

# CS107 / AC207

## SYSTEMS DEVELOPMENT FOR COMPUTATIONAL SCIENCE

### LECTURE 19

Tuesday, November 9th 2021

*Fabian Wermelinger*

Harvard University

# RECAP OF LAST TIME

- Generators
- Coroutines
- python internals: objects, bytecode and interpreter

# OUTLINE

- python internals:
  - Code objects
  - The interpreter and the evaluation loop
  - Frame objects
  - Generator objects
  - Traceback objects
- python lists and numpy arrays

# python INTERNALS: INTERPRETER

- All the data stored in a python program is built around the concept of an *object*. Code objects are *compiled bytecode*. The interpreter turns those code objects into *frame objects* and executes them (the *left column* you saw in [pythontutor](#)).
- What is still missing to execute these frame objects is *input* data (the *right column* you saw in [pythontutor](#)).
- The python interpreter obtains input data from a *value stack* and executes frame objects arranged in a *frame stack* in a central loop called the *evaluation loop*. In the *interactive* python shell this is called **REPL**: **R**ead, **E**valuate, **P**rint, **L**oop. The python interpreter is written in C (you can inspect the source code at <https://github.com/python/cpython>). At the very core of the evaluation loop is the `_PyEval_EvalFrameDefault` function. This is the function that brings everything together and makes your code come to life. *Everything that is executed in python must go through this function.*

# python INTERNALS: CODE OBJECTS

- Code objects is what the python interpreter executes. They represent raw bytecode.

- We can generate code objects with the `compile()` built-in function:

```
1 >>> a = 1
2 >>> co = compile('a + 1', '<string>', mode='eval')
3 >>> eval(co) # evaluate the code object
4 2
```

- The *raw bytecode* is contained in `co_code` :

```
1 >>> co.co_code
2 b'e\x00d\x00\x17\x00S\x00'
```

- We can *disassemble* bytecode into the *instructions* that python executes for a particular code object:

```
1 >>> import dis
2 >>> dis.dis(co)
3     1           0 LOAD_NAME           0 (a)
4     2           2 LOAD_CONST        0 (1)
5     4           4 BINARY_ADD
6     6           6 RETURN_VALUE
```

4 instructions are executed: 2 loads, 1 binary addition and returning the result.  
[A list of all bytecode instructions can be found here.](#)

# python INTERNALS: INTERPRETER

- We saw that a `LOAD_NAME` instruction pushes the object at given index in `co_names` onto the *value stack*. This instruction obtains the object from the *local scope*. The `LOAD_GLOBAL` instruction would be used to load a name from the *global scope*.
- In python, the *local* and *global* scopes can be inspected with the `locals()` and `globals()` built-ins, respectively:

```
1 >>> def f(x):
2 ...     l = x
3 ...     print(f'f() local variables: {locals()}')
4 ...     print(f'f() global variables: {globals()}')
5 ...
6 >>> g = 0
7 >>> f(g)
8 f() local variables: {'x': 0, 'l': 0}
9 f() global variables: {'__name__': '__main__', '__doc__': None, '__package__': None,
10 '__loader__': <class '_frozen_importlib.BuiltinImporter'>, '__spec__': None,
11 '__annotations__': {}, '__builtins__': <module 'builtins' (built-in)>,
12 'f': <function f at 0x7f6868f8f280>, 'g': 0}
```

# python INTERNALS: INTERPRETER

*Short detour:* low-level explanation of two things we have seen previously in the lecture: 1.) the `nonlocal` keyword

```
1 def f(x):
2     def g(y):
3         z = x + y
4         return z
5     return g
```

Read-only x!

```
1 >>> g = f(0)
2 >>> f.__code__.co_cellvars # tuple of vars referenced in nested functions
3 ('x',)
4 >>> g.__closure__
5 (<cell at 0x7f90da60fc10: int object at 0x7f90da769910>,)
6 >>> g.__closure__[0].cell_contents
7 0
8 >>> g(1)
9 1
10 >>> g.__closure__[0].cell_contents
11 0
```

```
1 def f(x):
2     def g(y):
3         z = x + y
4         x = y
5         return z
6     return g
```

```
1 >>> g = f(0)
2 >>> f.__code__.co_cellvars # tuple of vars referenced in nested functions
3 ()
4 >>> g.__closure__ # is None -- no cellvars in code object this time
5 >>> g(1)
6 Traceback (most recent call last):
7   File "<stdin>", line 1, in <module>
8   File "<stdin>", line 3, in g
9 UnboundLocalError: local variable 'x' referenced before assignment
10 >>> import dis; dis.dis(g)
11      3           0 LOAD_FAST           1 (x)
12           2 LOAD_FAST           0 (y)
13           4 BINARY_ADD
14           6 STORE_FAST          2 (z)
15 ... # more assembly ignored
```

The `UnboundLocalError` happens in line 3 (first access) because it is not defined in the scope!

