



Unladen Swallow:
Fewer coconuts, faster Python

Collin Winter
collinwinter@google.com
unladen-swallow@googlegroups.com



Why is Google doing this?

- Lots of Python at Google.
 - Python one of Google's three primary languages.
 - Python enables fast development, rapid prototyping.
 - Engineers should use the language they want...
 - ...and it should be fast.
- Biggest user: YouTube
 - YouTube is pure-Python.
 - #2 search site on the Internet, behind google.com.



Original Goals

- Make Python 5x faster.
- Source-compatible with existing Python code.
- Source-compatible with existing C extension modules.
- Focus on ease of migration.
- Open-source everything, merge back into CPython.



Why is Python slow?

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

- Everything is an object.
- Everything is a method call, eventually.
- Ducktyping means we can't predict receiver types.



Why is Python slow?

```
>>> x = dict()  
>>> len(x)
```



Why is Python slow?

```
>>> x = dict()  
>>> len(x)
```

...len() calls PyObject_Size()



Why is Python slow?

```
>>> x = dict()  
>>> len(x)
```

...len() calls PyObject_Size()

...which looks up x->ob_type->tp_as_sequence->sq_length



Why is Python slow?

```
>>> x = dict()  
>>> len(x)
```

...len() calls PyObject_Size()

...which looks up x->ob_type->tp_as_sequence->sq_length

...which is NULL, so call PyMapping_Size()



Why is Python slow?

```
>>> x = dict()  
>>> len(x)
```

...len() calls PyObject_Size()

...which looks up x->ob_type->tp_as_sequence->sq_length

...which is NULL, so call PyMapping_Size()

...which looks up x->ob_type->tp_as_mapping->mp_length



Why is Python slow?

```
>>> x = dict()  
>>> len(x)
```

...len() calls PyObject_Size()
...which looks up x->ob_type->tp_as_sequence->sq_length
...which is NULL, so call PyMapping_Size()
...which looks up x->ob_type->tp_as_mapping->mp_length
...which is a function pointer to dict_length()



Why is Python slow?

```
>>> x = dict()  
>>> len(x)
```

...len() calls PyObject_Size()
...which looks up x->ob_type->tp_as_sequence->sq_length
...which is NULL, so call PyMapping_Size()
...which looks up x->ob_type->tp_as_mapping->mp_length
...which is a function pointer to dict_length()
...which returns x->ma_used



Why is Python slow?



Sunday, February 21, 2010

Very late-binding

Can't even make assumptions about builtin functions, types

Very hard to predict, have to measure

Why is Python slow?

```
def foo(x):  
    yield len(x)  
    yield len(x)
```

```
>>> g = foo(range(5))
```

```
>>> g.next()
```

```
5
```

```
>>> len = lambda y: 8
```

```
>>> g.next()
```



Sunday, February 21, 2010

Very late-binding

Can't even make assumptions about builtin functions, types

Very hard to predict, have to measure

Why is Python slow?

```
def foo(x):  
    yield len(x)  
    yield len(x)
```

```
>>> g = foo(range(5))
```

```
>>> g.next()
```

```
5
```

```
>>> len = lambda y: 8
```

```
>>> g.next()
```

```
8
```



Sunday, February 21, 2010

Very late-binding

Can't even make assumptions about builtin functions, types

Very hard to predict, have to measure

Unladen Swallow



Sunday, February 21, 2010

Prior work:

- Rubinius (Ruby)
- V8 (Chrome), TraceMonkey (Firefox), SquirrelFish Extreme (Safari)
- Pysco, PyPy
- HotSpot JVM, StrongTalk
- Self 93
- SmallTalk

Unladen Swallow

- Compile to machine code.



Sunday, February 21, 2010

Prior work:

- Rubinius (Ruby)
- V8 (Chrome), TraceMonkey (Firefox), SquirrelFish Extreme (Safari)
- Pysco, PyPy
- HotSpot JVM, StrongTalk
- Self 93
- SmallTalk

Unladen Swallow

- Compile to machine code.
- Assume the program is less dynamic than the language.



