

Faster than C: Static Type Inference with Starkiller

Michael Salib

March 26, 2004

Abstract

Pure Python code is slow, primarily due to the dynamic nature of the language. I have begun building a compiler to produce fast native code from Python source programs. While compilation can improve performance, any such gains will be modest unless the compiler can statically resolve most of the dynamism present in source programs. Static type inference for Python programs would enable the safe removal of most type checks and most instances of dynamic dispatch and dynamic binding from the generated code. Removing dynamic dispatch and binding leads to large performance benefits since their existence precludes many traditional optimization techniques, such as inlining. I have built a static type inferencer for Python called Starkiller. Given a Python source file, it can deduce the types of all expressions in the program without actually running it. Starkiller's primary goal is to take a Python source program as input and deduce the information needed to make native code compilation of that program easy. This paper describes Starkiller's design and operation. It is partially adapted from a draft of my Master's thesis.

Contents

1	Executive Summary	3
2	Introduction	4
3	Core TI Algorithm	5
3.1	Nodes and Constraints	7
3.2	Functions	8
3.3	Classes and Objects	10
4	Advanced Language Features	12
4.1	Operators	12
4.2	Exceptions	13
4.3	Iterators and Generators	14
4.4	Modules	15

5	Foreign Code Interactions	15
5.1	External Type Descriptions	15
5.2	The builtins Module Type Description	16
6	Known Problems and Limits	18
6.1	Eval, Exec, and Dynamic Module Loading	18
6.2	Static Error Detection	23
6.3	Partial Evaluation	24
6.4	Integer Promotion	25

1 Executive Summary

Pure Python code is slow. Most of the time, that doesn't matter because application performance generally doesn't depend on pure Python very much. Nevertheless, there are situations where it does matter, so I've spent time trying to discover how to improve the performance of pure Python code. I believe that static compilation is the answer.

The greatest obstacle to efficient compilation and high performance is Python's dynamic nature: every operation must verify that its operands are of the correct type, every access requires going through layers of indirection, etc. The combination of dynamic binding and dynamic typing is particularly problematic for optimization since it postpones all dispatch decisions until runtime, thereby eliminating most benefits of static compilation.

This dynamism makes programming in Python a joy, but generating optimal code a nightmare. Yet while the presence of such abundant dynamism makes traditional static optimization impossible, in most programs, there is surprisingly little dynamism present. For example, in most Python programs:

1. all class and function objects are defined exactly once
2. class inheritance relationships do not change at run time
3. methods are not added after a class object has been created
4. most expressions have exactly one type; the vast majority of those that have more than one type have only a few types

The flip side is that what little dynamism a particular program makes use of is often absolutely vital.

I have developed a type inference algorithm for Python that is able to resolve most dispatches statically. This algorithm is based on Ole Agesen's Cartesian Product Algorithm for type inference of Self programs. I have built a type inferencer for Python based on this algorithm called Starkiller.

Starkiller's type inference algorithm works by constructing a dataflow network for types that models the runtime behavior of values in an input program. Of course, since Starkiller only has access to information available at compile time, it must make conservative approximations. This means that Starkiller will occasionally infer overly broad type sets for an expression, but it also means that it will never fail to infer the existence of a type that appears at runtime.

Each variable or expression in the input program is modelled by a node that contains a set of possible types that expression can achieve at runtime. Initially, all sets are empty, except for constant expressions, which are initialized to contain the appropriate constant type. Nodes are in turn connected with constraints that model data flow between them. A constraint is a unidirectional link that forces the receiving node's type set to always have at least the elements of the sender node's type set. In other words, types flow along constraints.

To avoid both loss of precision from merging types in function bodies and an exponential growth in analysis time required, Starkiller takes the cartesian product of the list of argument

type sets seen by any function call. The result is a list of argument lists where each argument list contains exactly one type for each argument. Starkiller performs analysis on function bodies for each of these monomorphic argument lists, caching the results so they can be used at other call sites. This ensures that polymorphic function calls do not pollute the analysis of functions called later in the call chain.

In much the same way, Starkiller deals with class instances by maintaining polymorphic state describing the instance's attribute types at the instance construction site. Taking the cartesian product of this state allows us to generate monomorphic instance states that flow from a instance definition node. These monomorphic instance states are immutable; when they encounter a `setattr` instruction, they do not change their own state, but pass the request back to their polymorphic parent who then changes its state and generates new monomorphic instance state types as appropriate.

The key idea behind Starkiller is that it never removes or retracts type information. At any point in time, an expression's node's type set accurately, but perhaps incompletely, describes the types that expression will observe at run time. This monotonicity ensures that Starkiller will eventually converge on a final type state of the entire program.

The last piece of the puzzle is describing how external types interact with the type system. Starkiller provides a mini-language that allows extension type authors to describe the type behavior of their extensions. These descriptions are used during inference to model the run time dataflow. Since most extensions are quite simple from a type perspective, these descriptions tend to be short and simple. For example, the `str` type constructor takes any object and returns a string instance (possibly calling the object's `__str__` method if it exists).

2 Introduction

Starkiller's TI algorithm is designed to take as input a Python source program and generate as output a declaration of what set of types each expression in that source program will achieve at runtime. There are two unexpected features of Starkiller's output that bear mentioning. The first is that each function and method in the declaration will be repeated such that each copy is specialized for a particular sequence of monomorphic argument types. The second feature is that a single class may produce more than one class instance type after construction. These features are artifacts of the manner in which Starkiller deals with parametric and data polymorphism respectively. Nevertheless, they are useful artifacts since they naturally suggest particularly efficient compilation strategies.

From a type inference perspective, Python is a large and complex language. In contrast to other languages that rely heavily on type inference for performance, such as Eiffel, Haskell, or the many variants of ML, Python was designed with little thought as to how the language semantics would hinder or help type inference. Instead, Python's semantics evolved over several years in response to feedback from a community of skilled practitioners. Thus, while languages like Haskell suffer occasional design flaws that had to be imposed specifically to make type inference easier, Python makes no such compromises in its design, which

only makes Starkiller’s job that much harder. One example of type inference dictating language development is how Haskell forbids the polymorphic use of functions passed as arguments because the Hindley-Milner type inference algorithm that Haskell relies upon cannot handle such nonlocal usage. This limitation stems directly from Hindley-Milner’s focus on incremental inference rather than whole program analysis.

While Starkiller’s TI algorithm can analyze a large subset of the Python language, there are some language constructs that it cannot presently handle. Programs that use these constructs are rejected by the inferencer. Most unhandled constructs have a known algorithm for type inference but have not been implemented due to lack of time. There are, however, several language features that remain unimplemented because we do not know how to perform type inference in their presence. These features introduce new code into the system at runtime and thus necessarily render static type inference impossible. These features are dynamic module loading, and the functions `eval` and `exec`. In practice, most programs that would benefit from static type inference and compilation do not make use of these features, so their exclusion from Starkiller is of little concern. Nevertheless, we close by discussing techniques that may one day allow Starkiller to support these features.