# python INTERNALS: INTERPRETER

*Short detour:* low-level explanation of two things we have seen previously in the lecture: 1.) the `nonlocal` keyword

```
1 def f(x):
2     def g(y):
3         z = x + y
4         return z
5     return g
```

Read-only x!

```
1 >>> g = f(0)
2 >>> f.__code__.co_cellvars # tuple of vars referenced in nested functions
3 ('x',)
4 >>> g.__closure__
5 (<cell at 0x7f90da60fc10: int object at 0x7f90da769910>,)
6 >>> g.__closure__[0].cell_contents
7 0
8 >>> g(1)
9 1
10 >>> g.__closure__[0].cell_contents
11 0
```

```
1 def f(x):
2     def g(y):
3         x = y
4         z = x + y
5         return z
6     return g
```



Will this result in an `UnboundLocalError`



# python INTERNALS: INTERPRETER

*Short detour:* low-level explanation of two things we have seen previously in the lecture: 1.) the `nonlocal` keyword

```
1 def f(x):
2     def g(y):
3         nonlocal x
4         z = x + y
5         x = y
6         return z
7     return g
```

Read-Write captured x !

```
1 >>> g = f(0)
2 >>> f.__code__.co_cellvars # tuple of vars referenced in nested functions
3 ('x',)
4 >>> g.__closure__
5 (<cell at 0x7f7bdcf37a60: int object at 0x7f7bdd0a2910>,)
6 >>> g.__closure__[0].cell_contents
7 0
8 >>> g(1)
9 1
10 >>> g.__closure__[0].cell_contents
11 1
```

```
1 def f(x):
2     def g(y):
3         z = x + y
4         x = y
5         return z
6     return g
```

The `UnboundLocalError` happens in line 3 (first access) because it is not defined in the scope!

```
1 >>> g = f(0)
2 >>> f.__code__.co_cellvars # tuple of vars referenced in nested functions
3 ()
4 >>> g.__closure__ # is None -- no cellvars in code object this time
5 >>> g(1)
6 Traceback (most recent call last):
7   File "<stdin>", line 1, in <module>
8   File "<stdin>", line 3, in g
9 UnboundLocalError: local variable 'x' referenced before assignment
10 >>> import dis; dis.dis(g)
11      3           0 LOAD_FAST           1 (x)
12           2 LOAD_FAST           0 (y)
13           4 BINARY_ADD
14           6 STORE_FAST          2 (z)
15 ... # more assembly ignored
```

# python INTERNALS: INTERPRETER

*Short detour:* low-level explanation of two things we have seen previously in the lecture: 2.) positional and keyword only arguments

```
1 >>> def f(posonly, /, pos_or_kw=None, *, kwonly=None):
2     ...     pass
3     ...
4 >>> f.__code__.co_argcount
5 2
6 >>> f.__code__.co_posonlyargcount
7 1
8 >>> f.__code__.co_kwonlyargcount
9 1
```

- The compiled bytecode is aware of what arguments are expected.
- `co_argcount` does not include keyword-only arguments (<https://docs.python.org/3/library/inspect.html#types-and-members>)
- See [PEP 457](#) and [PEP 570](#) for the introduction of positional-only arguments and the [control flow tutorial](#) for more discussion.

# python INTERNALS: INTERPRETER

## *Important terms for the python interpreter:*

- The evaluation loop (or REPL in the interactive python shell) will take a *code object* and convert it into a series of *frame objects*.
- Frame objects are executed in a so called *frame stack* (what we saw in [pythontutor](#)).
- The interpreter manages referenced variables in a *value stack*.
- The interpreter has at least one thread but *at most one thread can run at a time*. The python interpreter uses an internal global interpreter lock (called **GIL**) which prevents [race conditions](#) and ensures thread safety. The GIL imposes a very strong constraint on multi-threaded execution and was subject to many discussions in the past. [A recent post \(10/07/2021\) on the python-dev mailing list](#) proposes a new design to remove the GIL which would mean a major change in the python interpreter, a possible change that will be reality in the next python 4 release.
- **Did you know:** for the first time in 20 years, [python became the worlds most popular programming language](#) this year.

# python INTERNALS: BACK TO OBJECTS

*Repeat:* everything in python is an object!

## Fixed size object base:

```
1 typedef struct _object {
2     _PyObject_HEAD_EXTRA
3     Py_ssize_t ob_refcnt;
4     PyTypeObject *ob_type;
5 } PyObject; // C code in cpython
```

- `_PyObject_HEAD_EXTRA` is a macro that is usually empty.
- `ob_refcnt` is the *reference* count for the object.
- `ob_type` is a pointer to the type object. Recall that python is *dynamically typed*.

## Variable size object base:

```
1 typedef struct {
2     PyObject ob_base;
3     Py_ssize_t ob_size;
4 } PyVarObject; // C code in cpython
```

- `ob_base` is a fixed size object instance.
- `ob_size` is the number of items in the variable part.
- Containers (e.g. `list`) are objects of this type.