Sunday, February 21, 2010

Prior work:

- Rubinius (Ruby)
- V8 (Chrome), TraceMonkey (Firefox), SquirrelFish Extreme (Safari)
- Pysco, PyPy
- HotSpot JVM, StrongTalk
- Self 93
- SmallTalk

Unladen Swallow

- Compile to machine code.
- Assume the program is less dynamic than the language.
 - Globals/builtins aren't overridden once set.



Sunday, February 21, 2010

Prior work:

- Rubinius (Ruby)
- V8 (Chrome), TraceMonkey (Firefox), SquirrelFish Extreme (Safari)
- Pysco, PyPy
- HotSpot JVM, StrongTalk
- Self 93
- SmallTalk

Unladen Swallow

- Compile to machine code.
- Assume the program is less dynamic than the language.
 - Globals/builtins aren't overridden once set.
 - Classes never change.



Sunday, February 21, 2010

Prior work:

- Rubinius (Ruby)
- V8 (Chrome), TraceMonkey (Firefox), SquirrelFish Extreme (Safari)
- Pysco, PyPy
- HotSpot JVM, StrongTalk
- Self 93
- SmallTalk

Unladen Swallow

- Compile to machine code.
- Assume the program is less dynamic than the language.
 - Globals/builtins aren't overridden once set.
 - Classes never change.
 - Objects never change type.



Sunday, February 21, 2010

Prior work:

- Rubinius (Ruby)
- V8 (Chrome), TraceMonkey (Firefox), SquirrelFish Extreme (Safari)
- Pysco, PyPy
- HotSpot JVM, StrongTalk
- Self 93
- SmallTalk

Unladen Swallow

- Compile to machine code.
- Assume the program is less dynamic than the language.
 - Globals/builtins aren't overridden once set.
 - Classes never change.
 - Objects never change type.
 - You never debug/profile your code.



Sunday, February 21, 2010

Prior work:

- Rubinius (Ruby)
- V8 (Chrome), TraceMonkey (Firefox), SquirrelFish Extreme (Safari)
- Pysco, PyPy
- HotSpot JVM, StrongTalk
- Self 93
- SmallTalk

Unladen Swallow

- Compile to machine code.
- Assume the program is less dynamic than the language.
 - Globals/builtins aren't overridden once set.
 - Classes never change.
 - Objects never change type.
 - You never debug/profile your code.
 - Locals aren't deleted.



Sunday, February 21, 2010

Prior work:

- Rubinius (Ruby)
- V8 (Chrome), TraceMonkey (Firefox), SquirrelFish Extreme (Safari)
- Pysco, PyPy
- HotSpot JVM, StrongTalk
- Self 93
- SmallTalk

Unladen Swallow

- Compile to machine code.
- Assume the program is less dynamic than the language.
 - Globals/builtins aren't overridden once set.
 - Classes never change.
 - Objects never change type.
 - You never debug/profile your code.
 - Locals aren't deleted.
 - `sys._getframe()` doesn't matter.



Sunday, February 21, 2010

Prior work:

- Rubinius (Ruby)
- V8 (Chrome), TraceMonkey (Firefox), SquirrelFish Extreme (Safari)
- Pysco, PyPy
- HotSpot JVM, StrongTalk
- Self 93
- SmallTalk

Unladen Swallow

- Compile to machine code.
- Assume the program is less dynamic than the language.
 - Globals/builtins aren't overridden once set.
 - Classes never change.
 - Objects never change type.
 - You never debug/profile your code.
 - Locals aren't deleted.
 - `sys._getframe()` doesn't matter.
 - `exec` is never used.



Sunday, February 21, 2010

Prior work:

- Rubinius (Ruby)
- V8 (Chrome), TraceMonkey (Firefox), SquirrelFish Extreme (Safari)
- Pysco, PyPy
- HotSpot JVM, StrongTalk
- Self 93
- SmallTalk

Unladen Swallow

- Compile to machine code.
- Assume the program is less dynamic than the language.
 - Globals/builtins aren't overridden once set.
 - Classes never change.
 - Objects never change type.
 - You never debug/profile your code.
 - Locals aren't deleted.
 - `sys._getframe()` doesn't matter.
 - `exec` is never used.
 - ...