Like Agesen [1], we necessarily take an operational approach in describing Starkiller’s TI algorithm. We thus forego mathematically intensive descriptions of the denotational or operational semantics in favor of simpler step-by-step descriptions of what the TI algorithm does.

To ease exposition, we begin by describing Starkiller’s TI algorithm for the most basic language features: assignments, function definitions and calls, and the object system. We then examine the TI algorithm as it relates to more advanced but peripheral features of Python, such as generators and exception handling. Having completed a thorough description of how Starkiller’s TI algorithm works for pure Python code, we then turn our attention to Starkiller’s foreign code support and review how Starkiller integrates the analysis of foreign code extensions into its analytic framework. We explore the subject in some detail by examining how Starkiller handles the built-in functions and types. Finally, we review the most serious of Starkiller’s problems and limitations.

3 Core TI Algorithm

The core idea behind Starkiller’s type inference algorithm is to find concrete types in the program source and then propagate them through a dataflow network that models the dynamics of runtime data transfer. Concrete types are those types that are actually created at runtime, as opposed to an abstract type that is either never created (such as an abstract base class) at runtime or that is a model for any number of types (as in Hindley-Milner type inference). Because Starkiller deals exclusively in concrete types, it must have access to the entire input program source at once. In other words, it must perform whole program analysis rather than incremental one module at a time analysis.

Traditionally, whole program analysis techniques have been looked down upon compared to their incremental cousins. However, as the modern computing environment has evolved,

many of the reason for favoring incremental analysis techniques have evaporated. In part due to the open source movement, we live in a world where source code is universally available. Even large commercial libraries ship as source code. The vast majority of applications that benefit from static compilation and type inference have all their source code available for analysis at compile time. For those limited cases where programs must integrate with pre-compiled libraries, the mechanism Starkiller uses for interfacing with foreign code (see Section 5) can be readily adapted.

Starkiller’s type inference algorithm is flow-insensitive, meaning that it assigns types to variables in the input program source. In contrast, flow-sensitive algorithms assign types to individual variable accesses. In theory, flow-insensitive algorithms are less precise than their flow-sensitive counterparts. However, they are also much simpler to implement. To see how flow-insensitivity can hamper precision, consider the code in Figure 1. The variable `num` is assigned an integer value and then later a floating point value. Flow-insensitive algorithms like the one Starkiller uses combine both of these writes together; consequently both reads of `num` in the calls to `doSomething` see `num` as having a type set of integer, float. A flow-sensitive algorithm would have been smart enough to segregate the two writes so that the first call to `doSomething` would see `num` as having a typeset of integer while the second call would see `num` as having a type set of float. The increased precision has obvious implications for generating faster code.

```
num = 4 # an integer
doSomething(num)
num = 3.14159 # a float
doSomething(num)
```

Figure 1: The trouble with flow-insensitive analysis

Starkiller could be augmented to perform flow-sensitive analysis by converting input source programs into Single Static Assignment (SSA) form. In this form, variables are systematically renamed so that each read corresponds to exactly one write. Another approach would involve calculating the set of reachable definitions at every variable read point in the program. Starkiller does not use either of these approaches because they only solve trivial cases like the example shown in Figure 1 while being unable to handle more serious sensitivity precision losses. These algorithms are unsuitable for handling anything more complex than multiple writes to the same variable in one scope. Yet the only way that problem can occur is for programmers to reuse the same variable for different purposes. Introducing significant costs simply to provide for better inference of “bad” code is untenable. Moreover, neither SSA nor reachable definitions can handle attribute accesses easily, especially not in the presence of threads, generators, or continuations. The real benefit of flow-insensitivity is freedom from having to care about time.

3.1 Nodes and Constraints

Starkiller’s type inference algorithm works by constructing a dataflow network for types that models the runtime behavior of values in an input program. This dataflow network consists of nodes linked together by constraints. Of course, since Starkiller only has access to information available at compile time, it must make conservative approximations. This means that Starkiller will occasionally infer overly broad type sets for an expression, but it also means that it will never fail to infer the existence of a type that appears at runtime.

Each variable or expression in the input program is modelled by a node that contains a set of possible types that expression can achieve at runtime. Initially, all sets are empty, except for constant expressions, which are initialized to contain the appropriate constant type. Nodes are in turn connected with constraints that model data flow between them. A constraint is a unidirectional link that forces the receiving node’s type set to always contain at least the elements of the sender node’s type set. In other words, constraints impose a subset relationship between the pairs of nodes they connect. As a result, types flow along constraints. When a node receives a new type, it adds the type to its type set and promptly dispatches that type to all other nodes connected to it. The algorithm continues until type flow has stabilized.

The preceding discussion raises a question as to how different elements in input program source code generate nodes and constraints. We examine the simplest case, an assignment statement presently and defer more substantive cases for the following sections. Consider the assignment statement `x = exp` which binds the expression `exp` to the name `x`. Starkiller processes this statement by building a node for `x` if one is not already present and building a constraint from the node corresponding to `exp` to the node corresponding to `x`. That constraint ensures that any types in `exp`’s type set eventually propagate to `x`’s type set.

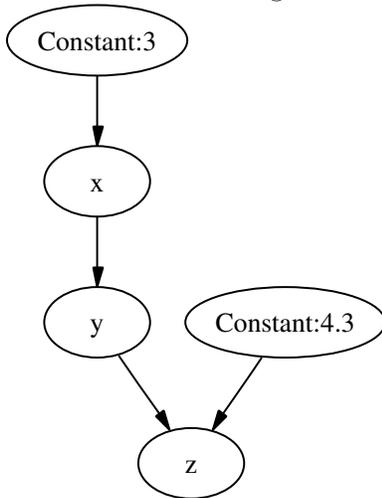
```
x = 3
y = x
z = y
z = 4.3
```

Figure 2: Assignment in action.

As a further example, consider the source code in Figure 2 and the corresponding constraint network shown in Figure 3.1. Initially, Starkiller creates nodes with empty type sets for the variables `x`, `y`, and `z`. It also creates nodes for the constant expressions `3` and `4.3`. However, those two nodes have type sets that consist of either the integer type or float type respectively. The assignments indicated dictate that Starkiller place constraints between `3` and `x`, `x` and `y`, `y` and `z`, and `4.3` and `z`. As a result, once processing has completed, the type set for nodes `x` and `y` will contain exactly one element: the integer type. The type set for node `z` will contain two types: the integer type and the float type.

It is important to realize that there are a variety of nodes, each with different behavior. The simplest nodes are variable nodes whose behavior is as described above. There are

Figure 3: Constraint network for code shown in Figure 2



more complicated nodes for function calls, function definitions, class definitions and instance definitions. Another important feature of the system is that constraints can be named. This allows a function call node to distinguish between constraints representing different arguments, for example.