No object in python is a direct instance of `PyObject`. BUT, every object in python can be cast to a `PyObject`; if it is variable size it can be cast to `PyVarObject` in addition.

# python INTERNALS: FRAME OBJECTS

A frame object is a PyObject with the following additional properties:

The PyFrameObject:

```
1 struct _frame {
2     PyObject ob_base;
3     struct _frame *f_back;
4     struct _interpreter_frame *f_frame;
5     PyObject *f_trace;
6     int f_lineno;
7     char f_trace_lines;
8     char f_trace_opcodes;
9     char f_own_locals_memory;
10 } PyFrameObject;
```

- `ob_base` is the base instance (as before).
- `f_back` is a pointer to the previous PyFrameObject towards the caller (enables the frame stack).
- `f_frame` is a pointer to the frame data.
- Other fields are used for debugging.

# python INTERNALS: FRAME OBJECTS

A frame object is a PyObject with the following additional properties:

The PyFrameObject:

```
1 struct _frame {
2     PyObject ob_base;
3     struct _frame *f_back;
4     struct _interpreter_frame *f_frame;
5     PyObject *f_trace;
6     int f_lineno;
7     char f_trace_lines;
8     char f_trace_opcodes;
9     char f_own_locals_memory;
10 } PyFrameObject;
```

- `ob_base` is the base instance (as before).
- `f_back` is a pointer to the previous PyFrameObject towards the caller (enables the frame stack).
- `f_frame` is a pointer to the frame data.
- Other fields are used for debugging.

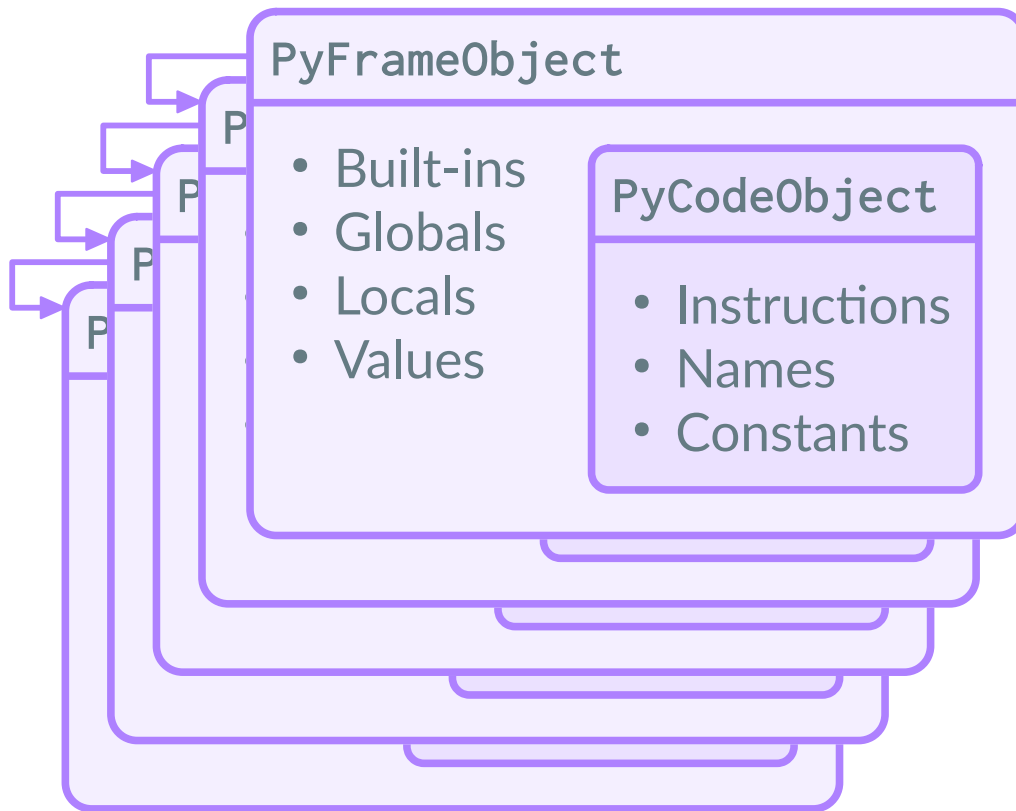
The frame data (some code not shown):

```
1 typedef struct _interpreter_frame {
2     PyObject *f_globals;
3     PyObject *f_builtins;
4     PyObject *f_locals;
5     PyCodeObject *f_code;
6     PyFrameObject *frame_obj;
7     PyObject *generator;
8     int f_lasti;
9     int depth;
10 } InterpreterFrame;
```