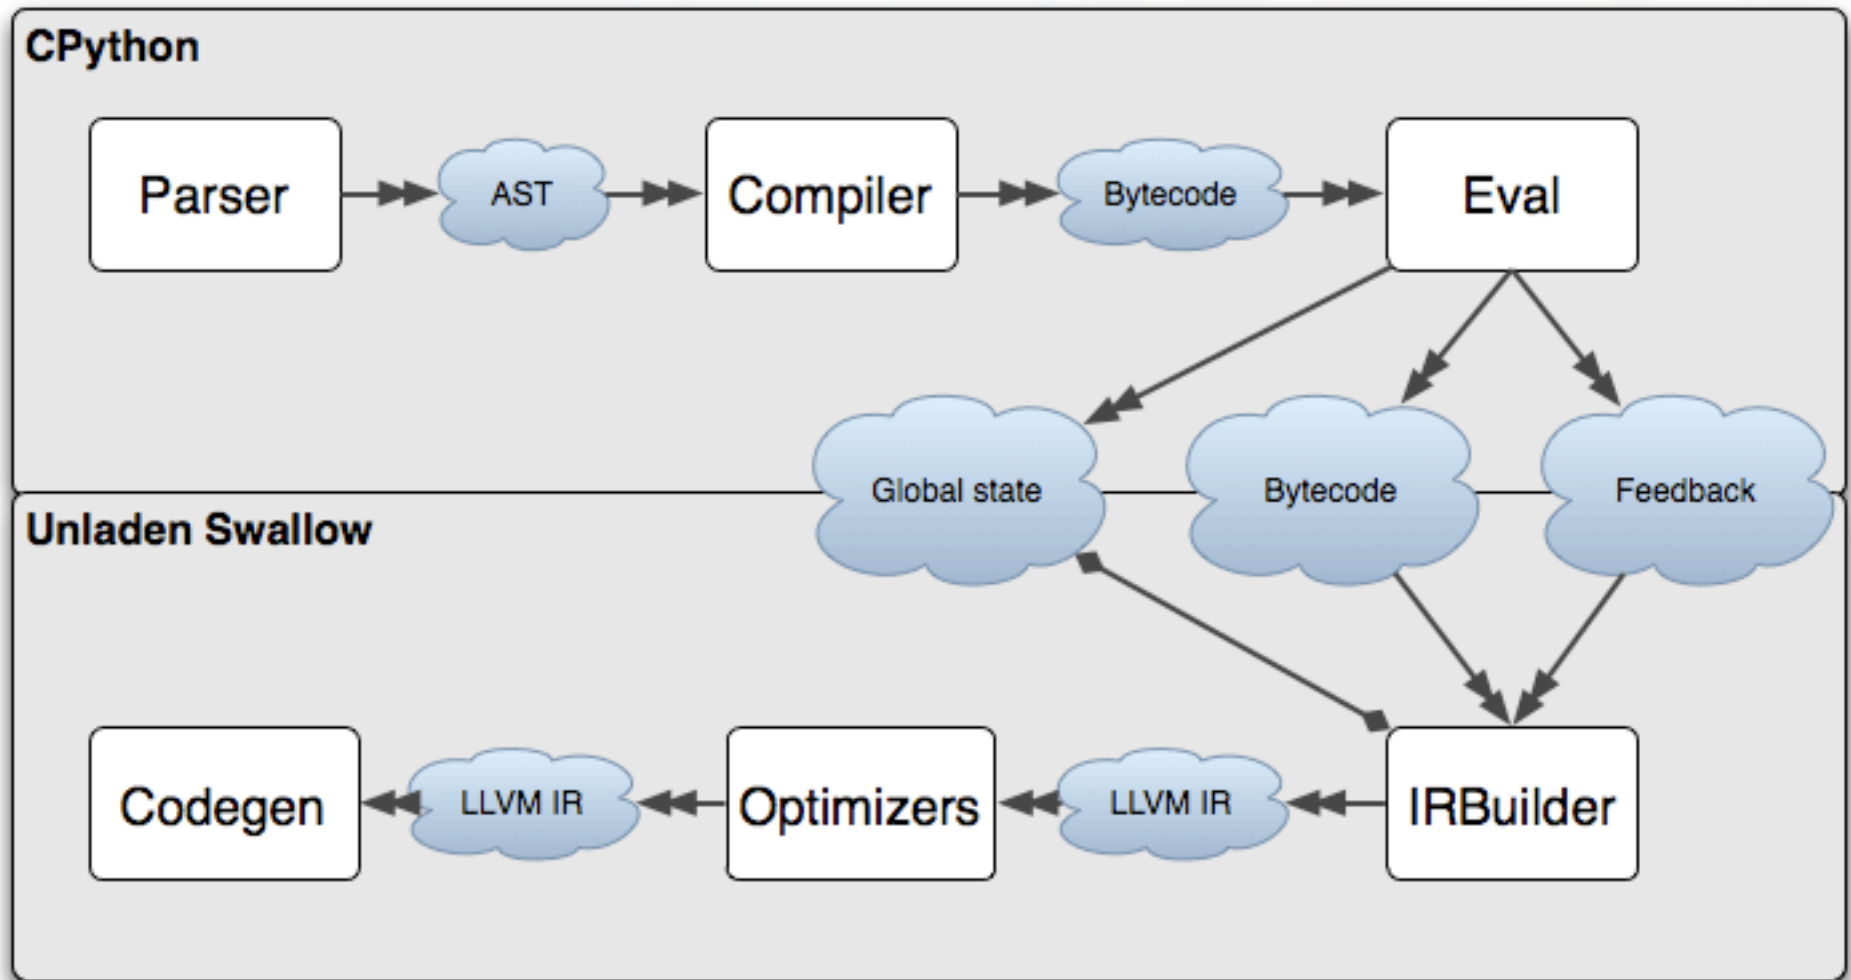


Sunday, February 21, 2010

Prior work:

- Rubinius (Ruby)
- V8 (Chrome), TraceMonkey (Firefox), SquirrelFish Extreme (Safari)
- Pysco, PyPy
- HotSpot JVM, StrongTalk
- Self 93
- SmallTalk

Unladen Swallow



Sunday, February 21, 2010

- Bytecode compiler based on IRBuilder
- Clang used to compile parts of the runtime to IR for inlining

Make it faster: specialize dynamically

```
def escape(value):  
    from django.utils.safestring import mark_for_escaping  
    return mark_for_escaping(value)
```



0	LOAD_CONST	1	(-1)
3	LOAD_CONST	2	(('mark_for_escaping',))
6	IMPORT_NAME	0	(django.utils.safestring)
9	IMPORT_FROM	1	(mark_for_escaping)
12	STORE_FAST	1	(mark_for_escaping)
15	POP_TOP		
16	LOAD_FAST	1	(mark_for_escaping)
19	LOAD_FAST	0	(value)
22	CALL_FUNCTION	1	
25	RETURN_VALUE		



