Garbage collection in Python: things you need to know

Last updated on August 10, 2020, in python

This article describes garbage collection (GC) in Python 3.7.

Usually, you do not need to worry about memory management. When objects are no longer needed, Python automatically reclaims memory from them. However, understanding how GC works can help you write better and faster Python programs.

Memory management

Unlike many other languages, Python does not necessarily release the memory back to the Operating System. Instead, it has a dedicated object allocator for objects smaller than 512 bytes, which keeps some chunks of already allocated memory for further use in the future. The amount of memory that Python holds depends on the usage patterns. In some cases, all allocated memory could be released only when a Python process terminates.

If a long-running Python process takes more memory over time, it does not necessarily mean that you have memory leaks. If you are interested in Python's memory model, you can read my article on memory management.

Since most objects are small, custom memory allocator saves a lot of time on memory allocations. Even simple programs that import third-party libraries can allocate millions of objects during the program lifetime.

Garbage collection algorithms

In Python, everything is an object. Even integers. Knowing when to allocate them is easy. Python does it when you need to create a new object. Unlike allocation, automatic deallocation is tricky. Python needs to know when your object is no longer needed. Removing objects prematurely will result in a program crash.

program crash.

Garbage collections algorithms track which objects can be deallocated and pick an optimal time to deallocate them. Standard CPython's garbage collector has two components, the reference counting collector and the generational **garbage collector**, known as gc module.

The **reference counting** algorithm is incredibly efficient and straightforward, but it cannot detect reference cycles. That is why Python has a supplemental algorithm called generational cyclic GC. It deals with reference cycles only.

The reference counting module is fundamental to Python and can't be disabled, whereas the cyclic GC is optional and can be triggered manually.

Reference counting

Reference counting is a simple technique in which objects are deallocated when there is no reference to them in a program.

Every variable in Python is a reference (a pointer) to an object and not the actual value itself. For example, the assignment statement just adds a new reference to the right-hand side. A single object can have many references (variable names).

This code creates two references to a single object:

```
a = [1, 2, 3]
b = a
```

An assignment statement itself (everything on the left) never copies or creates new data.

To keep track of references, every object (even integer) has an extra field called reference count that is increased or decreased when a pointer to the object is created or deleted. See Objects, Types and Reference Counts section, for a detailed explanation.

EXAMPLES, WHERE THE REFERENCE COUNT INCREASES:

- assignment operator
- argument passing
- appending an object to a list (object's reference count will be increased).

If the reference counting field reaches zero, CPython automatically calls the object-specific memory deallocation function. If an object contains references to other objects, then their reference count is automatically decremented too. Thus other objects may be deallocated in turn. For example, when a list is deleted, the reference count for all its items is decreased. If another variable references an item in a list, the item won't be deallocated.

Variables, which are declared outside of functions, classes, and blocks, are called globals. Usually, such variables live until the end of the Python's process. Thus, the reference count of objects, which are referred by global variables, never drops to zero. To keep them alive, all globals are stored inside a dictionary. You can get it by calling the `globals()` function.

Variables, which are defined inside blocks (e.g., in a function or class) have a local scope (i.e., they are local to its block). When Python interpreter exits from a block, it destroys local variables and their references that were created inside the block. In other words, it only destroys the **names**.

It's important to understand that until your program stays in a block, Python interpreter assumes that all variables inside it are in use. To remove something from memory, you need to either assign a new value to a variable or exit from a block of code. In Python, the most popular block of code is a function; this is where most of the garbage collection happens. That is another reason to keep functions small and simple.

You can always check the number of current references using `sys.getrefcount` function.

Here is a simple example:

```
import sys

foo = []

# 2 references, 1 from the foo var and 1 from getrefcount
print(sys.getrefcount(foo))
```

```
def bar(a):
    # 4 references
    # from the foo var, function argument, getrefcount and Python's function stack
    print(sys.getrefcount(a))


bar(foo)
# 2 references, the function scope is destroyed
print(sys.getrefcount(foo))
```

IIn the example above, you can see that function's references get destroyed after Python exits it.