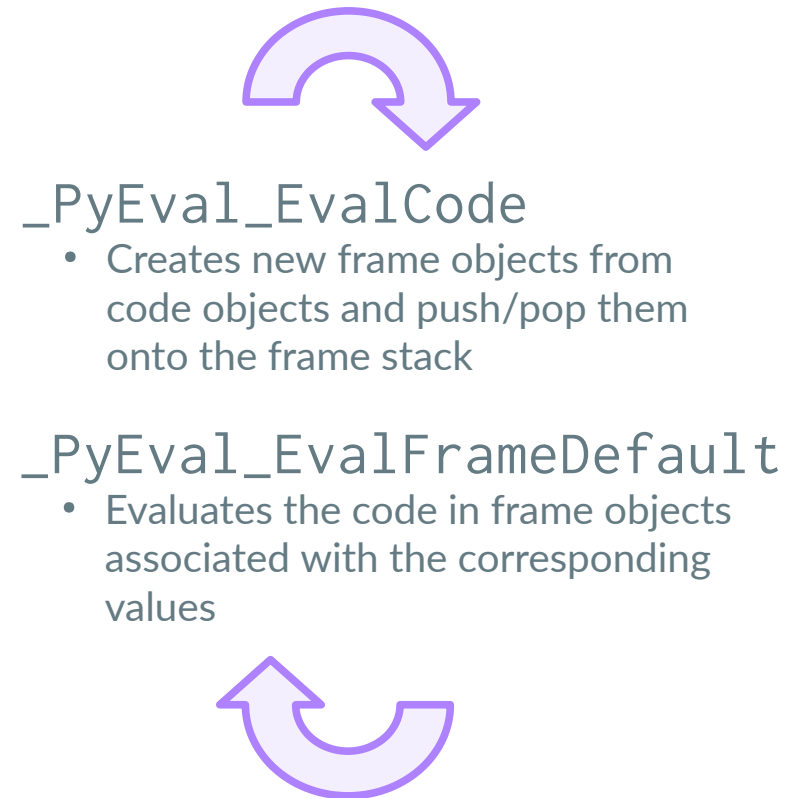
- *Is not an object* (has no `ob_base`).
- `f_globals` and `f_locals` point to data.
- `f_code` is the bytecode object that will be executed by the frame.
- `f_lasti` index of the last instruction executed. *Where is this index used?*

# python INTERNALS: FRAME OBJECTS

## Frame Stack



## Evaluation Loop



<https://github.com/python/cpython>

# python INTERNALS: GENERATOR OBJECTS

The `PyGenObject` (some code not shown):

```
1 typedef struct {
2     /* The gi_ prefix is intended to
3     remind of generator-iterator. */
4     PyObject ob_base;
5     /* Note: gi_frame can be NULL if
6     the generator is "finished" */
7     struct _interpreter_frame *gi_xframe;
8     /* The code object backing
9     the generator */
10    PyCodeObject *gi_code;
11 } PyGenObject;
```

- `gi_xframe` points to the current frame object for the generator.
- `gi_code` bytecompiled code object of the *generator function*.
- `PyGenObject` 's are flagged when created with `CO_GENERATOR` (this class), `CO_COROUTINE` (PEP 492) or `CO_ASYNC_GENERATOR` (PEP 525).

- Frame objects have a pointer to generator objects and they store the index of the last instruction in the bytecode of the frame.
- This allows to *resume* a generator object that is associated with a frame. **Why?** The frame data has this code:

```
1 PyObject *generator;
```

This is a *pointer* to a `PyObject` "somewhere" in memory. **Aside:** pointers in C/C++ are used for *dynamic memory management*.

- Although the frame object is in the stack, the generator pointer allows to obtain the generator object from somewhere else in memory which can then be resumed (e.g. the frame evaluates `next()` on a generator).
- In fact: python 's frame stack is maintained in *dynamic memory* (*heap*), something that is not true for standard program execution.



# python INTERNALS: DYNAMIC MEMORY

- Processes share CPU and memory among each other.
- Sharing memory is a non-trivial task when designing operating systems.
- Each physical memory cell (byte-sized cells) can be addressed uniquely. For example:
  - **32-bit system:** 4294967296 addresses; can handle 4GB (gigabyte) of RAM at most.
  - **64-bit system:** 18446744073709551616 addresses; can handle 16EB (exabyte) of RAM at most. (This is A LOT!)
- **Virtual memory** simplifies memory management in operating systems by making processes "think" that their memory space always starts at address 0. The **virtual address** is then translated to the real physical address by a hardware component called **memory management unit (MMU)**.