Make it faster: specialize dynamically

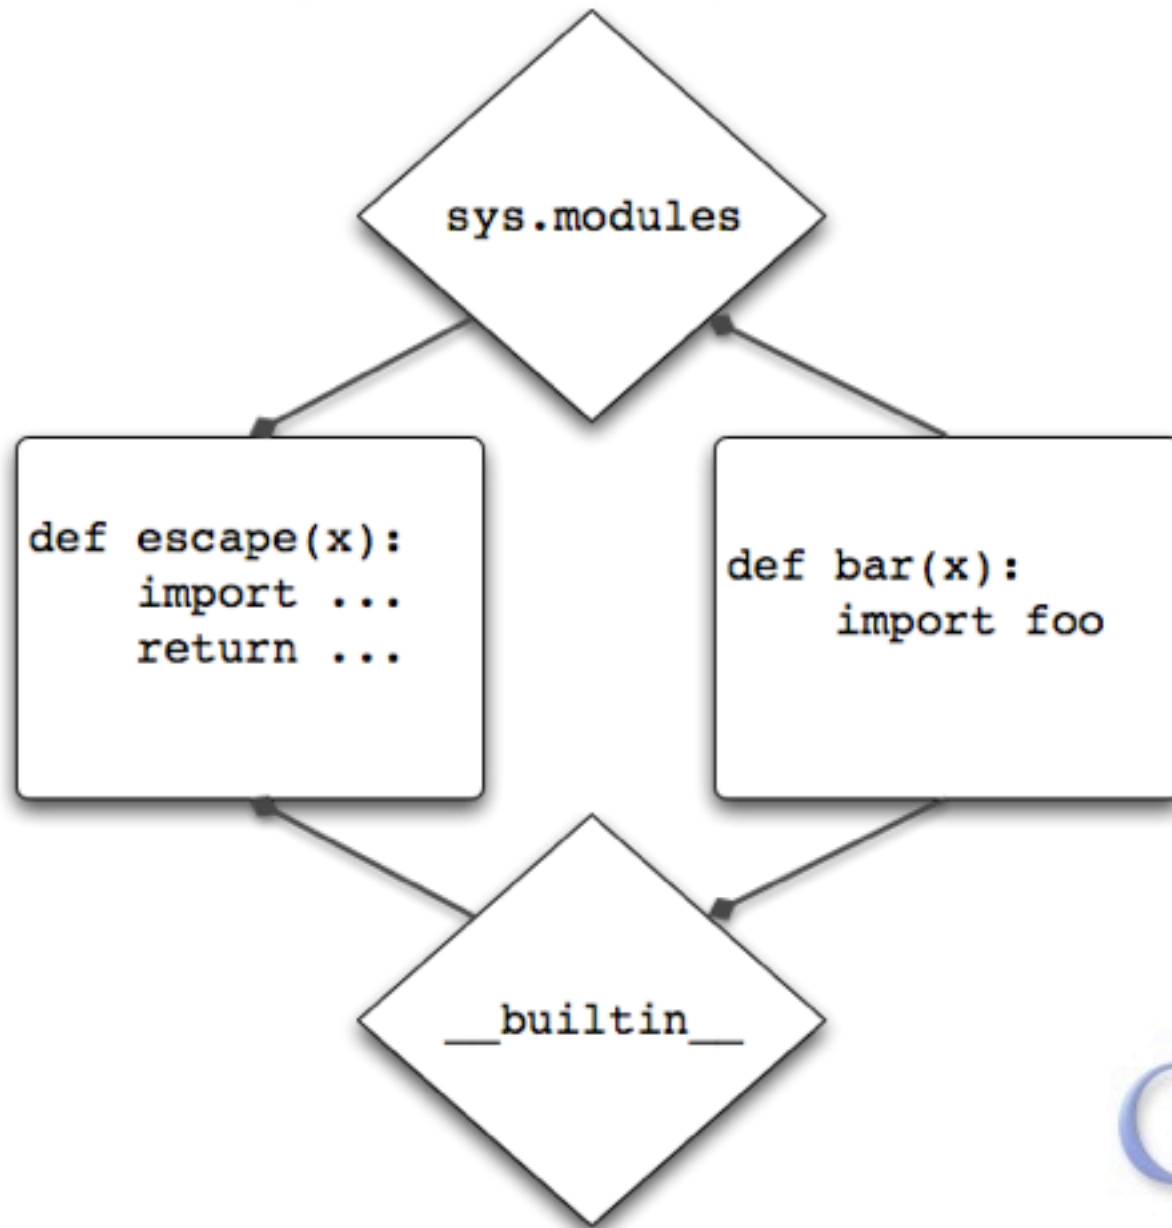
```
6  IMPORT_NAME                0 (django.utils.safestring)
```



```
if (world_has_not_changed) {  
    x = (PyObject *)11712432  
}  
else {  
    goto bail_to_interpreter  
}
```