Sometimes you need to remove a global or a local variable prematurely. To do so, you can use the `del` statement that removes a variable and its reference (not the object itself). This is often useful when working in Jupyter notebooks because all cell variables use the global scope.

The main reason why CPython uses reference counting is historical. There are a lot of debates nowadays about the weaknesses of such a technique. Some people claim that modern garbage collection algorithms can be more efficient without reference counting at all. The reference counting algorithm has a lot of issues, such as circular references, thread locking, and memory and performance overhead. Reference counting is one of the reasons why Python can't get rid of the GIL.

The main advantage of such an approach is that the objects can be immediately and easily destroyed after they are no longer needed.

Generational garbage collector

Why do we need additional garbage collector when we have reference counting?

Unfortunately, classical reference counting has a fundamental problem — it cannot detect reference cycles. A reference cycle occurs when one or more objects are referencing each other.

Here are two examples:

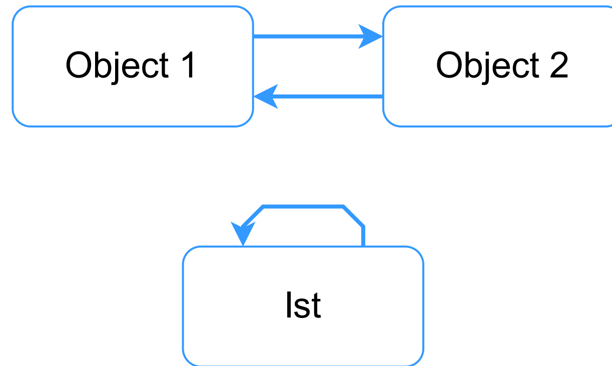

As we can see, the 'lst' object is pointing to itself, moreover, `object 1` and `object 2` are pointing to each other. The reference count for such objects is always at least 1.

To get a better idea, you can play with a simple Python example:

```
import gc

# We use ctypes moule  to access our unreachable objects by memory address.
class PyObject(ctypes.Structure):
    _fields_ = [("refcnt", ctypes.c_long)]


gc.disable()  # Disable generational gc

lst = []
lst.append(lst)

# Store address of the list
lst_address = id(lst)

# Destroy the lst reference
del lst

object_1 = {}
object_2 = {}
object_1['obj2'] = object_2
object_2['obj1'] = object_1

obj_address = id(object_1)

# Destroy references
del object_1, object_2

# Uncomment if you want to manually run garbage collection process
# gc.collect()

# Check the reference count
print(PyObject.from_address(obj_address).refcnt)
print(PyObject.from_address(lst_address).refcnt)
```

In the example above, the `del` statement removes the references to our objects (i.e., decreases reference count by 1). After Python executes the `del` statement, our objects are no longer accessible from Python code. However, such objects are still sitting in memory. That happens because they are still referencing each other, and the reference count of each object is 1. You can visually explore such relations using objgraph module.

To resolve this issue, the additional cycle-detecting algorithm was introduced in Python 1.5. The gc module is responsible for this and exists only for dealing with such a problem.

Reference cycles can only occur in container objects (i.e., in objects that can contain other objects), such as lists, dictionaries, classes, tuples. The garbage collector algorithm does not track all immutable types except for a tuple. Tuples and dictionaries containing only immutable objects can also be untracked depending on certain conditions. Thus, the reference counting technique handles all non-circular references.

When does the generational GC trigger

Unlike reference counting, cyclic GC does not work in real-time and runs periodically. To reduce the frequency of GC calls and micro pauses CPython uses various heuristics.

The GC classifies container objects into three generations. Every new object starts in the first generation. If an object survives a garbage collection round, it moves to the older (higher) generation. Lower generations are collected more often than higher. Because most of the newly created objects die young, it improves GC performance and reduces the GC pause time.

In order to decide when to run, each generation has an individual counter and threshold. The counter stores the number of object allocations minus deallocations since the last collection. Every time you allocate a new container object, CPython checks whenever the counter of the first generation exceeds the threshold value. If so, Python initiates the collection process.