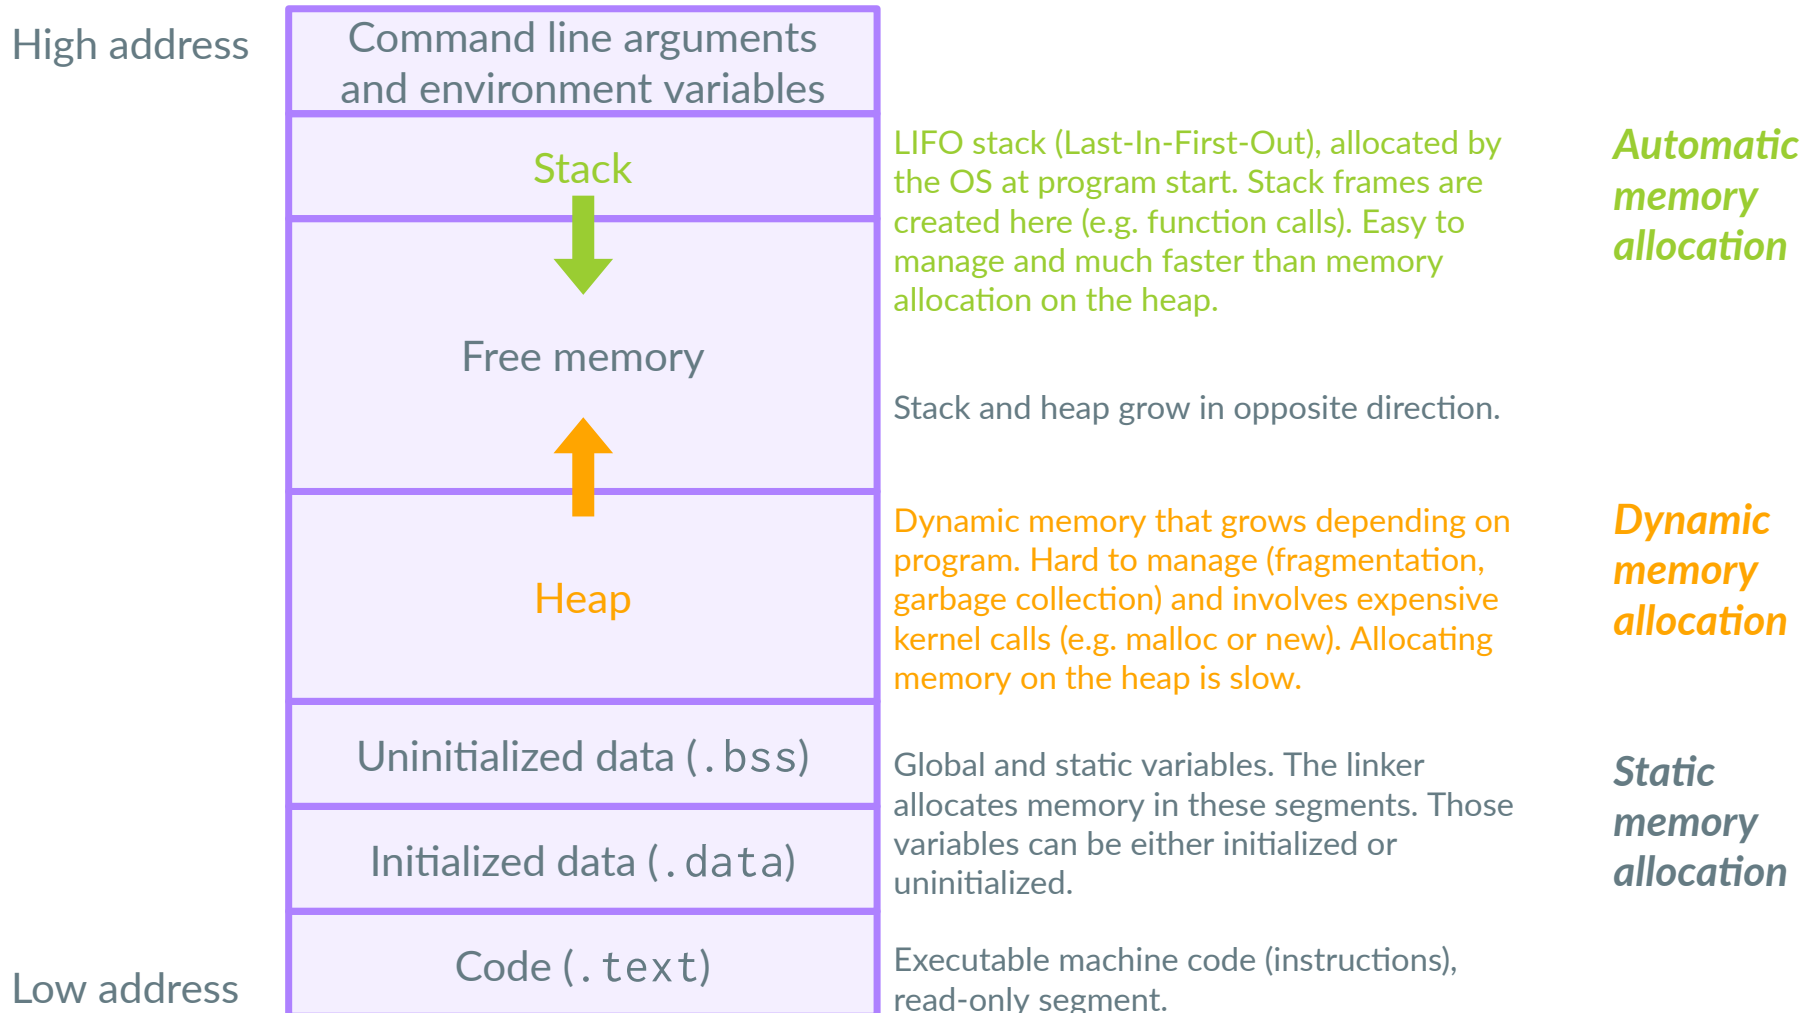
# python INTERNALS: DYNAMIC MEMORY

- **Static memory allocation:** variables with known size at compile time. The compiler allocates this memory inside the executable.
- **Automatic memory allocation:** similar to static memory allocation; the allocation requirements are known at compile time. Allocation is carried out on the stack when the code executes.
- **Dynamic memory allocation:** when it is not possible for the compiler to determine a specific memory request, e.g. the allocation size depends on user input at runtime, the memory must be allocated dynamically. A dynamic memory segment is allocated on the heap.

