

Regular Paper

Dynamic Version Checking for Gradual Updating

SATSUKI KASUYA^{1,a)} YUDAI TANABE^{1,b)} HIDEHIKO MASUHARA^{1,c)}Received: December 6, 2024, Revised: March 11, 2025,
Accepted: March 21, 2025

Abstract: Programming with Version (PWV) is a programming paradigm that allows programmers to safely utilize multiple versions of the same package within a single program, facilitating flexible version updates of dependent packages. Existing PWV languages ensure consistent version usage so as not to break software behaviors by leveraging the type system of the base language. However, dynamically typed languages need a mechanism to support multiple versions with an efficient method of ensuring consistent version usage without a type system. To introduce PWV features into dynamically typed languages, we propose a dynamic version checking (DVC) mechanism. It records version information in a value, propagates it during evaluation, and checks inconsistency using version information recorded in values. When an inconsistency is detected, the mechanism suggests how to modify the program to resolve potential semantic errors from the inconsistency. We develop Vythion, a Python-based PWV language with DVC, and implement its compiler. The compiler translates a Vythion program into a Python program with bitwise operations. Our performance measurement shows the DVC mechanism's overhead is scalable and acceptable for small programs but requires further optimization for real-world use. Additionally, we conduct a case study and discuss future directions to facilitate smoother updates in practical development.

Keywords: Software maintenance, Software migration, Dependency management, Compiler, Python

1. Introduction

Updating the version of upstream packages is one of the most troublesome tasks for downstream developers [17], [20]. An incompatible new version can break the behavior of downstream programs [10], [16]. Each new release of upstream packages requires downstream developers to assess its impact and modify their source code accordingly.

Replacing an upstream package with its new version is automated by package managers such as `pip`^{*1} in Python. For example, developers using NumPy [11] can automatically install the latest version by running `pip install --upgrade numpy`. Many developers benefit from this automation, as packages like NumPy are widely used across various domains, such as data analysis, deep learning, and image processing, in libraries like Pandas [30], PyTorch [25], and OpenCV [4].

Downstream developers carefully coordinate existing programs to update fundamental packages such as NumPy. The first major update of NumPy, version 2.0.0, was released in 2024. If any of the packages in use depends on NumPy 1.x series, the automatic installation of NumPy 2.0.0 via `pip` will fail. Manual installation, which is possible from the source, can break the existing behavior of downstream programs unintentionally, as some NumPy functions are incompatible with the old ones (see Ap-

pendix A.1).

Programming with Versions (PWV) [19], [27], [28] is a recent proposal designed to enable a gradual transition to new versions, thereby reducing update costs. The key ideas of PWV are (1) the simultaneous use of multiple versions, and (2) language mechanisms (i.e. types) that check version compatibilities. PWV languages ensure that programs use values created by compatible versions.

While previous research realized PWV in statically-typed languages, this research explores methods implementing PWV functionalities in dynamically-typed languages. To achieve this, we propose *dynamic version checking* (DVC) to alert when values of incompatible versions are used together at runtime. The DVC mechanism facilitates developers' communication regarding incompatibilities [18]; upstream developers specify compatibility for each function, allowing downstream developers to assess the impact of updates on their software through warnings.

The contributions of this paper are summarized as follows:

- We developed DVC, which records version information within values, propagates it during an evaluation, and detects inconsistencies based on the recorded version data.
- We implemented a Vythion compiler, a PWV language with DVC. The Vythion compiler translates Vythion programs into Python programs, and the DVC functions are compiled into efficient bitwise operations.
- We evaluated the runtime performance of the Vythion compiler. The results showed that Vythion is scalable and acceptable for debugging small programs as an offline analysis tool but requires further optimization for real-world use.
- We conducted a case study using gradual update scenarios

¹ Institute of Science Tokyo
Ookayama 2-12-1, Meguro, Tokyo 152-8552, Japan

a) satsuki.kasuya@prg.is.titech.ac.jp

b) yudaitnb@prg.is.titech.ac.jp

c) masuhara@acm.org

*1 `pip`: The PyPA recommended tool for installing Python packages.
<https://pip.pypa.io/> (Accessed January 10, 2025)




Fig. 1: Dependencies of the User program.

```

1 class SciPy: # SciPy 1.12.0:
2   def place_poles(A, B, poles):
3     return NumPy().solve(..) # Using NumPy 1.26.4
4
5 def my_place_poles(A, B, poles): # User Program
6   return NumPy().solve(..) # Using NumPy 2.0.0
7 NumPy().array_equal(
8   my_place_poles(A, B, poles),
9   scipy.place_poles(A, B, poles) ) # => False

```

Fig. 2: A program that uses NumPy and SciPy in Python

with the Vython compiler. Based on these findings, we discussed how Vython could work effectively in more practical scenarios.

The rest of the paper is organized as follows. Section 2 introduces inflexible update scenarios that motivate our research, and Section 3 offers an overview of the proposed features. Section 4 explains the semantics of the DVC mechanism, and Section 5 details the implementation of the Vython compiler. Sections 6 and 7 presents the performance evaluation and case study conducted using the compiler. Finally, Sections 8 and 9 discuss related work and provide concluding remarks.

2. Motivating Example

Consider a scenario where we update a User program that reimplements a function for solving pole placement problem^{*2} and test its behavior against SciPy [29] implementation. Figure 1 shows the dependencies of the User program. User depends on SciPy version 1.12.0, which indirectly depends on NumPy 1.26.4, and User directly depends on NumPy and attempts to update it from version 1.26.4 to 2.0.0.

As shown in Figure 2, both SciPy and User use the `solve` function from NumPy^{*3}. In User (Figure 2 bottom), `my_place_poles` is implemented using the `solve` function, and its results are compared against the existing implementation in SciPy. `place_poles` in SciPy 1.12.0 (Figure 2 top) directly returns the result of the `solve` function. We try to update NumPy in the User project.

Updating NumPy via pip

This attempt fails as follows.

```

1 $ pip install numpy==2.0.0
2 ERROR: scipy 1.12.0 requires numpy<1.29.0,>=1.22.4, but
   you have numpy 2.0.0 which is incompatible.

```

The error message indicates that this attempt resulted in broken

^{*2} This is a common task in control theory, placing closed-loop poles in desired locations to control the system response. The SciPy implementation can be found at SciPy Manual. https://docs.scipy.org/doc/scipy/reference/generated/scipy.signal.place_poles.html (Accessed January 10, 2025)

^{*3} These programs are simplified, but are essentially identical to the actual implementation. For more details, see Appendix A.2.

dependencies, as the installed SciPy is locked to NumPy versions below 1.29.0. Similarly, other Python package managers, such as poetry^{*4}, rye^{*5}, and uv^{*6}, also conservatively reject the installation of multiple versions of the same package within the development environment.

Updating NumPy from the Source

A potential workaround for using NumPy 2.0.0 without waiting for SciPy updates is to use NumPy 1.26.4 for SciPy and 2.0.0 for User independently. As mentioned earlier, while standard Python package managers do not support the installation of multiple versions of a package simultaneously, the `importlib` package^{*7} in the Python standard library allows dynamic switching to specific package versions (see Appendix A.2).

However, the workaround using the `importlib` package involves dynamically modifying module objects in `sys.modules`, which can result in unpredictable program behavior due to the subtle differences between the two or more versions of the same module. For example, the `solve` implementation was incompatibly changed in the NumPy 2.0.0 release. As explained in Appendix A.1, the ambiguous broadcasting rule was corrected in 2.0.0, so the `solve` function in the two versions may return different outputs even with the same input. As a result, the test of `my_place_poles` against `place_poles` fails, even if both implementations are identical. To mitigate such unfavorable situations, re-importing NumPy through the `importlib` package issues a warning due to the risk of unpredictable behavior.

```

1 UserWarning: The NumPy module was reloaded (imported a
   second time). This can in some cases result in
   small but subtle issues and is discouraged.
2 spec.loader.exec_module(numpy)

```

Downstream developers thus face with a dilemma: either undertake the extensive refactoring work required to adapt programs to a new version, or contend with hard-to-identify incompatibility errors while partially testing new features. The former choice necessitates tedious tasks including reading release notes and reviewing the implementations of all upstream packages to identify and resolve incompatibility errors through refactoring. Alternatively, the latter choice exposes programmers to limitations in current programming language systems, which lack mechanisms to detect the mixed use of incompatible implementation versions. [27], i.e., incorrect version usage is often reported as Python semantic errors, which fail to pinpoint the root cause stemming from version incompatibilities.

3. Safely Use Multiple Versions in Vython

Vython is a Python subset^{*8} that implements the PWV update

^{*4} Poetry, <https://python-poetry.org/> (Accessed January 10, 2025)

^{*5} astral-sh/rye, <https://github.com/astral-sh/rye> (Accessed January 10, 2025)

^{*6} astral-sh/uv, <https://github.com/astral-sh/uv> (Accessed January 10, 2025)

^{*7} importlib — The implementation of import, <https://docs.python.org/3/library/importlib.html> (Accessed January 10, 2025)

^{*8} Current Vython does not support some Python features, and the compiler covers about 60% of Python's AST node types. We believe that most unsupported features can be implemented without extra design considerations. However, introducing class inheritance may require extensions to the DVC design, making it a non-trivial task.

```

1 class NumPy!1.26.4():
2     def solve(self, A, B):
3         return res

```

```

1 class SciPy!1.12.0():
2     def place_poles(A, B, poles):
3         return NumPy!1.26.4().solve(..) # Using NumPy 1.26.4

```

```

1 def my_place_poles(A, B, poles): # User Program
2     return NumPy!2.0.0().solve(..) # Using NumPy 2.0.0
3 array_equal(
4     my_place_poles(A, B, poles),
5     SciPy!1.12.0().place_poles(A, B, poles) ) # => Warning!

```

Fig. 3: A program that uses NumPy and SciPy in Vython

model (*gradual updating*), which splits the burden of program modifications caused by package updates. Vython is designed to mitigate the version-locking problem by enabling version selection at individual code sites and provides debugging information to help address incompatibility errors with mixed package versions through the following features:

- **Using multiple versions in a code:** The programmer can selectively use multiple versions of a class definition by specifying a version when instantiating.
- **Dynamic version checking (DVC):** Vython records information about the class and its version used for creating a value, ensuring that programs use values created with consistent versions.