If we have two or more generations that currently exceed the threshold, GC

chooses the oldest one. That is because the oldest generations are also

collecting all previous (younger) generations. To reduce performance degradation for long-living objects, the third generation has additional requirements in order to be chosen.

The standard threshold values are set to (700, 10, 10) respectively, but you can always check them using the `gc.get_threshold` function. You can also adjust them for your particular workload by using the `gc.get_threshold` function.

How to find reference cycles

It is hard to explain the reference cycle detection algorithm in a few paragraphs. Basically, GC iterates over each container object and temporarily removes all references to all container objects it references. After full iteration, all objects which reference count lower than two are unreachable from Python's code and thus can be collected.

To fully understand the cycle-finding algorithm, I recommend you to read an original proposal from Neil Schemenauer and collect function from CPython's source code. Also, the Quora answers and The Garbage Collector blog post can be helpful.

Note that, the problem with finalizers, which was described in the original proposal, has been fixed since Python 3.4. You can read about it in the PEP 442.

Performance tips

Cycles can easily happen in real life. Typically you encounter them in graphs, linked lists or in structures, in which you need to keep track of relations between objects. If your program has an intensive workload and requires low latency, you need to avoid reference cycles as possible.

To avoid circular references in your code, you can use weak references, that are implemented in the `weakref` module. Unlike the usual references, the `weakref.ref` doesn't increase the reference count and returns `None` if an object was destroyed.

In some cases, it is useful to disable GC and use it manually. The automatic collection can be disabled by calling `gc.disable()`. To manually run the collection process, you need to use `gc.collect()`.

How to find and debug reference cycles

Debugging reference cycles can be very frustrating especially when you use a lot of third-party libraries.

The standard gc module provides a lot of useful helpers that can help in debugging. If you set debugging flags to `DEBUG_SAVEALL`, all unreachable objects found will be appended to `gc.garbage` list.

```
import gc

gc.set_debug(gc.DEBUG_SAVEALL)

print(gc.get_count())
lst = []
lst.append(lst)
list_id = id(lst)
del lst
gc.collect()
for item in gc.garbage:
    print(item)
    assert list_id == id(item)
```

Once you have identified a problematic spot in your code you can visually explore object's relations using objgraph.