All objects in python (including frame objects and code objects) are *allocated dynamically on the heap*.

# python INTERNALS: DYNAMIC MEMORY

## Virtual memory of a Linux process



# python INTERNALS: DYNAMIC MEMORY

- The python interpreter *emulates* a frame stack using *dynamic memory*. Frame objects are pushed and popped to and from the frame stack on the *heap* (dynamic memory pool).
- These stack operations are more expensive than the ones used with *automatic memory allocation* in a x86\_64 executable.
- Since all python objects are allocated on the heap, *generator objects persist* until the interpreter explicitly removes them from the heap. This allows to easily resume a suspended generator including its state. Because a PyFrameObject stores the last instruction in *f\_lasti*, it will be used to index into the bytecode of a generator object to resume execution with instruction  $f\_lasti + 1$ .

# python INTERNALS: TRACEBACK OBJECTS

- The python interpreter exposes a number of internal objects to the user of which we have discussed three so far:
  - Code objects for byte compiled code
  - Frame objects to execute code.
  - Generator objects for suspension and resumption of code execution.
- It is rare that you will need to manipulate these objects directly in your code. We have used them here to understand the low-level python internals without going too deep into the interpreter source code.
- The last python internal object we want to look at are *traceback objects*. They are created when *exceptions* are raised and used for debugging purposes or non-standard exception handling.

# python INTERNALS: TRACEBACK OBJECTS

Obtain a traceback object from an exception:

```
1 import dis
2
3 def g():
4     raise Exception
5
6 def f():
7     g()
8
9 def main():
10    try:
11        f()
12    except Exception as e:
13        tb = e.__traceback__
14        i = 0
15        while tb is not None:
16            frame = tb.tb_frame
17            li = frame.f_code.co_code[frame.f_lasti]
18            print(f'frame {i}: line frame={frame.f_lineno}; ' +
19                f'line trace={tb.tb_lineno}; ' +
20                f'last instruction={dis.opname[li]}')
21            tb = tb.tb_next
22            i += 1
```

Output:

```
1 frame 0: line frame=18; line trace=11; last instruction=LOAD_ATTR
2 frame 1: line frame=7; line trace=7; last instruction=CALL_FUNCTION
3 frame 2: line frame=4; line trace=4; last instruction=RAISE_VARARGS
```

- Traceback objects can only be obtained through an exception.
- They are similar to frame objects, except that `tb_next` points *towards the frame where the exception has been thrown*.
- Control on the left is in the `main()` function after the exception was thrown. The line number in the frame and trace objects are different because we execute this frame.

# python INTERNALS: TRACEBACK OBJECTS

Standard traceback:

```
1 def g():
2     raise Exception("A useful description goes here")
3
4 def f():
5     g()
6
7 def main():
8     f()
9
10 if __name__ == "__main__":
11     main()
```

Output:

```
1 Traceback (most recent call last):
2   File "/home/fabs/CS107/traceback/tb.py", line 17, in <module>
3     main()
4   File "/home/fabs/CS107/traceback/tb.py", line 13, in main
5     f()
6   File "/home/fabs/CS107/traceback/tb.py", line 9, in f
7     g()
8   File "/home/fabs/CS107/traceback/tb.py", line 5, in g
9     raise Exception("A useful description goes here")
10 Exception: A useful description goes here
```

- Read python traceback from bottom up.
- Each line that starts with **File** corresponds to a new traceback object.
- The defaults contain information about the file where the code is executed, the current line and the name of the frame.
- Recall that code blocks in python are among the following: modules, function bodies, class definitions or commands typed interactively.

# python LIST OBJECTS AND NUMPY ARRAYS

- python has the reputation of being slow.
- This is true. It is not due to bad design however (the GIL is debatable), but rather prioritizing flexibility and the possibility for fast prototyping.
- One of the reasons for this performance penalty is that python objects are not necessarily near by in memory due to dynamic memory allocation.
- How does this matter since memory cells in random access memory (RAM) can be accessed in constant time?
  - Additional pointer dereferences until you get to the data. (Everything must be referenced by pointers.)
  - Spatial and temporal locality of the data is not optimal. Results in many cache misses when reading or writing data.