3.2 Functions

We now examine how Starkiller performs type inference for code that defines or calls Python functions. But first, we review Python semantics regarding function definitions and calls. In Python, functions are first class objects and function definitions are imperative statements that are evaluated at runtime. Consequently, the definition shown in Figure 4 serves to create a function object and bind it to the name “factorial” when encountered at runtime. Because definitions are evaluated at runtime, the same function definition can easily be instantiated into different function objects. This point is illustrated by the source code in Figure 5. Successive calls to `makeAdder` will return distinct function objects since the definition of `add` triggers the creation of a new function object for each invocation of `makeAdder`. Python functions can have default arguments; any such expressions are evaluated once at function definition time. This feature has important implications for type inference.

Upon encountering a function definition, Starkiller creates a special function definition node that encapsulates the code associated with the definition. It also creates a variable node representing the function’s name and a constraint from the definition node to the name node. The definition node generates function type objects which then become part of the variable node’s type set. In this way, Starkiller models the runtime behavior of function definition as assignment of a newly created function object. Starkiller also creates constraints from any default argument expressions to the definition node, so that as information about the types of default arguments reaches the definition node, it can produce updated function types that incorporate that information.

```
def factorial(n):  
    if n == 1:  
        return 1  
    else:  
        return n * factorial(n - 1)  
  
x = factorial(5)  
y = factorial(3.14)
```

Figure 4: The factorial function.

```
def makeAdder(a):  
    def add(b):  
        return a + b  
    return add
```

Figure 5: A nested function.

When Starkiller finds a function call, it creates a special function call node. It then creates a constraint from the variable node associated with the name of the function being called to the call node. It also creates constraints from each of the nodes associated with the actual arguments to the call node as well. Finally, Starkiller creates a constraint from the call node to wherever the return value of the function call lies. This constraint is used by the call node to transmit the type of the function’s return value.

The call node responds to incoming types by taking the Cartesian product over the list of sets of callee types and argument types. The result is a list of monomorphic types where the first entry is the type of a callee and successive entries represent the argument types. For each monomorphic entry, the function call node attempts to find a matching template. Templates are unique instantiations of the node and constraint graph associated with a single function’s source code. Each template associates a monomorphic type for each of the function’s arguments. Thus, for any function, there could be many templates that differ only in the argument types. Templates are shared across all call sites of a particular function, so each function in principle only needs to be analyzed once for each monomorphic call signature. If the call node can find a matching template, it adds a constraint from that template’s return node to itself so that it can propagate return value types to its caller. If no template exists, the call node creates one, building a set of nodes and constraints.

Starkiller does slightly more than the preceding discussion indicates in order to properly handle lexical scoping. Because nested definitions have access to variables defined in their enclosing scopes, function types must depend on the types of out-of-scope variables referenced by the function. In other words, function types must effectively include the types of all variables they reference that they do not define. For the example code in Figure 5, this

means that the type of the function `add` incorporates the type of the argument `n`. One benefit of this approach is that function types are completely self contained: they include all information needed to perform inference at a function call site. There are no hidden “lexical pointers” that connect function calls back to the original defining context.

In this way, out-of-scope variables are treated like default arguments. When analyzing functions, Starkiller keeps track of which names are read by each function and which names are written to. This information is used to statically determine a list of out-of-scope dependencies for each function. Starkiller compiles these nonlocal dependencies recursively, so a if a function’s lexical parent does not define a nonlocal name that function references, then that name is nonlocal for the parent as well. At function definition time, the types of all nonlocal names are immediately available. Starkiller adds constraints from each nonlocal name to the function definition node, which in turn produces function types that incorporate the types of nonlocal names being referenced. This process is somewhat akin to a common optimization technique for functional languages known as lambda lifting.

The inference algorithm described so far has difficulties analyzing recursive functions. Consider the factorial function shown in Figure 4. When Starkiller encounters the definition, it creates a definition node and a name node and links them together. However, because factorial is recursive, it references itself and lists its own name as a nonlocal dependency. Starkiller thus adds a constraint from the name node for factorial to its definition node. Since factorial is a nonlocal name for itself, the function type generated by the definition node must include the type of factorial itself. But now we have generated a cycle. The definition node produces a function type which flows to the name node and then back to the definition node as a nonlocal reference type. As a result, the definition node performs CPA on its list of nonlocal and default argument type sets and produces a new function type that includes that old function type. Rinse, wash, and repeat.

Starkiller solves this problem by detecting the cycle and stopping it. When a definition node sees a new type being sent for a nonlocal name or default argument, it checks to see if that type or any type encapsulated by it is one of the types the definition node itself has generated. If the incoming type contains a type generated by the definition node, there must be a cycle, and the definition node ignores it. This architecture is designed to alleviate the problems that plagued Agesen’s CPA, namely an inability to precisely analyze recursive functions without becoming vulnerable to recursive customization. While not perfect, Starkiller’s approach has significant improvements over the heuristics Agesen employed to detect and stop recursive customization.

3.3 Classes and Objects

Having explored how Starkiller handles functions, we now turn our attention to how it handles Python’s object system. As before, we briefly review Python’s semantics before diving into Starkiller’s type inference algorithm.

Classes in Python are defined using an imperative statement in much the same way functions are. Within a class body, variables defined are class specific, that is, they are accessible by all instances of the class. Functions defined inside the class are methods. Unlike

a function definition, the body of a class definition is executed when the class definition is first encountered. The result of a class definition is the creation of a new class object and the binding of that object to its name. Once created, a class' methods and attributes can be modified or added, even after instances have already been created for it. Class definitions explicitly include an ordered list of base classes. This list can be modified at any time.

Class instances are created by calling the class as a function. This produces an instance as a result and also calls the class' constructor method with the calling arguments as a side effect. As with classes, instances can have attributes added or modified at any time simply by assigning to them. Attribute references using the `self` parameter are first resolved in the instance, then in the class, and then in the class' base classes. Attribute writes always occur in the instance; writing to the class or one of its base classes is permitted via explicit naming (i.e., `self.age = 3` for instance attributes versus `self.__class__.age = 3` for class writes). When an attribute lookup fails for the instance and the class, the base classes are perused in the order indicated by a depth first traversal of the inverted base class tree. Attribute accesses in the source code generate get attribute and set attribute nodes in the constraint network. These nodes have constraints to them indicating what object should be queried, and, for set attribute nodes, what the new value should be. The attribute name is a constant string encoded upon node construction.