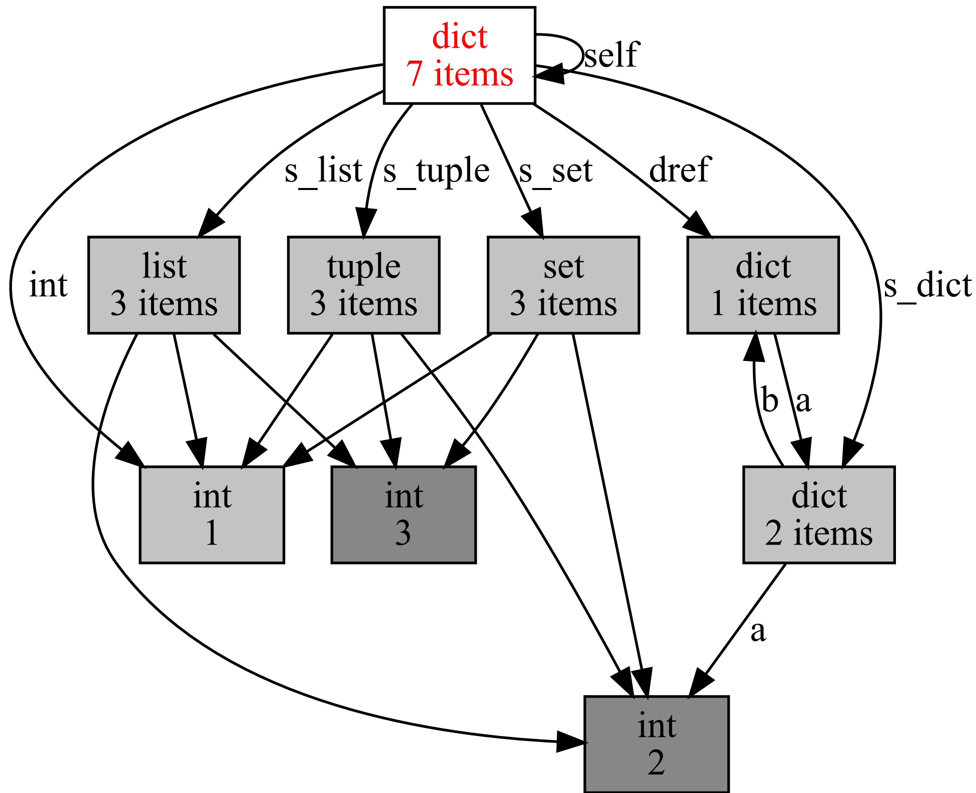

Conclusion

Most of the garbage collection is done by reference counting algorithm, which we cannot tune at all. So, be aware of implementation specifics, but don't worry about potential GC problems prematurely.

Hopefully, you have learned something new. If you have any questions left, I will be glad to answer them in the comments below.

> Popular posts in Python category

python, cpython internals, memory, advanced python

Want a monthly digest of these blog posts?

Email

Subscribe

No spam. No unnecessary emails.

Comments

Bob Hyman 3 years, 5 months ago (from disqus) #

Very interesting and useful description of the current state of affairs in Cython.

Are things different in other pythons? E.g, IPython.

Are there any projects underway to improve memory MGMT in python?

reply

> Artem 3 years, 5 months ago (from disqus) #
>
> Each implementation of Python uses its own collector. For example, Jyton uses standard Java's gc (since it running on the JVM) , and PyPy uses Mark and Sweep algorithm. The PyPy's gc is more complicated than CPython's and has additional optimizations http://doc.pypy.org/en/release-2.4.x/garbage_collection.html.
>
> I'm not aware of any changes in IPython since it's just an interactive shell running on CPython.
>
> Regarding memory management, there are tens of proposals in the PEP index, but only a few of them will be accepted in future.
>
> reply

Madison 3 years, 5 months ago (from disqus) #

Nice article. One small correction worth making is that not all block statements introduce a new execution scope. In particular a `with` statement does not have its own local scope any more than an `if` statement does. For this most part this only applies to `def` and `class` statements.

reply

> Artem 3 years, 5 months ago (from disqus) #
>
> Thanks, you are right. I was thinking about context's enter/exit.

> reply

ben 2 years, 1 month ago (from disqus) #

Thanks for the explanation. Can you please confirm then, that it would be pointless to import gc into jupyter notebook?

reply

Alice 8 months, 2 weeks ago #

Good write up.

reply

Simon 7 months, 3 weeks ago #

I don't understand why when you pass a variable to a function it has two more reference counts instead of one. You said that it has one from the function argument and one from Python's function stack. I don't know what is this function stack. Otherwise this is a great article!

reply

> Artem 7 months, 3 weeks ago #
>
> Function stack keeps track of all local variables that are avaliable inside the function.
>
> reply

kanch 7 months ago #

Hi there, your article help me a lot in reduce my python application memory usage.

Can I translate it into Chinese , and re-post with reference to my blog? I think it will help more people struggling with python memory usage in long run applications.

reply

> Artem 7 months ago #
>
> Sure, you can translate it. Just link my article in your translation.
>
> reply

jbo 5 months, 2 weeks ago #

Very well explained, just loved reading it and your other articles. Keep up the good work .

reply

ferdizera 3 months, 3 weeks ago #

Simple and objective. Nice article. Nice explanation.

reply

jb 1 month, 1 week ago #

Thank you, clear and useful

reply

Lokumcu Hayri 3 weeks, 6 days ago #

No lollygagging, just to-the-point explanations and examples. Thanks!

reply

NAME

MESSAGE

Plain text or markdown

Send

Back to top

© 2009-2020, Artem Golubin, me@rushter.com