Vython differentiates multiple class versions internally, allowing for their selective use. As shown in Figure 3, the current naive implementation requires version annotations in the surface language. Additionally, DVC is intended to be enabled only in debugging mode. Vython has a production mode that deploys programs without runtime checks for DVC.

Vython provides a mechanism for upstream developers to specify compatibility, which is utilized in DVC as follows.

Upstream Developer Specifies Compatibilities in Code

In Vython, upstream developers are responsible for specifying incompatibilities. In NumPy 2.0.0 (Figure 3 top right), the NumPy developer uses `incompatible()` (denoted as `incomp()` below for brevity) to mark an expression as incompatible with previous versions. Additionally, upstream developers can provide guidance (as shown below) to help downstream developers.^{*9} This information is recorded along with the class definition in the source code.

```

1 [Changed in version 2.0] `NumPy().solve(a,b)`
2 - If `b` is 1-dim, it is treated as a column vector (M,).
3 - Otherwise, it is treated as a stack of (M, K) matrices.
4 [NumPy 1.26.4 and earlier]
5 - `b` was treated as a stack of (M,) vectors if `b.ndim`
   equald `a.ndim - 1`.
6 [Refactoring hint]
7 - The old behaviour can be reconstructed by using np.solve(a,
   b[... , None])[... , 0]

```

^{*9} Efforts to provide incompatibility and refactoring hints are beyond the scope of this research and will not be explained in this paper. The incompatibilities and refactoring hints here are formatted information from the NumPy 2.0.0 release notes. <https://numpy.org/devdocs/release/2.0.0-notes.html#removed-ambiguity-when-broadcasting-in-np-solve> (Accessed January 10, 2025)

Notifying Downstream Developers of Incompatibility Causes

Downstream developers can follow a gradual updating model, selectively experimenting with new features introduced in 2.0.0 while using NumPy 1.26.4 in the User program (Figure 3 bottom). However, during program execution, objects from both NumPy versions may coexist, potentially violating implicit version assumptions of the old version of the user program. This often leads to difficult-to-diagnose errors, as discussed in Section 2.

In this situation, the downstream developer benefits from DVC and the guidance for updates specified by the upstream developer. The DVC mechanism reports runtime warnings on lines 3-5 in the User program, along with the Python runtime errors.

```

1 ... (Python runtime errors) ...
2 Incompatible version usage found in lines 3-5:
3 - NumPy 1.26.4
4 - from place_poles() in SciPy 1.12.0
5 - NumPy 2.0.0
6 - from my_place_poles() in User
7 [Changed in version 2.0.0] `NumPy().solve(a,b)`
8 - If `b` is 1-dim, it is treated as a column vector (M,).
9 - Otherwise, it is treated as a stack of (M, K) matrices.
10 [NumPy 1.26.4 and earlier]
11 - `b` was treated as a stack of (M,) vectors if `b.ndim`
   equald `a.ndim - 1`.
12 [Refactoring hint]
13 - The old behaviour can be reconstructed by using np.solve(a,
   b[... , None])[... , 0]

```

The Vython warning indicates that the `array.equal` call in lines 3-5 relies on values derived from incompatible versions of the `solve` function. This warning helps downstream developers identify violations of their pre-update assumptions about the specific version of an object, for instance, the downstream developer assumed that `array.equal` compares values using the same version of `NumPy.solve`. By following the refactoring hints provided by upstream developers, downstream developers can refactor their programs to ensure compatibility with version 2.0.0. By iterating this process, they can gradually migrate most parts of the User program to version 2.0.0.

4. Vython Semantics for DVC

4.1 Intuition to Vython Semantics

Vython associates version information with each object and ensures that method return values are created from a combination of consistent version implementations. This version information reflects the versions of implementation used to create the object. Version information is recorded in a format called version table (VT). Additionally, Vython considers any value derived from another object created using version *V* as also originating from version *V*. This subsection illustrates the principle behind the design decision through examples.

4.1.1 Version Tables Recorded in Objects

In Vython, all class instances will record the information of their instantiated class and its version when it is created. For example, the following program creates an instance of the NumPy class in Vython, specifying version 2.0.0.

```

1 x = NumPy!2.0.0()

```

The Vython runtime with DVC records version information in the NumPy instance, indicating that it was made from version 2.0.0 of the NumPy class. The VT of an instance of NumPy version 2.0.0 immediately after instantiation is represented as follows.

class	NumPy
version	2.0.0
flag	-

The `class` and `version` fields indicate the class and version of the implementation used to create the value, and the `flag` field indicates any incompatibility of the implementation with other versions. Here, the `flag` is set to `-`, and the VT is simply recording class and version information.

4.1.2 Version Information Propagation

At the time of returning a value from method invocation or field access, Vythion appends the version information of the receiver object to the return value. The following program performs a method invocation on an instance of NumPy version 2.0.0 on line 8.

```

1 class NumPy!2.0.0():
2     def add(self, a, b):
3         res = a + b
4         return res
5
6 x = NumPy!2.0.0()
7 a, b = 1, 2
8 y = x.add(a, b)

```

The version information recorded in the value created by `x.add(a, b)` merges the version information of the return value from the `add` method (on line 7) with the version information of the value bound to `x`, as follows.

class	Int	NumPy
version	1	2.0.0
flag	-	-

Here, the records for `Int` and `NumPy` come from the VTs of the return value of the `add` method and the `NumPy` instance bound to `x`, both of which are shown below.

class	Int
version	1
flag	-

VT of the return value of the `add` method

class	NumPy
version	2.0.0
flag	-

VT of the `NumPy` instance bound to `x`

4.1.3 Dynamic Version Checking

Vythion checks whether the value is created using consistent version implementations on the return values of method calls and field accesses. If the value is created using inconsistent version implementations, Vythion considers the method call to be a usage of potentially incompatible implementation and issues a warning to the programmer. The following program compares arrays created by incompatible implementations of the `solve` method.

```

1 class NumPy!1.26.4():
2     def solve(self, A, B):
3         return res
4
5 class NumPy!2.0.0():
6     def solve(self, A, B):
7         return res.incomp(..)
8
9 def array_equal(arr1, arr2):
10     ..
11     arr1 == arr2
12     ..
13     a, b = [...]
14     res1 = NumPy!1.26.4().solve(a, b)
15     res2 = NumPy!2.0.0().solve(a, b)
16     array_equal(res1, res2)

```

As described in Section 3, in Vythion, upstream developers can use the `incomp` method to explicitly declare the incompatibility introduced in version 2.0.0 of the `solve` function, thereby encouraging downstream developers to use the `solve` function with caution regarding its incompatibility. The design of specifying incompatibilities at the expression level within a method, rather than at the method definition level, allows fine-grained incompatibility specification for individual elements within a data structure. On lines 6 and 7, Vythion records the following VTs for the values bound to `res1` and `res2`, respectively.

class	NumPy	...
version	1.26.4	...
flag	-	...

VT of the value bound to `res1`

class	NumPy	...
version	2.0.0	...
flag	True	...

VT of the value bound to `res2`

Note that the `flag` field is set to `True` for the value bound to `res2` (right). For VTs where the `flag` field is set to `True`, Vythion performs consistency checking at the return point of method calls and field accesses. For example, on line 8, the `array_equal` function is called with `res1` and `res2` as arguments. Inside the `array_equal` function, on line 3, the equality of the vectors `arr1` and `arr2` is checked. At this point, Vythion performs a consistency checking on the return value of `==` on line 3 and outputs a warning and refactoring hints set by the upstream developer, as described in Section 3.

4.1.4 Difference between Vythion and Existing PWV Languages

An important design decision for DVC is that the version information of function arguments is not directly reflected in the return value. This design decision leads to differences in consistency-checking capabilities compared to the existing PWV languages. For example, in cases where an argument is merely discarded during the function body, the version information of the argument is not propagated into the version information of the return value.

```

1 class A!1.0.0(): ..
2 class A!2.0.0(): ..
3 class Choose!1():
4     def discard_snd(self, fst, snd):
5         return first
6
7 c = Choose!1.0.0()
8 x = c.discard_snd(A!1.0.0(), A!2.0.0())

```

As shown in line 5, the `discard_snd` method does not use `snd` in the actual computation, and therefore, the VT of the return value does not include the VT of the second argument. On the other hand, existing PWV languages such as VL [26], [27], [28] and BatakJava [19] conservatively reflect the version information of arguments in the return value.

Another interesting aspect of DVC is path sensitivity.

```

1 class A!1.0.0(): ..
2 class A!2.0.0(): ..
3 if rand():
4     x = A!1.0.0()
5 else:
6     x = A!2.0.0()
7 y = x

```

With the DVC mechanism, the recorded version information in the value bound to `y` in line 7 is only the one associated with the path actually taken. Therefore, upon the completion of the above program's execution, the VT of the value assigned to `y` is

the lower left VT if the then branch is executed, and lower right VT if the else branch is executed, as shown below.

class	A
version	1.0.0
flag	-

through then branch

class	A
version	2.0.0
flag	-

through else branch

In contrast, in existing PWV languages, a type-system-based analysis is performed for the above example; that is, the analysis enforces the result where the versions (and types) of the then and else branches completely match (meet), or become the join of both.

4.2 The Vython Semantics

This section presents the Vython semantics, which are defined through DVC functions over VT. Before detailing the DVC extensions applied to the underlying Python semantics, we begin by defining version tables.

4.2.1 DVC Functions for VT

We define VT as follows.

Definition 4.1 (Version Tables). A version table (VT) is a set of triples $\{(C, V, f)\}$, where C represents the class, V represents the version associated with a constructor, method, or field implementation, and f represents the flag that can take the value $-$ or \mathbf{T} (standing for True). f indicates incompatibility with other versions of the implementation when $f = \mathbf{T}$.

Note that each VT does not have duplicate triples and the order of recording triples does not relate to the meaning of the VT.

Next, we define DVC functions for VTs. All operations on the VT are performed using the following functions.