Upon encountering a class definition, Starkiller creates a variable node to represent its name, a new class definition node to generate appropriate class types, and a constraint linking the two together. It also creates constraints to the definition node from the nodes corresponding to the listed base classes as well as any nonlocal references. The definition node maintains type sets for each base class and nonlocal reference and takes the Cartesian product of those sets to generate a list of monomorphic types. These monomorphic types are packaged as class types and dispatched from the definition node. Reading or writing a class attribute is simple: when a class type reaches a get attribute node, the attribute type can be read directly from the class type. If it doesn't exist, then no type is returned. When a class type reaches a set attribute node, its state does not change. Class types are immutable and once created cannot be modified. Instead, they contain a reference back to the definition node that produced them. Rather than change their state, they inform the definition node of the state change, and if necessary, the definition node generates new class types incorporating that change which eventually propagate throughout the system.

When a class type reaches a function call node, it creates an instance definition node at the same place. Much like the class definition node, the instance definition node acts as the repository of the instance's polymorphic type state. The instance definition node generates monomorphic instance types that contain a single type for every attribute defined by applying the cartesian product algorithm over its polymorphic type state. It also creates a new function call node to represent the call to the class constructor, and adds constraints from the original function call to this new one to simulate argument passing. The class constructor is extracted using a get attribute node so that if later information adds a new constructor or new base class, all possible constructors will be called.

Because instance attributes shadow class attributes, precision can be lost. For example,

a common design scenario is to encode the initial value of an attribute as a class variable, but have instances write newer values as instance attributes directly. This technique can make the initialization logic simpler and cleaner since the instance always has a correct value available.

When an instance type reaches a get attribute node, the attribute name is looked up in the instance type itself, its class type, and its' class type's base classes. All types found as a result of these lookup operations are returned to the get attribute node. This parallel attribute lookup suggests a possible loss of precision since, at most, only one lookup result will be passed to the get attribute node at runtime. However, this loss in precision is necessary to properly deal with the fact that instance attributes shadow class attributes with the same name. However, because attribute writes must be explicitly qualified, writes do not impose a precision loss. When an instance type reaches a set attribute node, it informs its instance definition node of the resulting type change, but it does not change. Like class types, instance types are immutable. Note that Python classes can override the attribute read and write functions with custom methods. This means that, in addition to the process described above, Starkiller must create a call for the `__getattr__` method whenever that method exists and the attribute cannot be found. Calls to the `__setattr__` method must be invoked whenever it exists as well.

Instance method calls require some explanation. Python parses a method call like `inst.method(1,2)` into an AST that looks something like `CallFunction(GetAttribute(inst, 'method'), (1, 2))`. In other words, a bound method is first extracted from the instance using ordinary attribute lookup and that bound method object is then called just like a normal function. Bound methods are first class objects, just like functions. The instance attribute machinery in Starkiller packages both the instance type and the method type together into a bound method type that is returned to the get attribute node. This behavior only happens when the attribute name is found inside a class or base class and that resulting attribute type is a function.

4 Advanced Language Features

4.1 Operators

Python allows individual classes to define methods that specify the behavior of overloaded operators. For example, a class that wanted to implement the addition operator would define a method named `__add__` that takes the instance and the other object being added as parameters. While operator calls internally reduce to method calls, their semantics introduce sufficient complexity so as to preclude treating them as mere syntactic sugar for method calls. This complexity stems from Python's coercion rules, and their long storied history of informal specification, special casing, and evolutionary change.

These coercion rules specify which method will actually be executed for a particular operator call given a pair of operands. For an operator `op` the coercion rules specify selection of a method named `__op__` defined in the left operand. If such a method is not defined, the

rules mandate that a method named `__rop__` defined for the right operand be used, where the `r` prefix indicates right associativity. If that method is undefined, the implementation throws a `NotImplemented` exception. There are further complications, such as special case rules for operator precedence of new style classes where the right operand is an instance of a proper subclass of the left operand. In-place operators (e.g., `+=`) introduce further complexity since a method `__iop__` is searched first and used without coercion if found, but, if that method is not defined, execution falls back to a combination of the standard operator (`__op__`) and assignment. A complete copy of the coercion rules in all their hideous soul devouring glory can be found in Section 3.3.8 of [10].

Starkiller handles operator calls by creating an operator node that has constraints from the operand types and that generates the result type for that operation. Internally, the operator node performs CPA on the operand types just like a function call node, but unlike a function call node, it enumerates the resulting monomorphic operand type lists and uses the coercion rules to determine which method from which operand should be called for each monomorphic pair of operands. Having found the method, it builds a get attribute node to extract the correct method from the correct operand and a function call node to call that method so as to determine the result type of the operation. Since the result types of all of these function call nodes point back to the operator node using a named constraint, the operator node can easily generate a return type for the value of the operation expression.

4.2 Exceptions

Python provides structured exceptions not unlike those found in Java and C++. One notable difference is that functions and methods do not (and can not) declare a list of exceptions they can potentially throw. Exceptions can be any Python object, although exception comparisons in the catch clause of `try/except` statements are based on class and subclass comparisons. In other words, an exception will be caught by an exception handler only if it is an instance of a class or subclass of the handler's exception target.

Starkiller's type inference algorithm refuses to deal with exceptions in any way. This is because exception handling is complex and delicate and in many ways, represents a corner case in language implementation. In addition, with one small exception, exception handling provides no useful type information to a flow insensitive algorithm. Traditionally, Python exceptions have been understood as a slow mechanism for transferring control. This experience is also true in C++, Starkiller's target language. As a result, exception handling invariably occurs outside the critical path, and thus sees little if any benefit to aggressive optimization.

The one instance where exception handling provides necessary type information to a type inferencer is for programs that make use of named exceptions. Such programs use a variant of the `except` clause in a `try/catch` statement to bind the captured exception to a local variable and make it available to exception handler. Starkiller cannot perform complete type inference on programs that use this feature since it does not process type information for exceptions. One simple alternative would be to use the type of the exception handler's target as a surrogate for the type of the exception actually caught, but this approach would

violently break language compatibility. The crux of the problem is that the exception object that must be bound locally is generated in a later call frame, and we cannot determine that without exerting significant effort to track the movement of exception objects across call frames.

While the current implementation of Starkiller does not process exceptions, later versions easily could using the strategy outlined below. Doing so would involve using the static call graph, of which Starkiller has ready access to a conservative approximation. Each function call node already maintains constraints to itself from the return nodes of each template it makes use of. Exception handling could work in much the same manner, where function and method bodies maintained a hidden variable to record the types of all uncaught exceptions and passed them back to their callers using a named constraint in the same way that they pass return values to their callers.

In order to deal with the fact that a given function can contain multiple possibly nested try/except clauses, we introduce an exception node for each of them and build constraints between them over which exceptions flow based on their scoping relationships. Function call nodes have named constraints emanating from themselves to the most closely scoped exception node. When they encounter an exception type from one of their callee templates, they pass it along to the nearest exception node, which may make it available in a local binding if the exception matches the target or may propagate it onward to the next outer exception scope. This design supports all Python exception semantics while ensuring that inference remains complete, even in the face of named exceptions. Its major drawback is that precision for exception expressions may be reduced in some cases, but it should be more than adequate to improve upon the current Python implementation's performance.

