS, pp. 239–254, ACM (2014)
- [6] Buro, S., Crole, R.L., Mastroeni, I.: On multi-language abstraction - towards a static analysis of multi-language programs. In: SAS, Springer (2020)
- [7] Buro, S., Mastroeni, I.: On the multi-language construction. In: ESOP, Springer (2019)
- [8] Chipounov, V., Kuznetsov, V., Candea, G.: S2E: a platform for in-vivo multi-path analysis of software systems. In: ASPLOS, pp. 265–278, ACM (2011)

- [9] Cousot, P., Cousot, R.: Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: POPL, ACM (1977)
- [10] Duncan, C.: Native-code and shader implementations of perlin noise for Python. <https://github.com/caseman/noise> (2021), accessed: 2021-04
- [11] Fromherz, A., Ouadjaout, A., Miné, A.: Static value analysis of Python programs by abstract interpretation. In: NFM, Springer (2018)
- [12] Furr, M., Foster, J.S.: Checking type safety of foreign function calls. In: PLDI, ACM (2005)
- [13] Furr, M., Foster, J.S.: Polymorphic type inference for the JNI. In: ESOP, Springer (2006)
- [14] Furr, M., Foster, J.S.: Checking type safety of foreign function calls. ACM Trans. Program. Lang. Syst. (4) (2008)
- [15] Haapala, A., Määttä, E., CD, J., Ohtamaa, M., Necas, D.: Levenshtein Python C extension module. <https://github.com/ztane/python-Levenshtein/> (2021), accessed: 2021-04
- [16] Hu, M., Zhang, Y.: The Python/C API: evolution, usage statistics, and bug patterns. In: SANER, IEEE (2020)
- [17] Jakubek, A., Gałczyński, R.: Linked lists for CPython. <https://github.com/ajakubek/python-llist> (2021), accessed: 2021-04
- [18] Journault, M., Miné, A., Monat, R., Ouadjaout, A.: Combinations of reusable abstract domains for a multilingual static analyzer. In: VSTTE, Springer (2019)
- [19] Journault, M., Miné, A., Monat, R., Ouadjaout, A.: MOPSA: Modular Open Platform for Static Analysis. <https://gitlab.com/mopsa/mopsa-analyzer> (2021), accessed: 2021-04
- [20] Kondoh, G., Onodera, T.: Finding bugs in Java native interface programs. In: ISSTA, ACM (2008)
- [21] Kristensen, E.K., Møller, A.: Reasonably-most-general clients for JavaScript library analysis. In: ICSE, IEEE / ACM (2019), URL <https://doi.org/10.1109/ICSE.2019.00026>
- [22] Lee, S., Lee, H., Ryu, S.: Broadening horizons of multilingual static analysis: Semantic summary extraction from C code for JNI program analysis. In: ASE, IEEE (2020)
- [23] Li, S., Tan, G.: Finding bugs in exceptional situations of JNI programs. In: CCS, ACM (2009)
- [24] Li, S., Tan, G.: JET: exception checking in the Java native interface. In: SPLASH, ACM (2011)
- [25] Li, S., Tan, G.: Exception analysis in the Java Native Interface. Sci. Comput. Program. (2014)
- [26] Li, S., Tan, G.: Finding reference-counting errors in Python/C programs with affine analysis. In: ECOOP, Springer (2014)
- [27] Malcolm, D.: A static analysis tool for CPython extension code. <https://gcc-python-plugin.readthedocs.io/en/latest/cpychecker.html> (2018), accessed: 2021-04

- [28] Mao, J., Chen, Y., Xiao, Q., Shi, Y.: RID: finding reference count bugs with inconsistent path pair checking. In: ASPLOS, ACM (2016)
- [29] Matthews, J., Findler, R.B.: Operational semantics for multi-language programs. *ACM Trans. Program. Lang. Syst.* (3) (2009)
- [30] Meyer, M.: Distance library. <https://github.com/doukremt/distance> (2021), accessed: 2021-04
- [31] Miné, A.: Field-sensitive value analysis of embedded C programs with union types and pointer arithmetics. In: LCTES, ACM (2006)
- [32] Miné, A.: The octagon abstract domain. *High. Order Symb. Comput.* **19**(1), 31–100 (2006)
- [33] Monat, R., Ouadjaout, A., Miné, A.: Static type analysis by abstract interpretation of Python programs. In: ECOOP, Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020)
- [34] Monat, R., Ouadjaout, A., Miné, A.: Value and allocation sensitivity in static Python analyses. In: SOAP@PLDI, ACM (2020)
- [35] Monat, R., Ouadjaout, A., Miné, A.: A Multi-Language Static Analysis of Python Programs with Native C Extensions (Jul 2021), URL <https://doi.org/10.5281/zenodo.5141314>
- [36] Muła, W., Ombredanne, P.: Pyahocorasick library. <https://github.com/WojciechMula/pyahocorasick> (2021), accessed: 2021-04
- [37] Ouadjaout, A., Miné, A.: A library modeling language for the static analysis of C programs. In: SAS, Springer (2020)
- [38] Rinetzký, N., Poetzsch-Heffter, A., Ramalingam, G., Sagiv, M., Yahav, E.: Modular shape analysis for dynamically encapsulated programs. In: ESOP, Springer (2007)
- [39] van Rossum, G., Lehtosalo, J., Łukasz Langa: Python Enhancement Proposal 484. <https://www.python.org/dev/peps/pep-0484/> (2021), accessed: 2021-08-03
- [40] van Rossum, G., the Python development team: Python/C API Reference Manual. <https://docs.python.org/3.8/c-api/index.html> (2021), accessed: 2021-04
- [41] Schnell, I.: Bitarray library. <https://github.com/ilanschnell/bitarray> (2021), accessed: 2021-04
- [42] Tan, G., Croft, J.: An empirical security study of the native code in the JDK. In: USENIX, USENIX Association (2008)
- [43] Tan, G., Morrisett, G.: Ilea: inter-language analysis across Java and C. In: OOPSLA, ACM (2007)
- [44] Typeshed contributors: Typeshed. <https://github.com/python/typeshed/> (2021), accessed: 2021-04
- [45] Wei, F., Lin, X., Ou, X., Chen, T., Zhang, X.: JN-SAF: precise and efficient NDK/JNI-aware inter-language static analysis framework for security vetting of Android applications with native code. In: SIGSAC, ACM (2018)