Definition 4.2 (DVC functions). The mk function takes a class name C , a version number V , and a flag f , and creates a VT of size 1. The join function takes multiple VTs and creates a union set of VTs. The wf function takes a VT and checks whether the VT is well-formed or not.

$$\begin{aligned} \text{mk}(C, V, f) &= \{(C, V, f)\} & \text{join}(vt_1, vt_2) &= vt_1 \cup vt_2 \\ \text{wf}(vt) &= \begin{cases} \text{incomp} & \exists V_1, V_2. V_1 \neq V_2 \\ & \wedge (C, V_1, \mathbf{T}) \in vt \\ & \wedge (C, V_2, f) \in vt \\ \text{comp} & \text{otherwise} \end{cases} \end{aligned}$$

Here, incomp indicates that the value with the associated VT was produced by a computation involving a value originating from an incompatible class. At the implementation level, a warning is issued at the time when wf is called. Conversely, comp indicates that the value associated with the VT is derived from values generated with consistent class versions, allowing evaluation to proceed without any warnings.

To aid the reader's understanding, we provide several examples using DVC functions.

Example 4.1 (VT operations (join and mk)).

```
join(mk(NumPy, 2.0.0, T), mk(NumPy, 2.0.0, T))
= {(NumPy, 2.0.0, T)}
join(mk(NumPy, 2.0.0, T), mk(NumPy, 2.0.0, -))
= {(NumPy, 2.0.0, T), (NumPy, 2.0.0, -)}
join(mk(NumPy, 2.0.0, T), mk(Array, 1.0.3, -))
= {(NumPy, 2.0.0, T), (Array, 1.0.3, -)}
```

Example 4.2 (VT consistency checking (wf)).

```
wf({(NumPy, 2.0.0, T), (NumPy, 2.0.0, -)}) = comp
wf({(NumPy, 2.0.0, T), (NumPy, 1.26.4, -)}) = incomp
wf({(NumPy, 2.0.0, T), (NumPy, 1.26.4, T)}) = incomp
```

4.2.2 The Vython Semantics Using DVC Functions

Using DVC functions, we define the Vython semantics. The current Vython is a minimal object-oriented language because it lacks support for abstract classes, interfaces, class inheritance, and method overriding.

To highlight VT operations, we present the big-step operational semantics as an extension of standard object-oriented language semantics like Featherweight Java [13] in a somewhat informal manner. Figure 4 illustrates an excerpt of the Vython semantics. The meta-variables t, v, C, m , and f represent terms, values, method names, and field names, respectively. The meta-functions fields and mbody represent the lookup of field names and method arguments with their bodies for a given class. Method calls $t.m()$ and field accesses $t.f$ follow the standard notation, but note that $C!V(\bar{t}_i)$ and $C(\bar{v}_i)$ denotes instance creation (term) and class instance (value).

Furthermore, while the syntax allows the incompatible tag $t.\text{incomp}(C, V)$ to be applied to any terms, we assume it will be used specifically in the return statements of methods or constructors. The arguments C and V represent the class and version in which the method or constructor is defined, and they are expected to be automatically inserted during the compiler's preprocess.^{*10}

As discussed in Section 4.1, (E-FIELD) and (E-INVK) follow similar patterns since both field accesses and method calls involve accessing an instance's attributes. The VT of a method call's return value is determined by concatenating the VTs recorded in two components: the method's return value and the receiver object. Similarly, the VT for a field access result is derived by concatenating the VTs recorded in the field-assigned value and the target object. Lastly, both rules include a consistency check using wf .

(E-BUILTINOP) defines the rule for VT operations on builtin type values, such as boolean and numerical operations. In Vython, these operations are treated differently from method calls, as they are considered computations that always use the argument values. Therefore, the VTs of all arguments are joined into the VT of the resulting value.

(E-INCOMP) defines the VT semantics for the incompatibility tag. The incomp function is the only mechanism for assigning the

^{*10} The current Vython compiler does not yet support this feature. In the current implementation, incomp calls require both C and V in a surface program.

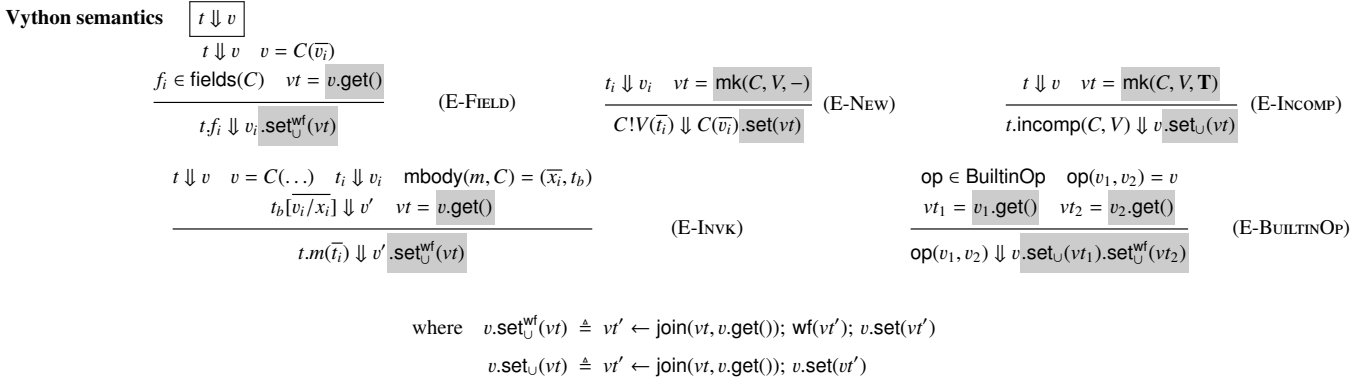


Fig. 4: An excerpt of the Vyhton semantics with the highlight of VT operations. `get` and `set` are getters and setters for the VT of each object, and they are assumed to be defined in all classes. `BuiltinOp` represents the set of operators (`op`) for Python's built-in types, such as `+`, `or`, and `=`, etc. The operator `op` also is a meta-level function for constant objects v_1 and v_2 . For example, when `op = +`, `op(Int(1), Int(2)) = Int(3)` that has an empty VT.

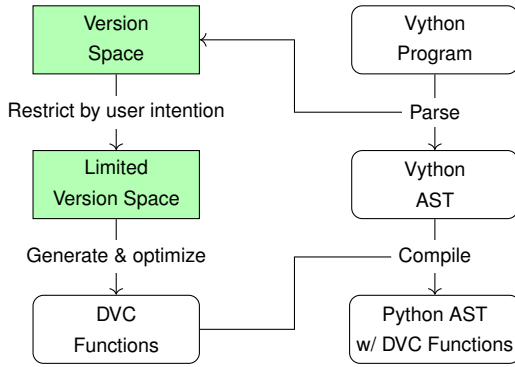


Fig. 5: Compilation flow of the Vyhton compiler.

T flag, which triggers consistency checking at a `wf` call. In contrast, (E-NEW) specifies the VT to be recorded for a newly created class instance, recording a `mk(C, V, -)`. Unlike (E-INCOMP), the attached checker flag here is `-`, meaning it only records version information without triggering DVC except in the cases where computations involve values with **T** flag assigned by `incomp`.

5. Implementation

5.1 The Vyhton Compiler

We implemented the Vyhton compiler which translates the Vyhton program into the Python program. In the Python programs generated after compilation, the DVC mechanism is achieved by representing a VT as an attribute of each object and operating and checking VTs by predefined DVC functions, which are inserted against class methods, field accesses, and class instantiation. Since we treat all values as objects, Vyhton literals are compiled into predefined Python classes.

Figure 5 shows the compilation flow of the Vyhton compiler. First, the Vyhton program is parsed into a Vyhton AST using the parser library `lark`^{*11}. During parsing, the compiler extracts version information for classes and constructs version space (available class and version pairs) globally. Programmers can further restrict the size of version tables by specifying the classes they

```

1 class A!1.0.3():
2     def __init__(self, value):
3         self.value = value
4     def get(self):
5         return self.value
6
7
8
9
10 class A_v_1_0_3:
11     def __init__(self, value):
12         self.vt = 1
13         self.value = value
14
15     @_vt_invk
16     def get(self):
17         return _vt_field(self, self.value)
18
19 ..
20 def _vt_invk(func):
21     def wrapper(*args, **kwargs):
22         result = func(*args, **kwargs)
23         if result is not None:
24             result = _vt_join(result, args[0])
25         if not _vt_well_formed(result):
26             _issue_warning(result, args[0])
27         return result
28     return wrapper
29 def _vt_field(receiver, result):
30     result = _vt_join(result, receiver)
31     if not _vt_well_formed(result):
32         _issue_warning(result, receiver)
33     return result
  
```

(a)

(b)

Fig. 6: Simple Vyhton program (a) before and (b) after compilation.

want to use with multiple versions as input to the compiler. The limited version space is then used to generate DVC functions. Finally, the Vyhton AST is compiled into Python AST, along with the insertion of generated DVC functions.

The Vyhton compiler inserts DVC functions by each evaluation rule. Figure 6 illustrate a program before and after compilation for a class `A` version 1.0.3, which only includes getter `get`. As shown in line 6 of Figure 6(b), the VT operation for method calls is implemented through the decorator function `_vt_invk`, and the VT operation for field accesses is implemented through `_vt_field` in line 8. Additionally, as shown in line 3, the con-

^{*11} `lark-parser/lark`, <https://github.com/lark-parser/lark> (Accessed January 10, 2025)

structor bit-encodes the initial VT and assigns it to the `vt` field (`mk`).

Among the implementation functions introduced so far, we will use the implementation of `_vt_invk`, which corresponds to (E-Invk), as an example to explain. When a method with this decorator is invoked, the `_vt_invk` function performs VT operations according to (E-Invk). First, the method is executed in line 12. Then, VT propagation to the return value is handled by `_vt_join` (`join`) in line 14, and the VT consistency of the return value is checked by `_vt_well_formed` (`wf`) in line 15.

5.2 Optimization

To avoid the significant runtime overhead anticipated, Vythion compiles VTs into bit sequences and DVC functions into combinations of bitwise operations.

5.2.1 Encoding VTs as Bit Sequences

Algorithm 1 shows the algorithm for encoding a VT into a bit sequence. The algorithm takes as input the version space \mathcal{V} (limited in size by the user) and the vt to be encoded, and outputs the bit sequence representation of vt under \mathcal{V} .

Note that the current implementation of Vythion is restricted to two versions per class within a program. This limitation is based on the assumption that, in most cases, programmers are primarily concerned with compatibility between two specific versions. By leveraging this assumption, a bit-encoded VT is a bit sequence whose length is equal to four times the number of classes in the limited version space. Every grouped set of four bits records what version of a certain class was used to create the value.

For example, consider a program whose resulting limited version space consists of NumPy versions 1.26.4 and 2.0.0.