4.3 Iterators and Generators

Recent Python releases have introduced iterators [11] and generators [8] into the language. Their design was heavily influenced by the Sather [7] and Icon [4] programming languages, where such objects play a vital role. Iterators are objects that follow a standard protocol whereby they provide clients with a stream of intermediate values using a `next` method and signal the end of iteration by raising a `StopIteration` exception. For loops act naturally over iterators. Generators are a particular type of iterator object that make use of the `yield` keyword. They allow functions to naturally suspend their entire execution state while providing intermediate values to their callers without terminating.

As with exceptions, Starkiller provides no support currently for generators and iterators. Due to the tight relationship between iterators and exceptions, support for the latter necessitates support for the former. Nevertheless, once Starkiller's type inference algorithm properly supports exceptions, it can be easily extended to support iterators and generators. Iterator support can be added by modifying how Starkiller analyzes for loops to check for and make use of `__iter__` and `next` methods when present. Generators can be supported by recognizing functions that contain the `yield` keyword and ensuring that any calls to them result in a generator object that implements the iterator protocol. Each invocation of the `next` method should be equivalent to executing the body of the generator, with the `yield`

statement being used to return values to the caller instead of the return statement as used in traditional functions.

4.4 Modules

The Python module system allows statements of the form `from email.encoders import encode_base64` and `import socket`. When using the former, objects named on the import statement are made immediately available to the importing module after the import statement has executed. When using the later form, module objects are bound to a variable with their name. A module's contents can be accessed using named attributes, such as `socket.gethostbyname`. Starkiller reflects these semantics by representing each module with a type. Note that the term module can refer to a Python source file or an extension module with a corresponding Starkiller type description.

Like function types, module types have templates associated with them, but unlike function types, there is only one template associated with each module type. Upon encountering an import statement, Starkiller locates the needed module and checks whether it has been analyzed before. Since Python caches module imports in order to ensure that each module is executed only once, Starkiller replicates this behavior by only analyzing a module once and thereafter referencing its template. For named imports, where some or all of the module's contents are imported into another module, Starkiller creates a variable node in the importing module for each name being imported. It then builds constraints from the original variable node in the imported module template to the newly created variable node of the same name in the importing module.

5 Foreign Code Interactions

Because it must perform type inference on programs that use both native Python source code as well as Python extensions written in other languages (typically C, C++, or Fortran), Starkiller is faced with a serious problem. Moreover, these external languages and the APIs they use to interact with Python are sufficiently complex so as to make type inference infeasible. In order to resolve this impasse, extension module authors are expected to write simple descriptions of the run time type behavior of their modules. These External Type Descriptions are written in a language that Starkiller specifically provides for just this purpose. In this section, we begin by exploring the External Type Description language. We conclude by examining the external type description for one particular module, the `__builtins__` module, which defines all of Python's core types and functions.

5.1 External Type Descriptions

External Type Descriptions are Python classes that subclass an `ExternalType` class provided by Starkiller. That base class provides its subclasses with a stylized interface into

Starkiller’s internals. There are also `ExternalFunction`, `ExternalMethod` and `ExternalModule` base classes provided to complement `ExtensionType`. The power of the extension type description system is difficult to understate; it allows extension authors to “plug into” Starkiller’s type inference machinery and make their own code part of that machinery. The relationship between Lisp and its macro system is analogous to the relationship between Starkiller and its external type description system.

Extension type description authors write description classes that describe the run time type behavior of their extension types. These classes are subclassed from `ExtensionType` and can use its `getState` and `addState` methods to manipulate the polymorphic state of that particular instance of the extension type. These description classes include methods that describe the behavior of the actual methods implemented by the extension type. For example, when an extension type instance’s `append` method is called, Starkiller will invoke the corresponding method in the extension type to determine the method call return type. The `append` method is presented not just with the monomorphic argument types of the method call, but also the result node that it will be directed to and the monomorphic state of extension type instance. The method can retrieve or add new type state to the extension type or any other type visible to it and it can also use the `callFunction` method to simulate arbitrary function calls. This feature can be used to model extension functions that return result types that depend on calling one of their arguments, such as `map`. In fact, the method can apply arbitrary transformations to the constraint network.

External types are treated in much the same way as class instances. In particular, their internal state is described completely by a mapping from state names to polymorphic variables. For example, if an external list type has method names `append` and `insert`, it will have a state element for each with their corresponding name and a type set consisting of a single external method type. As with class instances, when external type instances reach a get attribute node and are asked for a method, they generate a bound method type that encapsulates both the external type instance as well as the method name. When this bound external method type reaches a call function node in the network (i.e., when someone attempts to execute an external type’s method), the corresponding extension type method is called and asked to supply a return type for the operation. In lieu of supplying a return type, the method has the option of creating a link to the destination node for the return type so that it can be filled persistently. Polymorphic external type instances are associated with instance definition nodes that generate appropriate monomorphic external type instance state types that propagate through the network. Because their architecture mirrors that of native Python classes, external types see the same level of precision that native Python classes do.

5.2 The builtins Module Type Description

Python includes a special module called `__builtins__` that encompasses all the basic types and functions accessible by default. It includes named constructors for the basic types of machine integers, double precision floating point numbers, long integers, lists, tuples, and dictionaries. These basic datatypes are completely pervasive in all real Python programs; in

fact, certain syntactic constructs use them implicitly. For example, tuple unpacking happens automatically in assignment statements like `a, b = foo()`. Moreover, a class' base class list is represented as a tuple of classes stored in an attribute named `__bases__`. In order to properly analyze programs using these basic types, Starkiller includes an external type description for the `__builtins__` module. Those type descriptions will be described briefly, both to illustrate the issues involved in writing a proper external type description and to showcase the limits of inferencer precision when using basic datatypes. It is important to note that for a given external type, there can be many different external type descriptions that differ significantly in precision and run time inferencer performance.

Contrary to their name, Python lists are vector like data structures capable of holding any type of element while supporting fast appending at the end and constant time element writing and retrieval. Starkiller provides an external type description for lists that maximizes precision in the common case while degrading gracefully in the face of operations that necessitate imprecision. The crux of the problem is that while we want to be able to associate precise types with individual slots of a list, there are some list operations that modify the lists' type state in a way that is impossible for a flow insensitive inferencer to trace. For example, consider the case of a list initialized with the statement `L = [1, 2.2, z]`. Starkiller has enough information on the list to determine that the type of the first element of the list will be an integer. However, if at some later point in the program, `L.sort()` is called, we can no longer make that guarantee since `sort` swaps the list's elements in place.