# python LIST OBJECTS AND NUMPY ARRAYS

- Let us see how list objects are implemented in python :

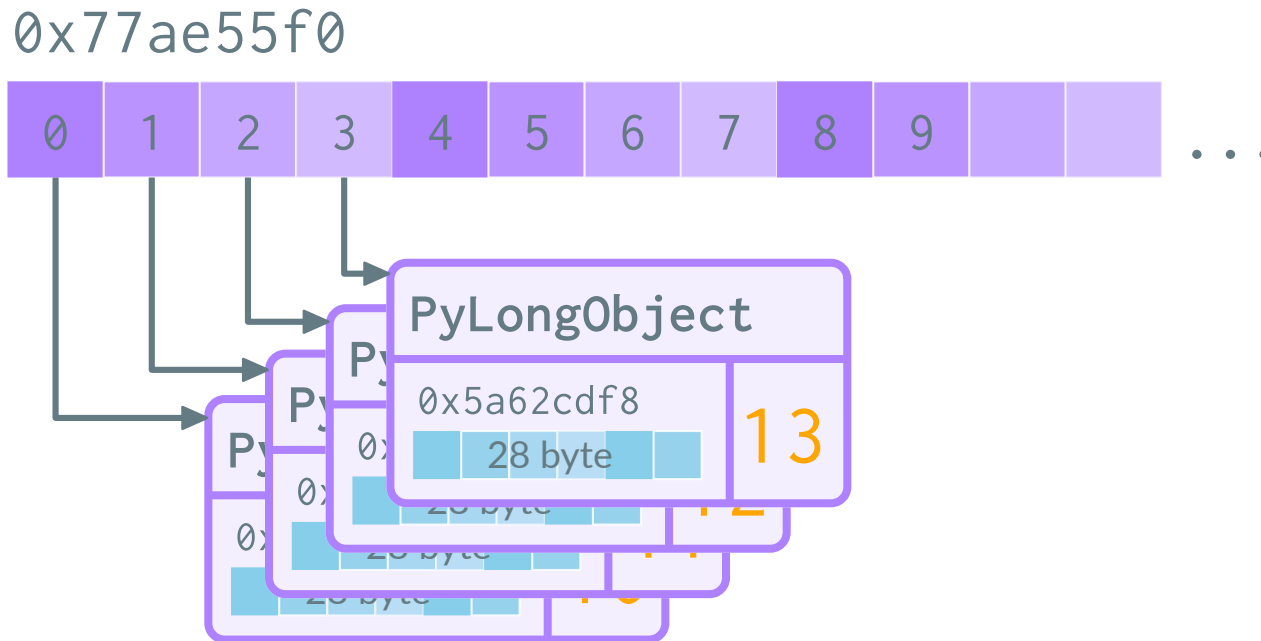
```
1 typedef struct {
2     PyVarObject ob_base;
3     // Vector of pointers to list elements. list[0] is ob_item[0], etc.
4     PyObject **ob_item;
5     Py_ssize_t allocated;
6 } PyListObject;
```

- `ob_item` is a pointer to pointer(s) to `PyObject` 's. For example, `ob_item[0]` returns a pointer to a `PyObject`, `ob_item[1]` returns the next pointer to the second `PyObject` and so on.
- In the following we assume `PyObject` represents a python integer:

```
1 struct _longobject {
2     PyObject ob_base;
3     Py_ssize_t ob_size;
4     digit ob_digit[1]; // the actual integer
5 } PyLongObject;
```

We assume that the `PyObject` takes 16 byte, `ob_size` is 8 byte and `ob_digit` is 4 byte. A `PyLongObject` then has a size of 28 byte.

# python LIST OBJECTS AND NUMPY ARRAYS



- Assume the following python list

```
1 li = [10, 11, 12, 13, 14, 15, 16, 17, 18, 19]
```

- The elements of `ob_item` are *coalesced* in memory. We can access the `PyObject` references in the list with  $\mathcal{O}(1)$  complexity.
- To obtain the actual value we must *dereference* the pointer and read `ob_digit[0]` for every item in the list.

# python LIST OBJECTS AND NUMPY ARRAYS

We can visualize this list on [pythontutor.com](http://pythontutor.com) showing all heap allocations:

The image shows a screenshot of the Python Tutor website. The main window displays the code `li = [10, 11, 12, 13, 14, 15, 16, 17, 18, ]` with a red arrow pointing to the first line, indicating it is the next line to execute. A green arrow points to the line below, indicating it has just been executed. Below the code is a horizontal scrollbar and two buttons labeled "< Prev" and "Next >". The text "Step 1 of 1" is displayed below the buttons. At the bottom left, there are links for "Rendered by Python Tutor" and "Customize visualization". To the right of the main window, there are two empty panels labeled "Frames" and "Objects".