Make it faster: `if (world_has_not_changed) {`



Make it faster: specialize dynamically

```
def template(data, rows):
    for row in rows:
        data.append("<tr>")
        for col in row:
            data.append("<td>%s</td>" % col)
        data.append("</tr>")
```

```
>>> our_data = []
>>> template(our_data, our_rows)
>>> print "".join(our_data)
```



Sunday, February 21, 2010

Highlights:

- data.append() lookups can be optimized to direct calls
- guard on type version
- "%s" % operator can be specialized

Zoom keys:

Apple option plus

Make it faster: call sites

```
data.append("<td>%s</td>" % col)
```



LOAD_FAST	0	(data)
LOAD_ATTR	0	(append)
LOAD_CONST	1	('<td>%s</td>')
LOAD_FAST	1	(col)
BINARY_MODULO		
CALL_FUNCTION	1	



Sunday, February 21, 2010

Highlights:

- data.append() lookups can be optimized to direct calls
- guard on type version
- "%s" % operator can be specialized

Zoom keys:

Apple option plus

Make it faster: call sites

```
LOAD_FAST      0 (data)
LOAD_ATTR     0 (append)
...
CALL_FUNCTION  1
```



```
data = locals[0];
if (Py_TYPE_VERSION(data) != EXPECTED_TYPE)
    goto bail_to_interpreter;
// LOAD_CONST
// LOAD_FAST
// BINARY_MODULO
retval = list_append(data, modulo_result);
```



Sunday, February 21, 2010

Highlights:

- data.append() lookups can be optimized to direct calls
- guard on type version
- "%s" % operator can be specialized

Zoom keys:

Apple option plus

Top-down Inlining Opportunities

```
>>> x = dict()  
>>> len(x)
```

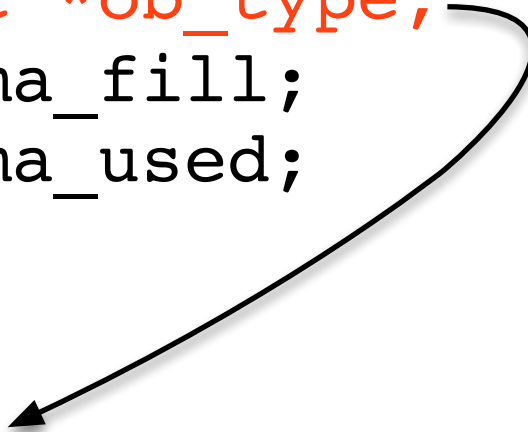
...len() calls PyObject_Size()
...which looks up x->ob_type->tp_as_sequence->sq_length
...which is NULL, so call PyMapping_Size()
...which looks up x->ob_type->tp_as_mapping->mp_length
...which is a function pointer to dict_length()
...which returns x->ma_used



Python Objects: constant(ish)

```
struct PyDictObject {
    Py_ssize_t ob_refcnt;
    PyTypeObject *ob_type;
    Py_ssize_t ma_fill;
    Py_ssize_t ma_used;
    ...
};

PyTypeObject PyDict_Type = {
    ...
    &dict_as_sequence,
    &dict_as_mapping,
    ...
};
```



Sunday, February 21, 2010

- Mirror locals[2] to LLVM IR by walking a pointer tree
- Use a custom AA pass to mark most of the object as constant.
- Constant propagation will remove the second test.

Top-down Inlining: len()

```
static PyObject *
builtin_len(PyObject *self, PyObject *o)
{
    Py_ssize_t res;
    {
        PySequenceMethods *m;
        if (o == NULL) {
            // Error
        }
        m = o->ob_type->tp_as_sequence;
        if (m && m->sq_length)
            res = m->sq_length(o);
        else {
            PyMappingMethods *m;
            if (o == NULL) {
                // Error
            } else {
                m = o->ob_type->tp_as_mapping;
                if (m && m->mp_length)
                    res = m->mp_length(o);
                else
                    // Error
            }
        }
    }
    if (res < 0 && PyErr_Occurred())
        return NULL;
    return PyInt_FromSsize_t(res);
}
```



Sunday, February 21, 2010

- Mirror locals[2] to LLVM IR by walking a pointer tree
- Use a custom AA pass to mark most of the object as constant.
- Constant propagation will remove the second test.