Consequently, the external type description for lists maintains a state variable indicating whether that particular list has been "tainted" by such operations. The list keeps a state variable which describes the type set of each slot and an extra state variable describing the type set of all slots combined. Reading from a particular slot of an untainted list yields the type set of that slot, while a slot read from a tainted list yields the combined type of all slots. Lists initially are untainted, but can become tainted when they become subject to a tainting operation, like the `sort` method described above. When that happens, additional types may be propagated through the constraint network to reflect the fact that previously single slot reads may actually refer to the contents of any slot. This protocol works well with CPA because it never violates the monotonicity principal: the set of types produced by reading a list item is always (partially) correct, but may not always be complete. Adding types is safe, but removing them is not.

Tuples use a very similar mechanism, namely maintaining state for each element. Of course, since tuples are immutable, they need not deal with tainting operations. Dictionaries use a similar method to associate key and value types together. Astute readers will notice one serious problem with the scheme described above: Starkiller deals in types, not values, but in order properly associate individual list indices with unique types, it must necessarily keep track of values. In general, Starkiller tracks types, however, in a limited number of special cases, it will actually process constant values. The two special cases are attribute names and indexing operations. Failing to provide special case value handling of constant attribute names would make Starkiller so imprecise as to be useless: the type of any attribute referenced would have to be equivalent to the combined type set off all attributes. Indexing

using the `__getitem__` and `__setitem__` methods plays just as serious role. Without special case support for constant index expressions, operations involving implicit tuple packing and unpacking would lose a great deal of precision. Unfortunately, such operations are utterly pervasive: they are simply too convenient not to use. Python code generally includes both constant and variable forms of both attribute access and indexing. Thus we see expressions like `x.attr` and `y[2]` in addition to expressions like `getattr(x, userInput())` and `y[z*2 - 1]`.

The challenge for extension type description authors lies in designing extension types descriptions that can make good use of constant information when available while degrading gracefully in the face of variable attribute and index accesses, all the while maintaining the monotonicity invariant that Starkiller relies upon.

6 Known Problems and Limits

6.1 Eval, Exec, and Dynamic Module Loading

The most glaring deficiency in Starkiller's type inference algorithm is its inability to handle dynamic code generation. Starkiller needs to be able to see all the code that could be executed at runtime in order to perform complete inference. As a result, constructs such as `eval` and `exec` which evaluate and execute source code potentially made available only at run time are problematic. For the same reason, Python's module system can theoretically be used to import code that does not exist at compile time since `import` statements are evaluated only at run time.

While Python offers several methods to execute code generated at runtime, in practice, `eval`, `exec`, and module imports that cannot be resolved statically are rarely used. In part, this is because dynamic code generation introduces security risks for the same reason that it makes static type inference difficult: code that has not been audited and verified prior to distribution cannot be trusted, especially if that code could have been contaminated by untrusted user data. Another reason is that code generated dynamically often runs slower than comparable static code. This slowdown is due to the fact that the dynamically generated code must be byte compiled before execution, while static code is byte compiled only once on application startup and then cached for successive executions.

Finally, dynamic code generation is infrequently used in the Python world because Python makes it easy to accomplish the same ends as achieved with dynamic code generation using other means. For example, in order to convert the string `'1'` into the corresponding integer, one could use the `eval` function or one could simply use the `int` function. The latter has the benefit that its return value will always be an integer should it return at all. In addition, the developer can rest secure in the knowledge that no extraneous code can be invoked by a call to `int`, regardless of the value of the input string. As another example, consider the case of a developer attempting to read an attribute from an object `instance` where the attribute name is derived from run time data. Solutions to this problem using `eval` such as `eval('instance.' + userInput())` compare poorly in sim-

plicity, performance, and aesthetics when compared to more traditional solutions such as `getattr(instance, userInput())` or even `instance.__dict__[userInput()]`.

To determine how serious a problem Starkiller's inability to support dynamic code generation may be, I performed a survey of the library code written in Python that was shipped with release 2.3 of Python. I examined all instances where `eval` and `exec` were used and recorded the difficulty involved in working around their use. The results are shown in the following table. One immediate observation is that `eval` and `exec` are infrequently used: there were less than 25 occurrences of both functions in over 123,000 lines of source code. Moreover, many of those occurrences can be easily removed in favor of alternative facilities afforded by the language. All of the remaining uses occur in programs or libraries that are simply not suited for Starkiller's target audience. These programs include heavily interactive applications like the Python Debugger and the interface to libreadline in addition to more esoteric applications, like the compiler infrastructure.

`bdb.py`, `pdb.py`

Debugger support

Not applicable for Starkiller. Implements run time support for the Python debugger, and uses `eval` to modify variables at run time after stopping at a breakpoint.

`dumbdbm.py`

Dumb DBM Clone

Trivial to work around. Uses `eval` to cheaply parse string into three numbers. Security hazard since that allows arbitrary code execution.

`gettext.py`

Internationalization and Localization support

Trivial to work around using named functions.

`gopherlib.py`

Gopher protocol support

Easy to work around using dictionaries instead of module variables as a static mapping, or, alternatively, a call to `globals`.

`mhlib.py`

Unix Mail Handling

Not applicable to Starkiller since `eval` is only used in testing code.

`os.py`

Generic Operating System interface

Moderately difficult to work around. `eval` is used to support conditional compilation. `eval` could be eliminated if `os.py` was better modularized with `autoconf`, but this is not for the faint of heart. Alternatively, `globals` could be used.

`random.py`

Random number generator

Not applicable to Starkiller since `eval` is only used in testing code.

`reconvert.py`

Conversion between old and new regular expressions

Not applicable to Starkiller since this code is only used during development.

`rexec.py`

Deprecated restricted execution framework

Very difficult to work around, but this code is going to be removed soon since it is terminally insecure.

`rlcompleter.py`

Word completion for libreadline

Not applicable to Starkiller since this code is only used in interactive applications.

`tzparse.py`

Experimental parser for timezone specifications

Trivial to work around. Uses `eval` to parse simple strings into numbers.

`warnings.py`

Generic warning framework

Moderately difficult to work around. Uses `eval` to check if user supplied names are defined. Could probably be replaced with dictionary mapping or a much safer call to `getattr`.

`compiler/ transformer.py`

Experimental Python parser

Moderately difficult to work around. Uses `eval` to parse literals in a way that precisely matches the C implementation of Python's parser. `eval` could probably be replaced if a more detailed literal parser was written in Python.

`logging/ config.py`

Configuration parser for logging system

Moderately difficult to work around. Uses `eval` to map class name strings to the corresponding class objects. `eval` could probably be replaced by an explicit dictionary mapping or a call to `getattr`. This is also a potential security vulnerability allowing arbitrary code execution.

`idlelib/ CallTrips.py, ObjectBrowser.py, PyShell.py`

Integrated Development Environment

Not applicable to Starkiller since this code uses `eval` to support dynamic introspection and live object modification for an IDE.