Example 5.1 (Bit-encoded VT). The following version table is encoded into the bit sequence 1101 under the version space $\{\text{NumPy} \mapsto [1.26.4, 2.0.0]\}$.

class	NumPy	NumPy	NumPy
version	2.0.0	2.0.0	1.26.4
flag	True	-	-

The length of the bit sequence encoding the VT, $|b|$ in Algorithm 1, is 4. The first, second, third, and fourth bits correspond versions 1.26.4 with flag -, 1.26.4 with flag T, 2.0.0 with flag -, and 2.0.0 with flag T, respectively. Following the three elements contained in the version table, the encoding algorithm returns a bit sequence 1101 with the first, third, and fourth bits set to 1.

5.2.2 Encoding Helper Functions as Bitwise Operations

Along with the encoding of VT into a bit string, DVC functions are also encoded into bitwise operations. The three DVC functions, `mk`, `join`, and `wf` are represented using bitwise operations as follows.

Definition 5.1 (DVC functions (bit-encoded)).

Algorithm 1: VT Encoding Algorithm

Data:

- \mathcal{V} : Limited version space, a map from class names to available versions $\{C_i \mapsto [\overline{V_{ij}}]\}^{*12}$ (Note: $\mathcal{V}[C_i]$ denotes the set of available versions for class C_i .)
- $vt = \{(C_i, V_{ij}, f_i) \mid C_i \in \text{dom}(\mathcal{V}), V_{ij} \in \mathcal{V}[C_i], f_i \in \{\mathbf{T}, -\}\}$

Result: A bit sequence b representing the vt under \mathcal{V} .

```

1 begin
2    $|b| \leftarrow 2 \times \sum_{C \in \text{dom}(\mathcal{V})} |\mathcal{V}[C]|;$ 
   //  $|\mathcal{V}[C]| = 2$  in our assumption
3   for  $i \leftarrow 0$  to  $|b| - 1$  do
4      $b[i] \leftarrow 0;$ 
5   foreach  $(C_i, V_{ij}, f_i) \in vt$  do
6      $offset \leftarrow 2 \times (j + \sum_{k=0}^{i-1} |\mathcal{V}[C_k]|);$ 
   //  $|\mathcal{V}[C_k]| = 2$  for any  $k$  in our assumption
7     if  $f_i = \mathbf{T}$  then
8        $index \leftarrow offset + 1;$ 
9     else if  $f_i = -$  then
10       $index \leftarrow offset;$ 
11     $b[index] \leftarrow 1;$ 
12  return  $b;$ 

```

$$\text{mk}(C, V, f) = \llbracket \{(C, V, f)\}, \mathcal{V} \rrbracket_{\text{bit}}$$

$$\text{join}(\hat{v}_1, \hat{v}_2) = \hat{v}_1 \mid \hat{v}_2$$

$$\text{wf}(\hat{v}t) = \begin{cases} \text{incomp} & ((\hat{v}t \gg 1) \& \hat{v}t) \gg 1 \\ & | (\hat{v}t \gg 3) \& \hat{v}t \\ & \& \text{mask} \neq 0 \\ \text{comp} & \text{otherwise} \end{cases}$$

where `mask` is a bit sequence of the form (0001)+,

\mathcal{V} is a limited version space,

and $\hat{v}t$ represents the bit-encoded vt .

The `mk` function is compiled into a bit sequence determined by $\llbracket * \rrbracket_{\text{bit}}$, the encoding function defined in Algorithm 1. This function is evaluated at compile time, producing hard coded bit sequences in the Python AST, thus incurring no runtime overhead with $\llbracket * \rrbracket_{\text{bit}}$. The `join` function is simply compiled into a bitwise OR operation (\mid).

Compared to the other two DVC functions, `wf` is less straightforward and requires a detailed explanation. The purpose of `wf` is to detect when a VT element (C, V_1, \mathbf{T}) exists and another element (C, V_2, f) is present in the input VT, where $V_2 \neq V_1$ and f can be either \mathbf{T} or $-$, returning `incomp` as a result. Assuming each class has only two versions, this detection corresponds to recognizing one of the following bit patterns.

- 1 _ _ 1
- 1 _ 1 _
- _ 1 1 _

Furthermore, we focus on (E-Incomp), which, as previously noted, is the sole mechanism for assigning the \mathbf{T} flag, and optimizing its behavior at the implementation level. As defined in Definition 4.2, if a VT element (C, V_1, \mathbf{T}) exists in vt , adding $(C, V_1, -)$ – an element with the same class name and version name – to vt does not affect the results of the `wf` functions with

^{*12} In the expression $C_i \mapsto [\overline{V_{ij}}]$ and V_{ij} , the bar notation is introduced as a shorthand to explicitly indicate the indexing of elements in an iterable set or vector. Specifically, it abbreviates the typical right-hand side of a set comprehension by making explicit that each element is associated with a specific index within the collection. For example, $\{C_i \mapsto [\overline{V_{ij}}]\}$ is an abbreviation for the set comprehension $\{C_i \mapsto [\overline{V_{ij}}] \mid i \in \text{Set}\}$.




Fig. 7: Overhead of DVC functions in Vythion (left) for simple algorithm benchmarks and (right) for repeating additions 2000 times with the number of VT entries.

other VT elements, such as (C, V_2, f) . Accordingly, Vythion implements (E-INCOMP) with $vt = \text{join}(\text{mk}(C, V, T), \text{mk}(C, V, -))$ in the premise.

By leveraging this optimization, the bit pattern $1_1_$ will always simplify to 1111 , which is subsumed by the other two patterns. Therefore, it suffices to detect the following two patterns:

- $1_1_$
- $_11_$

To achieve this, the $1_1_$ pattern is detected by $((vt \gg 1) \& vt) \gg 1 \& \text{mask} \neq 0$, and the $_11_$ pattern is detected by $((vt \gg 3) \& vt) \& \text{mask} \neq 0$. Further optimization is performed to eliminate redundant operations, resulting in the current wf definition in Definition 5.1.

6. Performance Evaluation

We conducted preliminary experiments on runtime performance. This performance evaluation focuses on the overhead introduced by DVC in Vythion's debugging mode. In contrast, the production mode (without the DVC feature) incurs no additional runtime cost, as it simply executes a Python program where multiple class versions are distinguished by their names.

6.1 Settings

We ran the following two benchmarks and calculated the average over 1000 iterations. The experiments were conducted with Python 3.12.1 on an Intel Core i5-10400F running Windows 11 23H2.

Benchmarks

- **Simple Algorithm:** Vythion programs implementing four simple algorithm insert, sort, fib, sieve, and is_prime^{*13} using a VT with a maximum of two entries.
- **Scalability:** a Vythion program that repeats additions 2000 times, with the number of VT entries doubling from 2^0 to 2^{11} .

The five algorithm in the simple algorithm benchmark were chosen to examine the overhead trends of the DVC feature in programs that can be written using the current Vythion. insert and sort rely heavily on user-defined class instances, whereas fib, sieve, and is_prime do not.

Each benchmark examines where the overhead occurs by dis-

abling certain VT operations through the following Vythion compiler options.

Compiler options

- (1) **python:** baseline, production mode.
- (2) **wrap-literals:** compiling literals as with VTs.
- (3) **mk:** (2) + VT Initialization at instantiations.
- (4) **join:** (3) + VT propagation.
- (5) **wf (vython):** (4) + consistency check.

6.2 Result and Discussion

In this subsection, we present and discuss the performance evaluation results for simple algorithm and scalability benchmarks.

6.2.1 Simple Algorithm

Result

Figure 7 (left) shows the execution times of simple algorithm relative to python. Focusing on each algorithm's case wf (vython), the program dominated by method calls to user-defined class instances, such as insert and sort, exhibit an overhead mostly within 20 to 30 times. In contrast, programs dominated by arithmetic or boolean operations, such as fib and sieve, show a larger overhead ranging from 60 to 70 times.

Focusing on individual DVC functions, it is evident that wrap-literals and join exhibit significant overhead across benchmarks. These two DVC functions account for 80% of the overhead in sort, is_prime, fib, and sieve.

Compared to other dynamic analysis tools for Python (i.e. DynaPyt [9]), which generally exhibit an overhead of up to 20x, the current overhead of Vythion is not practically acceptable and requires further optimization.

Discussion on the Overhead by wrap-literals

The current Vythion attaches VT to all Python primitive values, even when they are involved in computations entirely unrelated to the classes of interest to the programmer. This negates bytecode optimizations for constant calculations.

Overhead can be reduced by discontinuing compile-time literal wrapping and instead performing dynamic casts only when computations interact with classes in the version space. An ad-hoc implementation of this approach is marking pure Vythion functions with the @pure annotation and enclosing them in a try-catch block. When a DVC function is invoked on a Python value, any resulting error is caught and used to cast the value to a Vythion

^{*13} Please see Appendix A.3 for the implementation of each benchmark program.




Fig. 8: Overhead of DVC functions in Vython for simple algorithm benchmarks with dynamic literal casts.

value.

Figure 8 shows the results of a simple algorithm benchmark using this ad-hoc compilation method. The results demonstrate that, apart from `insert`, which could not be annotated with `@pure` due to its reliance on destructive updates, all benchmarks achieved performance comparable to `python`.

Another possible approach to reducing overhead is unboxing optimization, which replaces computations on boxed objects with equivalent operations on lightweight, unboxed values. Applying this to Vython is expected to significantly reduce the number of extractions from and insertions into wrapper class instances during primitive value computations.

Discussion on the Overhead by join

The current approach performs `join` on every field access, method call, and built-in operation. By incorporating static analysis, to compose DVC functions, we believe it will be possible to omit most `join`. The details of this idea are discussed in Section 8.2.

6.2.2 Scalability

Figure 7 (right) shows that overhead does not increase significantly as the VT size grows. However, it is notable that the overhead only for `wf` (red line) shows an increase around a VT size of $2^9 - 2^{11}$.

It is difficult to determine the number of class versions that must be considered in practical application update scenarios. As discussed in Section 5, the version space determining the maximum VT size can be managed and reduced by the programmer. However, updating a class may necessitate including additional upstream classes in its dependency.

Therefore, while additional case studies are necessary to confirm that the VT size remains manageable for practical scenarios, the results indicate that Vython debugging mode is scalable, particularly for non-large-scale applications.