Top-down Inlining: result

```
static PyObject *
builtin_len(PyObject *self, PyObject *v)
{
    return PyInt_FromSsize_t(((PyDictObject *)v)->ma_used);
}
```



Sunday, February 21, 2010

- Mirror locals[2] to LLVM IR by walking a pointer tree
- Use a custom AA pass to mark most of the object as constant.
- Constant propagation will remove the second test.

Measurement & Testing

- Benchmarks:
 - 32 benchmarks (and counting)!
 - YouTube hotspots: Spitfire templates, pickling.
 - Libraries: pickling, regexes, html5lib.
 - Apps: 2to3, Django, Rietveld, SpamBayes, etc.
 - Microbenchmarks: GC, IO, string operations.
 - Now in use by PyPy, soon by Jython.
- Correctness:
 - SWIGed code, extensions used by YouTube.
 - Google's large Python codebase.
 - Large Python projects: Django, NumPy, Twisted, etc.
 - Randomized testing.



Performance vs CPython 2.6.4

Benchmark	CPython	Unladen	Change
2to3	25.13 s	24.87 s	1.01x faster
django	1.08 s	0.68 s	1.59x faster
html5lib	14.29 s	13.20 s	1.08x faster
nbody	0.51 s	0.28 s	1.84x faster
rietveld	0.75 s	0.55 s	1.37x faster
slowpickle	0.75 s	0.55 s	1.37x faster
slowspitfire	0.83 s	0.61 s	1.36x faster
slowunpickle	0.33 s	0.26 s	1.26x faster
spambayes	0.31 s	0.34 s	1.10x slower



Performance vs Jython 2.5.1

Benchmark	Jython	Unladen	Change
ai	1.58	0.52	3.07x faster
django	2.0	0.8	2.48x faster
float	0.33	0.07	3.49x faster
rietveld	2.45	0.59	4.17x faster
slowpickle	0.86	0.56	1.54x faster
slowspitfire	2.76	0.63	4.39x faster
slowunpickle	0.38	0.27	1.37x faster



Performance vs PyPy

Benchmark	PyPy	Unladen	Change
ai	0.61	0.51	1.1921x faster
django	0.67	0.68	1.0136x slower
float	0.03	0.07	2.7108x slower
html5lib	20.04	16.42	1.2201x faster
pickle	17.70	1.09	16.2465x faster
rietveld	1.09	0.59	1.8597x faster
slowpickle	0.43	0.56	1.2956x slower
slowspitfire	2.50	0.63	3.9853x faster
slowunpickle	0.26	0.27	1.0585x slower
unpickle	28.45	0.78	36.6427x faster



Memory Usage vs CPython 2.6.4

Benchmark	CPython	Unladen	Change
2to3	26396 kb	46896 kb	1.77x
django	10028 kb	27740 kb	2.76x
html5lib	150028 kb	173924 kb	1.15x
nbody	3020 kb	16036 kb	5.31x
rietveld	15008 kb	46400 kb	3.09x
slowpickle	4608 kb	16656 kb	3.61x
slowspitfire	85776 kb	97620 kb	1.13x
slowunpickle	3448 kb	13744 kb	3.98x
spambayes	7352 kb	46480 kb	6.32x





Looking Back



Sunday, February 21, 2010

- Q1 release
- Tool detours: gdb, oProfile
- LLVM problems

- What we did right: code reuse,

A decorative header at the top of the page features four overlapping spheres: a green one on the left, and blue, red, and yellow ones on the right.

GIL?

Google™

Sunday, February 21, 2010



Merger



Sunday, February 21, 2010

- PEP is out for review
- Review going well
- Good feedback from community:
 - Prioritize -j never performance
 - Shared linking for LLVM
 - No C++ at all under --without-llvm build

Fin

Questions?

<http://code.google.com/p/unladen-swallow/>

collinwinter@google.com
unladen-swallow@googlegroups.com
#unladenswallow on OFTC