Despite the problems inherent in supporting dynamic code insertion in a high performance static compilation environment, Starkiller will eventually support some (perhaps limited) form of dynamic code insertion. We describe possible implementation strategies presently. The most immediate problem faced in implementing `eval` is the need for a parser and interpreter or compiler at run time. Once the dynamic code has been parsed and interpretation begins, other problems quickly emerge. Consider the following possibilities:

1. the evaluated code could call functions with argument types that do not match any existing template
2. the evaluated code could assign a value with a type not seen before as an object's attribute
3. the evaluated code could replace code or data that had previously been assumed to be constant and thus inlined

The first problem is easily solvable if we mandate that the compiler generate a generic template for each function and method in addition to whatever specialized templates the type inferencer indicates are required. Such generalized templates treat all arguments as polymorphic and perform explicit type checks throughout in much the same manner as the current implementation of Python does today. Including them imposes no performance overhead beyond increased executable size. The second and third problems are interrelated and pose far more difficult challenges. In and of itself, the second problem is not insurmountable since any data member could easily be replaced with a pointer to new tagged data. The larger difficulty stems from the existence of (possibly inlined) code throughout the system that accesses that data.

There are two possible solutions to this problem. One solution would involve placing layers of indirection so that any code or data that could be modified by dynamically inserted code could be easily replaced. This would impose a severe performance penalty unless developers were given some method to declare code “sealed” from further changes. The second solution requires that Starkiller recompile the entire world image while running whenever it needs to handle dynamically inserted code. The overhead involved in this recompilation would be substantial, making it unsuitable for frequent uses of the `eval` function and practical for little more than updating modules in the field. However, this recompilation process is amenable to several optimizations. For example, recompilation might only be used as a last resort if the evaluation interpreter ran into any of the three problems described above. In the common case, evaluation would proceed without incident. Once recompilation was triggered, the resulting world image would be generalized in ways that would make recompilation less likely to be needed again. Another optimization requires that Starkiller track dependencies between different functions, methods, and data definitions so that it can recompile only the parts of the program that are affected by the evaluated code. However,

this approach presents further performance challenges since it requires that Starkiller take full control of inlining away from the target compiler.

In general, the need to perform run time code patching is difficult to reconcile with the peak performance since very high performance compilation techniques make heavy use of inlining and other techniques that destroy program modularity. The conflict between high performance and dynamic code generation is not impossible to reconcile, but it remains quite challenging.

6.2 Static Error Detection

The standard Python implementation provides a consistent model for handling errors. A small class of errors, mostly syntactic in nature, is detected by the byte compiler when a program is compiled or run for the first time. These errors typically abort compilation or execution entirely. All other errors are detected at runtime and are dealt with using Python's exception handling mechanism. In practice, this means that many errors that are detected statically in other languages, such as calling a function with more arguments than are expected, become run time errors that are only detected upon execution. This error handling model poses a problem for Starkiller: what should one do with potential errors detected statically by the type inferencer?

```
x = "hi there"  
x = 4  
y = x + 3
```

Figure 6: An example of static detection of run time errors.

To illustrate the problem, consider the simple program in Figure 6. This code will compile and execute correctly in the standard Python implementation without issue. But when Starkiller analyzes this code, it will conclude that the variable `x` can have a type of either string or integer. Since `x` is then used as an operand for an addition with an integer, Starkiller cannot blindly invoke the addition without resolving the `x`'s polymorphism. Because it is flow insensitive (see Section 3), Starkiller cannot determine that the string value of `x` will not be present by the time control reaches the addition expression. Having failed to statically resolve `x`'s polymorphism, we must do so dynamically by inserting a run time type check on `x` before proceeding with the addition. If `x` is an integer, then the addition can proceed without harm. But if `x` is a string, we have a problem, since the addition operator is not defined for operand types of string and integer.

What should the type inferencer do at this point? Aborting compilation because an error is possible would be unwise in this case since we know that the string type associated with `x` is a spurious artifact of the type inference algorithm and not something that will hinder the addition operation. At the same time, we cannot simply ignore the issue since it could truly be a programmer error. Consider the same case as before, but with the first two lines

swapped. That program would compile correctly but would generate a runtime exception when fed into the standard Python implementation. Python’s inability to statically detect errors like this as well as simpler errors such as typographic errors in variable names has been a major complaint from its user base since its inception. That suggests that Starkiller should make static errors visible to the user so they can check and repair incorrect code. Unfortunately, doing so would also produce spurious warnings as seen in the first example. Moreover, due to the effects of templating, the same error may be detected many times, drowning the user in a sea of meaningless errors with no easy way to separate the spurious from the legitimate.

The core of the problem is that some of the errors that Starkiller statically detects represent real errors that the standard Python implementation cannot statically detect while others represent artifacts of the type inferencer’s imprecision and are completely harmless. There is no way, a priori, to discriminate between the two. Because Starkiller focuses on improving run time performance and not necessarily run time safety, it forgoes reporting statically detected errors such as those described above to the user. However, it does not ignore possible error cases. Instead, it inserts code to raise run time exceptions as needed when it statically detects that an error is possible and dynamically determines that an error has occurred. This approach preserves the standard Python semantics while bypassing the need to build an extensive error reporting and filtering system. Nevertheless, future research would be well directed at improving static error detection and presentation to developers.

6.3 Partial Evaluation

Starkiller’s implementation for dealing with attributes is rather inelegant. Essentially, support for explicit constant attribute names has been hard coded into the type inferencer as a special case, since the program source contains them directly. A more elegant approach would have been to treat all attribute accesses like invocations of the `getattr` function. Since this function cannot know statically what attribute name it will look up, it returns the types of all attributes together. Unfortunately, named attribute references are far too common for this approach to have been practical; its application would have rendered all non-trivial programs a morass of imprecision. However, it would be desirable to replace Starkiller’s special case logic for dealing with attribute access with a more generalized partial evaluation framework, since constant explicit attribute access is simply a special case of partial evaluation. This approach would provide other benefits as well, such as improving precision when analyzing some common programming idioms.

One such idiom describes how to build a class that deals with many other client classes without expanding those client classes. It works by implementing one method for each client class to be examined and uses Python’s introspection features to select the correct method to call based on one of the client class’ data attributes. For example, consider Figure 7 which portrays a stereotypical implementation of the Visitor design pattern [2] used to implement part of a compiler. The compiler defines classes for all possible nodes of an abstract syntax tree, including nodes to represent constants, if statements, and for loops. If we wanted to write code that would operate on instances of all those nodes without touching the original

classes, we might build a visitor class as shown in Figure 7. This class uses introspection to determine the class name of any instance presented to it and then uses that name to construct an appropriate method name to process that instance. This is a well known idiom in the Python community.

```
class Constant(Node):
    pass

class If(Node):
    pass

class For(Node):
    pass

class visit:
    def processNode(self, node):
        return getattr(self, 'visit' + node.__class__.__name__)(node)

    def visitConstant(self, node):
        pass

    def visitIf(self, node):
        pass

    def visitFor(self, node):
        pass
```

10

20

Figure 7: The visitor design pattern in Python.