7. Case Study

In the case study, we implement Keyword In Context (KWIC) according to the second modularization criteria presented by Parnas [24], and verify incompatible updates in downstream programs that use it. The KWIC implementation creates a dictionary that includes the context in which each word is used, as shown below.

¹ "reverse-engineering weep ReLU networks"
² "My Fair Bandit: Distributed Learning of Max-Min Fairness




Fig. 9: Class dependency diagram.

with Multi-player Bandits"
 3 "Scalable Differentiable Physics for Learning and Control"

KWIC input (ICML paper titles)

```

1 2 ferentiableScalable|andControl
2 1 FairMy|Bandit:Distrib
3 1 ibutedBandit:FairMy|Bandits
4 2 ferentiableScalable|Control
5 2 Scalable|Differentiable
6 1 Bandit:FairMy|DistributedLea
7 1 My|FairBandit:Dis
8 1 ibutedBandit:FairMy|FairnesswithMu
9 2 ferentiableScalable|forLearningand

```

KWIC output

7.1 Settings

We outline the experimental setup for the KWIC case study. We describe the overall class structure and program, as well as an update scenario in which selected call sites transition from version 1 to version 2 of the `String` class.

Class Structure

Figure 9 shows the class dependencies of the KWIC implementation. We define that class A depends on class B if class A's definition includes at least one of the following: a B instantiation, a method call of a B object, or field access of a B object. When class A depends on class B, any modifications to class B necessitate a code review of class A's implementation.^{*14} We implemented Parnas's module structure as individual classes shown in Figure 9. Additionally, we introduced the `String` class, which assumes standard library APIs and was implicitly used by each module in the original paper. The roles of the other classes remain the same as described in the original paper.

KWIC Implementation

Figure 10 shows the definition of a series of classes for implementing KWIC. Passing a list of `String` classes to `Integrate.main()` will output the corresponding KWIC index. These classes are implemented using version 1 of the `String` class. Notably, in the `sort` method of the `Sort` class, the program processes words one letter at a time, converting lowercase letters to uppercase. For further implementation details, please refer to the published GitHub repository.^{*15}

^{*14} It is important to note that the absence of a direct dependency from class A to class B does not guarantee that a review is unnecessary when class B is updated. Identifying such implicit errors is the purpose of DVC.

^{*15} prg-titech/kwic-vython, <https://github.com/prg-titech/kwic-vython> (Accessed January 10, 2025)

```

1 class LineStore!1():
2   def __init__(self):
3     self.rows = []
4
5   def set_char(self, row, word, offset, char):
6     current_word = self.get_word(row, word)
7     while offset >= len(current_word):
8       current_word.append("")
9     current_word[offset] = char
10
11 class Input!1():
12   def input(texts):
13     line_store = LineStore!1()
14     for row_i in range(len(texts)):
15       line = texts[row_i]
16       split_line = line.split()
17       for word_i in range(len(split_line)):
18         word = split_line[word_i]
19         for char_i in range(word.size()):
20           line_store.set_char(row_i, word_i, char_i, word.
21                               get(char_i))
22     return line_store
23
24 class Rotate!1():
25   def __init__(self, line_store):
26     self.line_store = line_store
27     self.shift_table = []
28     for row in range(line_store.num_rows()):
29       for word in range(line_store.num_words(row)):
30         self.shift_table.append((row, word))
31
32 class Sort!1():
33   def __init__(self, rotate):
34     self.rotate = rotate
35
36   def first_shift_to_str(self, shift):
37     keyword = String!1("")
38     for char_i in range(self.rotate.num_chars(shift, 0)):
39       char = (self.rotate.get_char(shift, 0, char_i))
40       if String!1("a").get(0) <= char <= String!1("z").get
41         (0):
42         char = chr(char - 32)
43         keyword.add(String!1(char))
44     return keyword
45
46   def do_sort(self):
47     self.row_indices = sorted(
48       range(self.rotate.num_rows()),
49       key=lambda r: self.first_shift_to_str(r),
50     )
51     return self
52
53 class Output!1:
54   def __init__(self, sort):
55     self.sort = sort
56
57   def output(self):
58     ..
59
60 class Integrate!1:
61   def main(self, titles):
62     line_store = Input.input(titles)
63     rotated = Rotate(line_store)
64     sorted_rotate = Sort(rotated).do_sort()
65     Output(sorted_rotate).output()

```

Fig. 10: Classes for KWIC

Update Scenario

We define an update scenario for the downstream User's program which uses the KWIC implementation as follows. Initially, the User program relied on only version 1 of the String class. However, a requirement arises to use a new method introduced `include_pattern` in version 2 of the String class for a specific process, while keeping the use of the KWIC implemented using version 1 of the String class. Therefore, the downstream programmer selectively updates some of the call sites of the String class in their program to use version 2.

Figure 11 shows the definitions of version 1 (top) and version 2 (bottom) of the String class. The update introduces two incompatible changes: (1) the addition of a new method, `include_pattern`, for checking the existence of a substring that

```

1 class String!1():
2   def __init__(self, value):
3     if type(value) == int:
4       self.value = chr(value)
5     elif type(value) == str:
6       self.value = value
7
8   def get(self, i):
9     if(i < 0):
10      return 0
11     if(len(self.value) <= i):
12      return 0
13     return ord(self.value[i])

```

```

1 class String!2():
2   def __init__(self, value):
3     if type(value) == int:
4       self.value = chr(value)
5     elif type(value) == str:
6       self.value = value
7
8   def get(self, i):
9     if(i < 0):
10      return 0
11     if(len(self.value) <= i):
12      return 0
13     return self.value[i].incomp(String, 2)
14
15   def include_patter(self, regex):
16     ..

```

Fig. 11: Class String before and after update

```

1 ..
2 titles = [String!2("reverse-engineering_weep_ReLU_
3             networks"), String!2("Hi"), String!2("Reverse-
4             engineering_deep_ReLU_networks")]
5 Integrate().main(titles)
6 ..
7 # operation using newly introduced method in version 2 of
8   the String class
9 .. title.include_pattern(r'..') ..

```

Fig. 12: Updated Downstream User's Program

matches a given regex, and (2) a modification to the behavior of the `get` method, which retrieves the character at a specified index. In the original implementation, `self.get(i)` returned the character code of the *i*th character in the string stored in `self`. After the update, it instead returns a single-character string. Notably, this incompatibility is intended to replicate the real-world incompatibility for `String#[]` introduced in Ruby 1.9.0 in the process of multilingualization (Ruby M17N).^{*16 *17}

Downstream User's Program

Figure 12 shows the user program after the gradual update; the part that originally used version 1 of the String class has been selectively updated to use version 2 of the String class, as shown in line 2. In line 3, a list of `String!2` instances is passed to the `main` method of the `Integrate` class, which outputs the KWIC index corresponding to the given list.

^{*16} ruby/doc/NEWS at ruby/ruby·v1.9.0.0, <https://github.com/ruby/ruby/blob/v1.9.0.0/doc/NEWS#L13> (Accessed January 10, 2025)

^{*17} Regarding the behavioral differences related to this incompatibility, please refer to the NEWS (an overview of the ChangeLog file for a new Ruby version) of the subsequent stable release, version 1.9.1. NEWS for Ruby 1.9.1 - Compatibility issues <https://docs.ruby-lang.org/en/3.4/NEWS/NEWS-1.9.1.html> (Accessed January 10, 2025)

7.2 Gradually Updating KWIC in Vython Inconsistent Version Usage in Sort

When we execute the updated user program in Figure 12, a value derived from the `get` method for a `String!2` instance is assigned to the `char` in line 38 in Figure 10. Then, at the execution of line 39, where `char` is passed to the inequality operation, the following message is notified by both Python runtime and DVC.

```
1 TypeError: '<=' not supported between instances of 'int'
  and 'str'
```

Error Message from Python Runtime

```
1 Incompatible version usage found:
2 - String 1
3 - String 2
4 [Changed in version 2] `String().get(i)`:
5 - returns a string whose length is 1 consisting of the i-
  th character.
6 - but returns the character code of the i-th character in
  version 1.
```

The Warning Issued from DVC

Here, Python runtime reports that its inequality operations are performed on a combination of values with different types (`int` and `str`). Specifically, the evaluating value of `String!1("a").get(0)` has a type of `int`, and the evaluating value of `char` has a type of `str`.

DVC reports that its inequality operations are performed on a combination of values created from an incompatible version of the implementation. Specifically, the evaluating value of `String!1("a").get(0)` has a VT of $\{(String, 1, -), \dots\}$ and the evaluating value of `char` has a VT of $\{(String, 2, T), \dots\}$. Additionally, DVC also reports how these values can be incompatible. By combining this information, we identify the cause of the program's failure: the implementation assumes that the return value of the `get` method from the version 1 `String` class (an `int` value) is assigned to a `char`. However, due to the update, the return value of the version 2 `get` method (a `char` value) is being assigned instead.

```
1 class Sort!1:
2 ..
3 def first_shift_to_str(self, shift):
4     keyword = String!1("")
5     for char_i in range(self.rotate.num_chars(shift, 0)):
6         char = self.rotate.get_char(shift, 0, char_i)
7         if String!2("a").get(0) <= char <= String!2("z").
           get(0):
8             char = chr(char - 32)
9             keyword.add(String!1(char))
10    return keyword
```

Refactored Definition of `first_shift_to_str`: 1

This cause analysis leads to a program revision, as exemplified in the above program. We refactor the program by changing the version of `get` method from version 1 to version 2 in the inequality operation. Then, executing the programs with this refactored program, the following results were produced.

```
1 TypeError: unsupported operand type(s) for -: 'str' and
  'int'
```

Error Message from Python Runtime

The Python runtime raises an error at line 40 of Figure 10.

This error reports that an `int` type value and a `str` type value were passed as operands to a subtraction operation. Specifically, the evaluating value of 32 has a type of `int`, and the evaluating value of `char` has a type of `str`.

However, no error is reported by the DVC mechanism in this case. This is because the return value of the version 2 `get` method assigned to `char` and the naive integer value 32 seem to be version consistent.

```
1 class Sort!1:
2 ..
3 def first_shift_to_str(self, shift):
4     keyword = String!1("")
5     for char_i in range(self.rotate.num_chars(shift, 0)):
6         char = self.rotate.get_char(shift, 0, char_i)
7         if String!2("a").get(0) <= char <= String!2("z").get
           (0):
8             char = chr(ord(char) - 32)
9             keyword.add(String!1(char))
10    return keyword
```

Refactored Definition of `first_shift_to_str`: 2

Therefore, in this case, we refactor the program as described above based solely on information provided by the Python runtime, without any feedback from the DVC mechanism. Specifically, we refactored the program to perform the subtraction of character codes using the `ord` method to explicitly obtain the character codes.