Python 3.6

```
→ 1 li = [10, 11, 12, 13, 14, 15, 16, 17, 18, ]
```

→ line that just executed

→ next line to execute

< Prev Next >

Step 1 of 1

Rendered by [Python Tutor](http://pythontutor.com)  
[Customize visualization](#)

Frames Objects

# python LIST OBJECTS AND NUMPY ARRAYS

*Example:* for-loop over iterable

Sum values in iterable x :

```
1 import timeit
2 import numpy as np
3
4 def pysum(x):
5     s = 0
6     for i in x:
7         s += i
8
9 def main():
10    x = np.array(list(range(1000000)))
11    t = timeit.timeit('f(x)',
12                    globals={'f': pysum, 'x': x},
13                    number=10
14    )
15
16 if __name__ == "__main__":
17    main()
```

Assembly:

```
1 5          0 LOAD_CONST      1 (0)
2          2 STORE_FAST      1 (s)
3
4 6          4 LOAD_FAST      0 (x)
5          6 GET_ITER
6          >> 8 FOR_ITER        12 (to 22)
7          10 STORE_FAST     2 (i)
8
9 7          12 LOAD_FAST     1 (s)
10         14 LOAD_FAST     2 (i)
11         16 INPLACE_ADD
12         18 STORE_FAST     1 (s)
13         20 JUMP_ABSOLUTE 8
14         >> 22 LOAD_CONST     0 (None)
15         24 RETURN_VALUE
```

Running this code with 1'000'000 elements takes **0.74 seconds**, averaged over 10 samples.

# python LIST OBJECTS AND NUMPY ARRAYS

- While PyObject 's in python are generic in the sense that they are very flexible in terms of describing data, the generality comes at a performance price.
- The python interpreter is designed to work with PyObject 's exclusively (*recall*: everything in python is an *object*).
- Performance oriented designs are centered around *data* rather than objects.
- Because the python interpreter is written in C , extensions can easily be implemented.
- **NumPy** is a python extension module designed for efficient numerical computation in python .

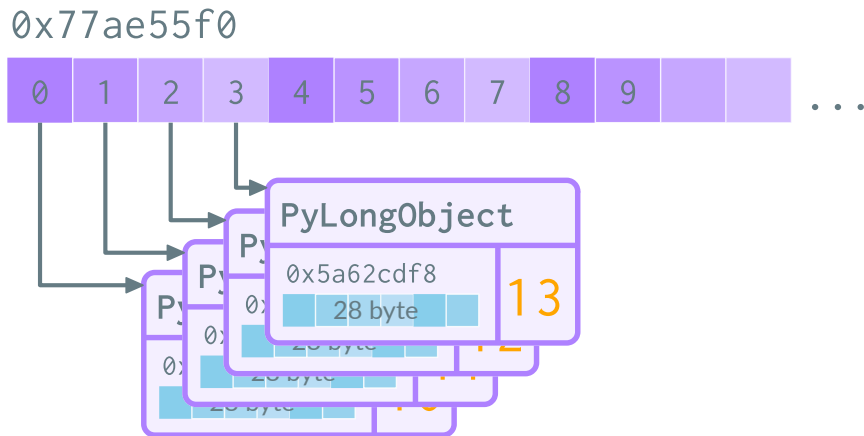
It operates on its own data structures for this reason.

# python LIST OBJECTS AND NUMPY ARRAYS

*Back to this list:*

```
1 li = [10, 11, 12, 13, 14, 15, 16, 17, 18, 19]
```

Object oriented  
python list:




- The list above is the same as

```
1 PyObject *ob_item[10]; // 10 contiguous pointer
```

- PyLongObject uses 32-bit integral type for integers, i.e., `int` in C.

Data oriented  
NumPy array:



- The elements of a NumPy array are **contiguous data items**, not pointers (references) to PyObject 's.
- The NumPy array above is similar to

```
1 int ob_item[10]; // 10 contiguous data items
```

- Reading the data in this format will saturate the memory bandwidth!

# python LIST OBJECTS AND NUMPY ARRAYS

*Example:* for -loop over iterable (same as before)

Sum values in iterable x :

```
1 import timeit
2 import numpy as np
3
4 def psum(x):
5     s = 0
6     for i in x:
7         s += i
8
9 def npsum(x):
10    s = x.sum()
11
12 def main():
13    x = np.array(list(range(1000000)))
14    t = timeit.timeit('f(x)',
15                    globals={'f': npsum, 'x': x},
16                    number=10
17    )
18
19 if __name__ == "__main__":
20    main()
```

Assembly:

```
1 10      0 LOAD_FAST      0 (x)
2         2 LOAD_METHOD  0 (sum)
3         4 CALL_METHOD   0
4         6 STORE_FAST   1 (s)
5         8 LOAD_CONST  0 (None)
6         10 RETURN_VALUE
```

Running this code with 1'000'000 elements averaged over 10 samples:

- Pure python : **0.74 seconds**
- NumPy array: **0.0046 seconds**

***Two orders of magnitude faster!***

**Note:** the `sum()` built-in function is only slightly faster (**0.60 seconds**) than the naive for -loop implementation.

# RECAP

- python internals:
  - Code objects
  - The interpreter and the evaluation loop
  - Frame objects
  - Generator objects
  - Traceback objects
- python lists and numpy arrays