Starkiller handles code like this, but not well. Because it performs no partial evaluation, it assumes that the `getattr` call can return any value that is an attribute of instances of the visitor class, when in fact, it can only return attributes whose names begin with the string “visit”. Starkiller’s conservative assumption leads to imprecision which in turn hinders performance and increases generated code size since extra type checks and error handling code must be synthesized to deal with the invocation of other attributes besides the visit methods. A more general solution would obviate the need to check for impossible combinations that cannot arise in practice.

6.4 Integer Promotion

Python offers two builtin integer datatypes: a standard “machine” integer that must be at least 32-bits wide and a long integer that can have unbounded length. Traditionally, these

two datatypes were kept distinct and isolated from one another. This meant that when a machine integer value overflowed, the Python Virtual Machine would raise an `OverflowError` exception. Long integers, of course, cannot overflow since they use arbitrary precision arithmetic to grow as needed. Starkiller’s problem stems from the fact that recent Python releases have adopted a change in semantics [12] designed to eventually unify the machine and long integer types. Newer Python releases now respond to overflow of a machine integer by silently promoting it into a long integer. The eventual goal of these changes is to eliminate machine integers completely, but that is not expected to happen for quite some time.

While unifying the two integer datatypes does solve some very real problems in the language, it introduces a host of new problems for Starkiller. There are four options for Starkiller to implement Python’s newer integer semantics. The first approach requires that Starkiller mirror the language definition precisely. The second option would be to implement sophisticated techniques for range analysis and propagation associated with previous work in using unboxed integers in higher level languages. The third option would be to uniformly use long integers everywhere, while the fourth option would be to retain the integer semantics found in older versions of Python. At the moment, Starkiller uses the fourth option, opting for high performance and limited implementation complexity at the cost of conformance with the language specification.

Precisely mirroring the language definition means that machine and long integers remain distinct types but that all operations that could overflow a machine integer must be checked at runtime for overflow and promoted to long integers as needed. From a type inference perspective, this means that all primitive integer operations such as addition and multiplication may return either machine integers or long integers. As a result, almost all integer variables will be polymorphic since they must hold machine integers as well as long integers. Making all integer variables polymorphic cripples the performance of generated code by inhibiting the use of unboxed arithmetic. The standard Python Virtual Machine does not face this problem since it uniformly uses boxed integers anyway, and, in any case, it has enough overhead to mask that caused by integer promotions.

The literature associated with compiler optimization for dynamic languages is full of techniques for partially reclaiming the performance lost to integer boxing [3, 9]. Unfortunately, many of these strategies necessitate a degree of implementation sophistication that is not presently available to Starkiller. Recent work on unboxing in statically typed languages [6] such as Haskell and the ML family of languages may prove more appropriate to Starkiller’s static compilation model while posing less of an implementation hazard.

In contrast, using long integers exclusively offers somewhat better performance than is possible with strict conformance to the standard since all integer variables can be at least monomorphic. Long integers could even be implemented in an unboxed manner, if they were implemented as a linked list of machine words. In that implementation, long integers would contain one word of the integer data and a pointer to the succeeding word. For most integers, the next pointer would be `NULL`, signifying that this was the last word comprising the integer. As a result, long integers would be twice the size of machine integers in the common case and could be stack allocated but would incur extra overhead needed to constantly

check if the next pointer was null. This overhead would manifest itself in repeated branch points inserted into the instruction stream for code that was heavily laden with arithmetic operations. Excess branching hinders performance by confounding processor pipelining in modern architectures [5]. Some of this overhead could probably be ameliorated by judiciously marking such branches using the GCC extension function `__builtin_expect`, which directs the compiler to insert a hint to the target processor's branch prediction unit that one end of the branch is unlikely to be taken. The primary problem with this approach is that while it does represent the direction in which the Python language is moving towards, Python is not there yet, so it would represent a divergence from the official language definition.

A related solution involves boxing integers directly, where all integers are represented by a machine word that contains either a pointer to a long integer, or a 31-bit integer. This representation uses one bit of the integer to determine whether the object is a machine integer or a pointer. While well accepted in the Smalltalk community, this approach has been harshly criticized in the Python community, making its introduction into Starkiller problematic. Much of the criticism has centered on difficulties achieving efficient implementations across many different processor architectures; the Alpha architecture in particular imposes a severe performance penalty on the bit twiddling operations required to support this approach. Many of the other concerns raised center around implementation complexity and would not be relevant to Starkiller.

An even greater divergence from the language specification is implied by the third option, namely keeping the original Python semantics of isolating machine from long integers. Machine integer overflow would result in an exception rather than silent promotion to long. Since Python's machine integers are signed, Starkiller can use the SIGFPE signal to detect overflows and raise exceptions without cluttering the instruction stream with conditional checks. This option offers the best performance while introducing significant semantic differences compared to the Python language specification.

References

- [1] Ole Agesen. *Concrete Type Inference: Delivering Object-Oriented Applications*. PhD thesis, Stanford University, 1996.
- [2] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley Professional Computing Series. Addison-Wesley Publishing Company, New York, NY, 1995.
- [3] Jean Goubault. Generalized boxings, congruences and partial inlining.
- [4] Ralph E. Griswold and Madge T. Griswold. *The Icon Programming Language*. Prentice Hall, 1990.
- [5] John L. Hennessey and David A. Patterson. *Computer Organization and Design: The Hardware/Software Interface*. Morgan Kaufmann, 1998.

- [6] Xavier Leroy. The effectiveness of type-based unboxing. In *Workshop Types in Compilation '97*. Technical report BCCS-97-03, Boston College, Computer Science Department, June 1997.
- [7] Stephan Murer, Stephen Omohundro, David Stoutamire, and Clemens Szyperski. Iteration abstraction in sather. *ACM Trans. Program. Lang. Syst.*, 18(1):1–15, 1996.
- [8] Neil Schemenauer, Tim Peters, and Magnus Lie Hetland. Simple generators. Python Enhancement Proposal 255, May 2001.
- [9] Peter J. Thiemann. Unboxed values and polymorphic typing revisited. In *Proceedings of the seventh international conference on Functional programming languages and computer architecture*, pages 24–35. ACM Press, 1995.
- [10] Guido van Rossum and Francis L. Drake, editors. *Python Language Reference*. Python-Labs, 2003.
- [11] Ka-Ping Yee and Guido van Rossum. Iterators. Python Enhancement Proposal 234, January 2001.
- [12] Moshe Zadka and Guido van Rossum. Unifying long integers and integers. Python Enhancement Proposal 237, March 2001.