Executing the programs with this refactored program, it outputs the correct KWIC index as same as before an update along with countless warnings from Vython. However, these warnings are due to DVC's conservative design, and the results confirm that there is in fact no problem. Thus, we successfully completed the update of the downstream User's program from version 1 of the `String` class to version 2.

7.3 Discussion

Through this case study, we demonstrate that Vython with DVC can support gradual updates. However, several limitations were also observed. For instance, during the execution of the program after the first refactoring, DVC did not issue any warnings about a type error caused by version-related incompatibilities. While such issues fall outside the current scope of DVC, they represent a class of problems that are likely to occur frequently in practice due to version incompatibilities. It will be necessary to extend DVC to address these structural incompatibility issues.

The current implementation of DVC is not able to identify which evaluation step each version information was attached in when inconsistent version information is detected. In addition to reporting inconsistency, it would be more helpful to be able to report what piece of program caused the inconsistency.

Moreover, unlike static compatibility checking such as VL [27] and BatakJava [19], we recognize that DVC is ineffective against refactoring source code outside the execution path. Therefore, it would be challenging to migrate the entire source code following the gradual updating model using DVC alone. Considering that Python also has a lightweight typing mechanism^{*18}, exploring static mechanisms for dynamic languages could be a reason-

^{*18} PEP 484 - Type Hints <https://peps.python.org/pep-0484/> (Accessed January 10, 2025)

able direction for future work.

8. Related Work

8.1 Language-Based Approach Using Multiple Version in a Program

Programming with Versions is a programming paradigm proposed by Tanabe et al. that enables the use of multiple versions of packages, modules, or classes within a program, facilitating gradual updates. They highlight the incompatibility issues that arise when multiple versions of functions or classes coexist in a program and emphasize the importance of handling version consistency [27], [28] and structural compatibility [19] in their languages.

λ_{VL} [26], [27] is a functional core calculus with a type system that ensures version consistency. Its implementation, VL [28], enables the automatic selection of function versions using consistency-based type inference. On the other hand, Batak-Java [19] is an extension of Featherweight Java [14] with a type system that ensures structural compatibility. Through Batak-Java's type inference, the version of the class used for all instance creations is automatically determined. Programs that pass Batak-Java's type inference are guaranteed to avoid method call or field access failures caused by version incompatibilities. Vythion can be viewed as a dynamic checking mechanism that ensures version consistency. Exploring language extensions to guarantee structural compatibility presents an interesting direction for future work.

Carvalho and Seco propose a mechanism that allows version updates of programs to be expressed within the programming language itself, from the perspective of software evolution. Versioned Featherweight Java [5], [7] is an extension of Featherweight Java that introduces multi-branching and merge operations. Using program slicing, it extracts a well-typed single-version Featherweight Java code from a version-controlled code-base at compile time. They have also implemented a similar approach for Python [6], reducing the effort required for refactoring by automatically inserting some compatibility-absorbing code at compile time. While this approach cannot execute programs involving incompatible versions without pre-defined conversions, Vythion allows programmers to run programs with any version combination and has a mechanism that dynamically detects issues arising from incompatibilities.

8.2 Dynamic Analysis Tools With an Implementation Similar to Vythion

Some offline tools performing dynamic checks exhibit implementation methods that are partially similar to DVC. TaintCheck [23] is a dynamic analysis tool built on Valgrind [22], designed to detect vulnerabilities such as buffer overflows and format string exploits. TaintCheck marks input data from untrusted sources as tainted, tracks the propagation of tainted attributes during program execution (i.e., identifying which other data becomes tainted), and detects instances where tainted data is used in unsafe operations.

Similar approaches are found in the literature on dynamic information flow analysis. Austin and Flanagan [2], [3] introduced

the semantics of λ_{info} [1], a variant of lambda calculus designed to derive the evaluation rules for Featherweight JavaScript, a subset of JavaScript. This calculus assigns security labels to values and dynamically verifies these labels to prevent the leakage of sensitive information by malicious JavaScript programs.

While these tools differ in purpose from Vythion, they share a key characteristic: attaching metadata to values and defining how this metadata propagates. In Vythion, version tables are used in place of security labels. These tools have been reported to incur runtime overheads ranging from several times (as seen in λ_{info}) to as much as 25 times (as observed in TaintCheck) compared to implementations without additional checks. Although the level of optimization varies significantly, incorporating techniques such as sparse labeling—used in λ_{info} to minimize the addition of check metadata—could potentially enhance the performance of Vythion.

8.3 Dynamic Software Updating

Dynamic Software Updating (DSU) [12], [15], [21] is a technique that enables software to be updated while it is running. In DSU, different versions of programs may be executed consecutively, and objects created by the old version may be passed to the new one. This scenario resembles the Programming With Versions paradigm, where multiple program versions are executed simultaneously. To prevent destructive behavior arising from incompatibilities between different versions, existing DSU systems employ various strategies. For instance, some systems restrict update timing (e.g. updating functions only when they are inactive) [12] and perform state transformations on objects created by the old program to make them compatible with the new program [21].

Although DSU and DVC serve different purposes, both require techniques to ensure that compatible versions are used consistently during execution. This overlap suggests that techniques from DSU and DVC could be mutually beneficial. For instance, integrating DVC into DSU might allow for the early detection of harmful dynamic updates, thereby improving overall system stability.

9. Conclusion and Future Work

We implement Vythion and conduct a preliminary evaluation. The results indicate that while the current implementation is prototypical, its performance is scalable and acceptable for debugging small programs but requires further optimization for real-world applications. We plan to undertake the following future work.

Compatibility Management Toward Better Feedback

We plan to extend the DVC mechanism to accommodate the diverse compatibility requirements of real-world software. Practical software packages often have numerous versions and evolve non-linearly [8] through experimental features and language extensions. Currently, however, the DVC mechanism assumes a linear evolution involving only two versions of classes. The sole incompatibility annotation available to upstream developers, `incomp`, simply indicates that a version is incompatible with all

others. This limitation it impossible to express incompatibilities between specific versions in scenarios involving three or more versions, as commonly encountered in real-world development. Leveraging tools to manage source code differences and incompatibilities, it becomes possible to synthesize feedback that accounts for the history of updates.

Surface Language Design

The current Vythion requires specifying class versions in the surface program. We plan to develop a method to automatically infer versions working on Python programs. Specifically, we are considering a mechanism to explore a version combination that can be executed successfully by automatically iterating the cycle of switching a version of each expression and executing it with Vythion. This will help minimize the annotations given by downstream developers, identify dependencies on old versions, and automate updates.

Acknowledgments The authors would like to thank the members of the PRG lab for their daily discussions. The authors also thank the reviewers of APLAS 2024 SRC, the audiences of APLAS & ATVA 2024 and PPL 2024 for their valuable comments on the preliminary stages of this research. Additionally, the authors are grateful to CRuby committers for their efforts in investigating two-decades-ago incompatibilities in the case study. This work was supported by JSPS KAKENHI Grant Numbers JP23K19961 and JP23K28058.

References

- [1] Austin, T. H.: Dynamic Information Flow Analysis for Featherweight JavaScript Technical Report # UCSC-SOE-11-19, (online), available from <https://api.semanticscholar.org/CorpusID:16308311> (2011).
- [2] Austin, T. H. and Flanagan, C.: Efficient purely-dynamic information flow analysis, *Proceedings of the ACM SIGPLAN Fourth Workshop on Programming Languages and Analysis for Security*, PLAS '09, New York, NY, USA, Association for Computing Machinery, p. 113–124 (online), DOI: 10.1145/1554339.1554353 (2009).
- [3] Austin, T. H. and Flanagan, C.: Permissive dynamic information flow analysis, *Proceedings of the 5th ACM SIGPLAN Workshop on Programming Languages and Analysis for Security*, PLAS '10, New York, NY, USA, Association for Computing Machinery, (online), DOI: 10.1145/1814217.1814220 (2010).
- [4] Bradski, G.: The OpenCV Library, *Dr. Dobbs' Journal of Software Tools* (2000).
- [5] Carvalho, L. and Costa Seco, J. a.: Deep Semantic Versioning for Evolution and Variability, *Proceedings of the 23rd International Symposium on Principles and Practice of Declarative Programming*, PPDP '21, New York, NY, USA, Association for Computing Machinery, (online), DOI: 10.1145/3479394.3479416 (2021).
- [6] Carvalho, L. and Costa Seco, J. a.: A Language-Based Version Control System for Python, *38th European Conference on Object-Oriented Programming (ECOOP 2024)* (Aldrich, J. and Salvaneschi, G., eds.), Leibniz International Proceedings in Informatics (LIPIcs), Vol. 313, Dagstuhl, Germany, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, pp. 9:1–9:27 (online), DOI: 10.4230/LIPIcs.ECOOP.2024.9 (2024).
- [7] Carvalho, L. and Costa Seco, J.: Software Evolution with a Typeful Version Control System, *Software Engineering and Formal Methods* (Ölveczky, P. C. and Salaün, G., eds.), Cham, Springer International Publishing, pp. 145–161 (2019).
- [8] Conradi, R. and Westfechtel, B.: Version models for software configuration management, *ACM Comput. Surv.*, Vol. 30, No. 2, p. 232–282 (online), DOI: 10.1145/280277.280280 (1998).
- [9] Eghbali, A. and Pradel, M.: DynaPyT: a dynamic analysis framework for Python, *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ESEC/FSE 2022, New York, NY, USA, Association for Computing Machinery, p. 760–771 (online), DOI: 10.1145/3540250.3549126 (2022).
- [10] Foo, D., Chua, H., Yeo, J., Ang, M. Y. and Sharma, A.: Efficient static checking of library updates, *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, ESEC/FSE 2018, New York, NY, USA, Association for Computing Machinery, p. 791–796 (online), DOI: 10.1145/3236024.3275535 (2018).
- [11] Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del R'io, J. F., Wiebe, M., Peterson, P., G'érard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C. and Oliphant, T. E.: Array programming with NumPy, *Nature*, Vol. 585, No. 7825, pp. 357–362 (online), DOI: 10.1038/s41586-020-2649-2 (2020).
- [12] Hayden, C. M., Smith, E. K., Hardisty, E. A., Hicks, M. and Foster, J. S.: Evaluating Dynamic Software Update Safety Using Systematic Testing, *IEEE Transactions on Software Engineering*, Vol. 38, No. 6, pp. 1340–1354 (online), DOI: 10.1109/TSE.2011.101 (2012). Conference Name: IEEE Transactions on Software Engineering.
- [13] Igarashi, A., Pierce, B. C. and Wadler, P.: Featherweight Java: a minimal core calculus for Java and GJ, *ACM Trans. Program. Lang. Syst.*, Vol. 23, No. 3, p. 396–450 (online), DOI: 10.1145/503502.503505 (2001).
- [14] Igarashi, A., Pierce, B. C. and Wadler, P.: Featherweight Java: a minimal core calculus for Java and GJ, *ACM Transactions on Programming Languages and Systems*, Vol. 23, No. 3, pp. 396–450 (online), DOI: 10.1145/503502.503505 (2001).
- [15] Islam, C., Prokhorenko, V. and Babar, M. A.: Runtime software patching: Taxonomy, survey and future directions, *Journal of Systems and Software*, Vol. 200, p. 111652 (online), DOI: 10.1016/j.jss.2023.111652 (2023).
- [16] Jayasuriya, D., Terragni, V., Dietrich, J. and Blincoe, K.: Understanding the Impact of APIs Behavioral Breaking Changes on Client Applications, *Proc. ACM Softw. Eng.*, Vol. 1, No. FSE (online), DOI: 10.1145/3643782 (2024).
- [17] Kula, R. G., German, D. M., Ouni, A., Ishio, T. and Inoue, K.: Do developers update their library dependencies?, *Empirical Software Engineering*, Vol. 23, No. 1, pp. 384–417 (online), DOI: 10.1007/s10664-017-9521-5 (2018).
- [18] Lam, P., Dietrich, J. and Pearce, D. J.: Putting the semantics into semantic versioning, *Proceedings of the 2020 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*, Onward! 2020, New York, NY, USA, Association for Computing Machinery, p. 157–179 (online), DOI: 10.1145/3426428.3426922 (2020).
- [19] Lubis, L. A., Tanabe, Y., Aotani, T. and Masuhara, H.: BatakJava: An Object-Oriented Programming Language with Versions, *Proceedings of the 15th ACM SIGPLAN International Conference on Software Language Engineering*, Auckland New Zealand, ACM, pp. 222–234 (online), DOI: 10.1145/3567512.3567531 (2022).
- [20] Mirhosseini, S. and Parmin, C.: Can automated pull requests encourage software developers to upgrade out-of-date dependencies?, *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering*, ASE '17, Urbana-Champaign, IL, USA, IEEE Press, pp. 84–94 (2017).
- [21] Neamtiu, I., Hicks, M., Stoyke, G. and Oriol, M.: Practical dynamic software updating for C, *SIGPLAN Not.*, Vol. 41, No. 6, pp. 72–83 (online), DOI: 10.1145/1133255.1133991 (2006).
- [22] Nethercote, N. and Seward, J.: Valgrind: a framework for heavyweight dynamic binary instrumentation, *SIGPLAN Not.*, Vol. 42, No. 6, p. 89–100 (online), DOI: 10.1145/1273442.1250746 (2007).
- [23] Newsome, J. and Song, D. X.: Dynamic Taint Analysis for Automatic Detection, Analysis, and Signature Generation of Exploits on Commodity Software, *Network and Distributed System Security Symposium*, (online), available from <https://valgrind.org/docs/newsome2005.pdf> (2005).
- [24] Parnas, D. L.: On the criteria to be used in decomposing systems into modules, *Commun. ACM*, Vol. 15, No. 12, p. 1053–1058 (online), DOI: 10.1145/361598.361623 (1972).
- [25] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J. and Chintala, S.: PyTorch: An Imperative Style, High-Performance Deep Learning Library, *Advances in Neural Information Processing Systems 32*, Curran Associates, Inc., pp. 8024–8035 (online), available from <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf> (2019).
- [26] Tanabe, Y., Aotani, T. and Masuhara, H.: A Context-Oriented Programming Approach to Dependency Hell, *Proceedings of the 10th ACM International Workshop on Context-Oriented Programming: Ad-*

- vanced Modularity for Run-Time Composition, COP '18, New York, NY, USA, Association for Computing Machinery, p. 8–14 (online), DOI: 10.1145/3242921.3242923 (2018).
- [27] Tanabe, Y., Lubis, L. A., Aotani, T. and Masuhara, H.: A Functional Programming Language with Versions, *The Art, Science, and Engineering of Programming*, Vol. 6, No. 1, pp. 5:1–5:30 (online), DOI: 10.22152/programming-journal.org/2022/6/5 (2021).
- [28] Tanabe, Y., Lubis, L. A., Aotani, T. and Masuhara, H.: Compilation Semantics for a Programming Language with Versions, *Programming Languages and Systems* (Hur, C.-K., ed.), Singapore, Springer Nature Singapore, pp. 3–23 (2023).
- [29] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P. and SciPy 1.0 Contributors: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, *Nature Methods*, Vol. 17, pp. 261–272 (online), DOI: 10.1038/s41592-019-0686-2 (2020).
- [30] Wes McKinney: Data Structures for Statistical Computing in Python, *Proceedings of the 9th Python in Science Conference* (Stéfan van der Walt and Jarrod Millman, eds.), pp. 56 – 61 (online), DOI: 10.25080/Majora-92bf1922-00a (2010).

Appendix

A.1 Incompatible Behaviours Between NumPy 2.0.0 and 1.26.4

This section outlines some of the major incompatibilities between NumPy 1.26.4 and NumPy 2.0.0. All 11 examples we collect, along with the scripts to reproduce them, are available on the GitHub repository (https://github.com/prg-titech/numpy_diff).

A.1.1 Incompatibilities in `numpy.linalg.solve`

The `numpy.linalg.solve` function solves a linear matrix equation. It solves the equation $ax = b$ for x , where a is a square matrix and b is a vector or matrix provided as arguments to the function. The implementation was changed in NumPy 2.0.0. The following note is from the official NumPy documentation:

Changed in version 2.0: The b array is only treated as a shape $(M,)$ column vector if it is exactly 1-dimensional.

In all other instances it is treated as a stack of (M, K) matrices. Previously b would be treated as a stack of $(M,)$ vectors if $b.ndim$ was equal to $a.ndim - 1$.

As a result, when b is not strictly one-dimensional, the output of the `solve` function differs for the same input. For example, consider the following program run with NumPy 1.26.4 and NumPy 2.0.0.

```
1 import numpy as np
2
3 # Shape (2, 2, 2)
4 a = np.array(
5     [[ [3, 1], [1, 2] ]
6      , [[2, 1], [1, 3] ] ])
7 # Shape (2, 2)
8 b = np.array(
9     [ [9, 8]
10      , [7, 10] ])
11
12 x = np.linalg.solve(a, b)
13 print(x)
```

When we run the above program with NumPy versions 1.26.4 and 2.0.0, we get the following different outputs due to incompatibility in broadcasting rules.

```
1 @ Running linalg_solve.py with numpy 1.26.4
2 [[2. 3. ]
3  [2.2 2.6]]
4 @ Running linalg_solve.py with numpy 2.0.0
5 [[[2.2 1.2]
6   [2.4 4.4]]
7
8   [[4. 2.8]
9   [1. 2.4]]]
```

The reason for this difference lies in how the b array is treated in different versions of NumPy. In version 1.26.4, if b 's number of dimensions ($b.ndim$) is equal to one less than the number of dimensions of a ($a.ndim - 1$), b is interpreted as a stack of $(M,)$ vectors. This means that in version 1.26.4, the b array is treated as a stack of 1-dimensional vectors, each corresponding to a 2×2 matrix in a . Therefore, the program is interpreted as follows:

$$\begin{pmatrix} 3 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 9 \\ 8 \end{pmatrix},$$

$$\begin{pmatrix} 2 & 1 \\ 1 & 3 \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \end{pmatrix} = \begin{pmatrix} 7 \\ 10 \end{pmatrix}.$$

However, in version 2.0.0, the behavior was modified such that the b array is treated as a column vector only if it is strictly 1-dimensional. In all other cases, it is treated as a stack of (M, K) matrices. Consequently, for the given input, b is treated as a stack of 2-dimensional matrices. Therefore, the program is interpreted as follows:

$$\begin{pmatrix} 3 & 1 \\ 1 & 2 \end{pmatrix} \begin{pmatrix} x_1 & x_2 \\ x_3 & x_4 \end{pmatrix} = \begin{pmatrix} 9 & 8 \\ 7 & 10 \end{pmatrix},$$

$$\begin{pmatrix} 2 & 1 \\ 1 & 3 \end{pmatrix} \begin{pmatrix} y_1 & y_2 \\ y_3 & y_4 \end{pmatrix} = \begin{pmatrix} 9 & 8 \\ 7 & 10 \end{pmatrix}.$$

A.1.2 Incompatibilities in Other Functions

In addition to `numpy.linalg.solve`, NumPy 2.0.0 introduces several other backward-incompatible modifications. Among the programs we collected that produce different outputs solely due to version differences in NumPy, we list some notable input-output pairs below. For other examples where downstream developers might easily notice incompatibilities due to Python runtime errors, such as differences in output types, please refer to the repository.

`numpy.nonzero`

The function return the indices of the elements that are non-zero. The function previously ignored whitespace so that a string only containing whitespace was considered `False`, however, whitespace is now considered `True` in string arrays newly in NumPy 2.0.0.

```
1 import numpy as np
2
3 arr = np.array([' ', 'a', ' '])
4 print(np.nonzero(arr))
```

```
1 @ Running nonzero.py with numpy 1.26.4
2 (array([1]),)
3 @ Running nonzero.py with numpy 2.0.0
4 (array([0, 1]),)
```


numpy.linalg.lstsq

The function returns the least squares solution to a linear matrix equation. The default value of the `rcond` (cut-off ratio) parameter in `lstsq` was changed in NumPy 2.0.0. This change introduces a subtle incompatibility: while most inputs yield the same output regardless of the NumPy version, inputs with elements near machine precision can produce different results depending on the NumPy version. The following example illustrates such a case.

```
1 import numpy as np
2
3 a = np.zeros((10**2, 2))
4 a[0, 0] = 1
5 a[m-1, 1] = 2.22e-16
6 b = np.zeros(m)
7 b[m-1] = 1
8
9 x, res, rank, s = np.linalg.lstsq(a, b)
10 print(...)
```

```
1 @ Running linalg_lstsq.py with numpy 1.26.4
2 Solution with default rcond: [0.00000000e+00 4.5045045e
+15]
3 Residuals: [4.93038066e-32]
4 Rank: 2
5 Singular values: [1.00e+00 2.22e-16]
6 @ Running linalg_lstsq.py with numpy 2.0.0
7 Solution with default rcond: [0. 0.]
8 Residuals: []
9 Rank: 1
10 Singular values: [1.00e+00 2.22e-16]
```

numpy.loadtxt and numpy.genfromtxt

The functions provide readers for simply formatted files. Default encoding for these functions was changed in NumPy 2.0.0. Previously, these two functions selected `encoding=bytes` as the default parameter, but starting from version 2.0.0, it has been changed to `encoding=string`. As a result, programs that expect custom converters assuming a byte value will be broken by the update.

```
1 import numpy as np
2 import io
3 def custom_converter(byte_string):
4     return float(byte_string.decode('utf-8'))
5
6 data = b"1.1\n2.2\n3.3\n"
7 with open('data.txt', 'wb') as f:
8     f.write(data)
9
10 # Load the data using loadtxt with the custom converter
11 try:
12     data = np.loadtxt('data.txt', converters={0:
13         custom_converter})
14     print(f"Data loaded successfully: {data}")
15 except Exception as e:
16     print(f"An error occurred: {e}")
```

```
1 @ Running loadtxt_genfromtxt.py with numpy 1.26.4
2 Data loaded successfully: [1.1 2.2 3.3]
3 @ Running loadtxt_genfromtxt.py with numpy 2.0.0
4 An error occurred: could not convert string '1.1' to
float64 at row 0, column 1.
```

A.2 Dynamically Switching NumPy Versions

This section describes the reproduction of the motivating examples from Section 2 in actual Python programs. The complete source code and instructions to reproduce the results of this paper are available on the GitHub repository (<https://github.com/prg-titech/use-multi-versions>).

A.2.1 Installing Multiple NumPy Versions from Sources

For example, to install numpy version 1.26.4 into a directory named `numpy-1.26.4` using pip on a Linux OS, use the following command.

```
1 $ mkdir numpy-1.26.4
2 $ pip download numpy==1.26.4
3 $ pip install numpy-1.26.4- ... .whl -t numpy-1.26.4
```

A.2.2 Simultaneously Using Multiple NumPy Versions in Code

The following `load_numpy` function dynamically loads a specified version of NumPy. It takes a string representing the version, sets the appropriate NumPy path, and removes any cached instances of NumPy from `sys.modules`. The function then temporarily modifies the system path to include the specified version's path installed in the last section and imports the NumPy module from its initialization file. Finally, `load_numpy` returns the module object for the specified version of NumPy.

```
1 # version_dispatch.py
2 def load_numpy(version):
3     if version == '1.26.4':
4         numpy_path = os.path.abspath('numpy-1.26.4')
5     elif version == '2.0.0':
6         numpy_path = os.path.abspath('numpy-2.0.0')
7     else:
8         raise ValueError(f"Unsupported numpy version: {version}")
9
10    # Clear cache
11    if 'numpy' in sys.modules:
12        del sys.modules['numpy']
13    for mod_name in list(sys.modules):
14        if mod_name.startswith('numpy'):
15            del sys.modules[mod_name]
16
17    # Set environment paths
18    original_path = sys.path.copy()
19    sys.path.insert(0, numpy_path)
20
21    try:
22        numpy_init_path = os.path.join(numpy_path, 'numpy'
23            , '__init__.py')
24        spec = importlib.util.spec_from_file_location("
25            numpy", numpy_init_path)
26        if spec is None:
27            raise ImportError(f"Cannot find numpy module in
28                {numpy_path}")
29        numpy = importlib.util.module_from_spec(spec)
30        spec.loader.exec_module(numpy)
31    finally:
32        # Restore sys.path
33        sys.path = original_path
```

The following program shows the full version of the program shown in Figure 2. The implementation of the pole placement problem (`place_poles` and `my_place_poles`) has been simplified, as it is not the focus of this section. Using `./version_dispatch.py`, which defines the `load_numpy` function described in the previous subsection, `place_poles` is evaluated with NumPy 1.26.4 on line 26, and `my_place_poles` is evaluated with NumPy 2.0.0 on line 27. Finally, the results of the two functions are compared on line 29.

As mentioned in Section 2, despite `place_poles` and `my_place_poles` being identical implementations except for the NumPy version they use, the result evaluates to `False`.

```
1 from version_dispatch import load_numpy
2
3 # SciPy
```

```

4 class SciPy():
5     def place_poles(self, A, B, desired_poles):
6         np = load_numpy('1.26.4')
7         res = np.linalg.solve(A, B)
8         return res
9
10 # User Program
11 def my_place_poles(A, B, desired_poles):
12     np = load_numpy('2.0.0')
13     res = np.linalg.solve(A, B)
14     return res
15
16 def main():
17     np = load_numpy('2.0.0')
18     A = np.array(
19         [ [3, 1], [1, 2]]
20         , [ [2, 1], [1, 3]] ] )
21     B = np.array(
22         [ [9, 8]
23         , [7, 10]] )
24     desired_poles = np.array([-1.0, -2.0])
25
26     expect = SciPy().place_poles(A,B,desired_poles).tolist()
27     actual = my_place_poles(A,B,desired_poles).tolist()
28
29     test = np.array_equal(expect, actual) # => False
30
31 main()

```

A.3 Programs Used for Simple Benchmarks

sort performs a merge sort on a Python list of 1000 elements, is_prime uses a simple algorithm to determine the primality of 128456903, fib recursively computes the 20th Fibonacci number, sieve finds all prime numbers up to 5000 using "Sieve of Eratosthenes" algorithm, and insert inserts one thousand Node instances to a binary tree.

A.3.1 insert

```

1 class Node!1():
2     def __init__(self, value):
3         self.value = value
4         self.left = None
5         self.right = None
6
7     def insert_right(self, v):
8         if self.right == None:
9             self.right = Node!1(v)
10        else:
11            self.right.insert(v)
12
13    def insert_left(self, v):
14        if self.left == None:
15            self.left = Node!1(v)
16        else:
17            self.left.insert(v)
18
19    def insert(self, v):
20        if(self.value <= v):
21            self.insert_right(v)
22        else:
23            self.insert_left(v)
24
25    root = Node!1(5)
26    a = [...] # Array of 1000 elements, random numbers
27           # between 1 and 10000
28    for i in a:
29        root.insert(i)

```

A.3.2 sort

```

1 def sort(list):
2     if len(list) < 1:
3         return []
4     elif len(list) == 1:
5         return list
6     pivot = list[0]
7     lower_list = []
8     upper_list = []
9     middle_list = []
10

```

```

11     for item in list:
12         if item < pivot:
13             lower_list.append(item)
14         elif item > pivot:
15             upper_list.append(item)
16         else:
17             middle_list.append(item)
18
19     sorted_lower_list = sort(lower_list)
20     sorted_upper_list = sort(upper_list)
21
22     return sorted_lower_list + middle_list +
23           sorted_upper_list
24
25 a = [...] # Array of 1000 elements, random numbers
26           # between 1 and 10000
27 sort(a)

```

A.3.3 prime

```

1 def is_prime(n):
2     if n <= 1:
3         return False
4     if n == 2 or n == 3:
5         return True
6     if n % 2 == 0 or n % 3 == 0:
7         return False
8     return is_prime_recursive(n, 5)
9
10 def is_prime_recursive(n, i):
11     if i * i > n:
12         return True
13     if n % i == 0 or n % (i + 2) == 0:
14         return False
15     return is_prime_recursive(n, i + 6)
16
17 is_prime(128456903)

```

A.3.4 fib

```

1 def fib(n):
2     if n<=2:
3         return 1
4     else:
5         return fib(n-1) + fib(n-2)
6
7 fib(20)


```

A.3.5 sieve


```

1 def sieve(flags, size):
2     prime_count = 0
3
4     i = 2
5     while i < size + 1:
6         if flags[i - 1]:
7             prime_count = prime_count + 1
8             k = 2 * i
9             while k <= size:
10                 flags[k - 1] = False
11                 k = k + i
12             i = i + 1
13
14     return prime_count
15
16 flags = [True, ...]
17 sieve(flags, 5000)

```




Satsuki Kasuya He is a graduate student in Department of Mathematical and Computing Science, School of Computing, Institute of Science Tokyo (formerly Tokyo Institute of Technology). He received B.S. degree from Tokyo Institute of Technology in 2024. He is interested in programming languages.



Yudai Tanabe He is an assistant professor in Department of Mathematical and Computing Science, School of Computing, Institute of Science Tokyo (formerly Tokyo Institute of Technology.) He received B.S., M.S., and D.S. degrees from Tokyo Institute of Technology in 2018, 2020, and 2023 respectively. Formerly, he

was a JSPS Research Fellow at Tokyo Institute of Technology until 2023, a program-specific researcher at Kyoto University until 2024, and an assistant professor in Department of Mathematical and Computing Science, School of Computing, Tokyo Institute of Technology until September 2024. His research interest is programming language and software engineering, especially in programming language theory, type systems, and software maintenance.



Hidehiko Masuhara He is a Professor of Mathematical and Computing Science, School of Computing, Institute of Science Tokyo (formerly Tokyo Institute of Technology.) He received his B.S., M.S., and Ph.D. in Computer Science from Department of Information Science, the University of Tokyo in 1992, 1994, and 1999, respectively, and served as an assistant professor, lecturer, and as-

sociate professor at Department of Graphics and Computer Science, Graduate School of Arts and Sciences, the University of Tokyo from 1995 until 2013. He was the Dean of the School of Computing of Tokyo Institute of Technology from 2022 until 2024. His research interest is programming languages, especially in aspect- and context-oriented programming, partial evaluation, computational reflection, meta-level architectures, parallel/concurrent computing, and programming environments.