

Profiling and optimizing RAM and CPU in RMG-Py

Richard West



[Northeastern.edu/comocheng](https://northeastern.edu/comocheng)

**“Premature optimization
is the root of all evil.”**

- Donald Knuth, Turing Award Lecture 1974

“Premature optimization is the root of all evil.”

- Donald Knuth, Turing Award Lecture 1974

Not an excuse to write bad code!

1. Get it working
2. Get it working right
3. Test it's working right
4. Profile it
5. Improve the slow parts, then go to 3

How Python does memory management

Reference counting...

...with a garbage collector.

Memory Allocator is in cPython...

...and in the OS.

Every object has a dictionary of attributes...

...unless you use `__slots__` or Cython.

What makes Python slow?

Attribute lookups!

Slow

```
for species1 in long_list:
    for species2 in long_list:
        distance = ( species1.molecule[0].atom[0].x -
                    species2.molecule[0].atom[0].x )
```

Faster

```
for species1 in long_list:
    x1 = species1.molecule[0].atom[0].x
    for species2 in long_list:
        distance = x1 - species2.molecule[0].atom[0].x
```

Slow

```
newlist = []
for word in oldlist:
    newlist.append(word.upper())
```

Faster

```
upper = str.upper()
newlist = []
append = newlist.append
for word in oldlist:
    append(upper(word))
```

How to profile RMG-Py for CPU use

```
$ python rmg.py -p hexadiene/input.py
```

```
MODEL GENERATION COMPLETED
```

```
The final model core has 9 species and 8 reactions  
The final model edge has 140 species and 317 reactions
```

```
RMG execution terminated at Wed Jan 22 14:42:49 2014
```

```
=====  
Profiling Data  
=====
```

```
Sorted by internal time
```

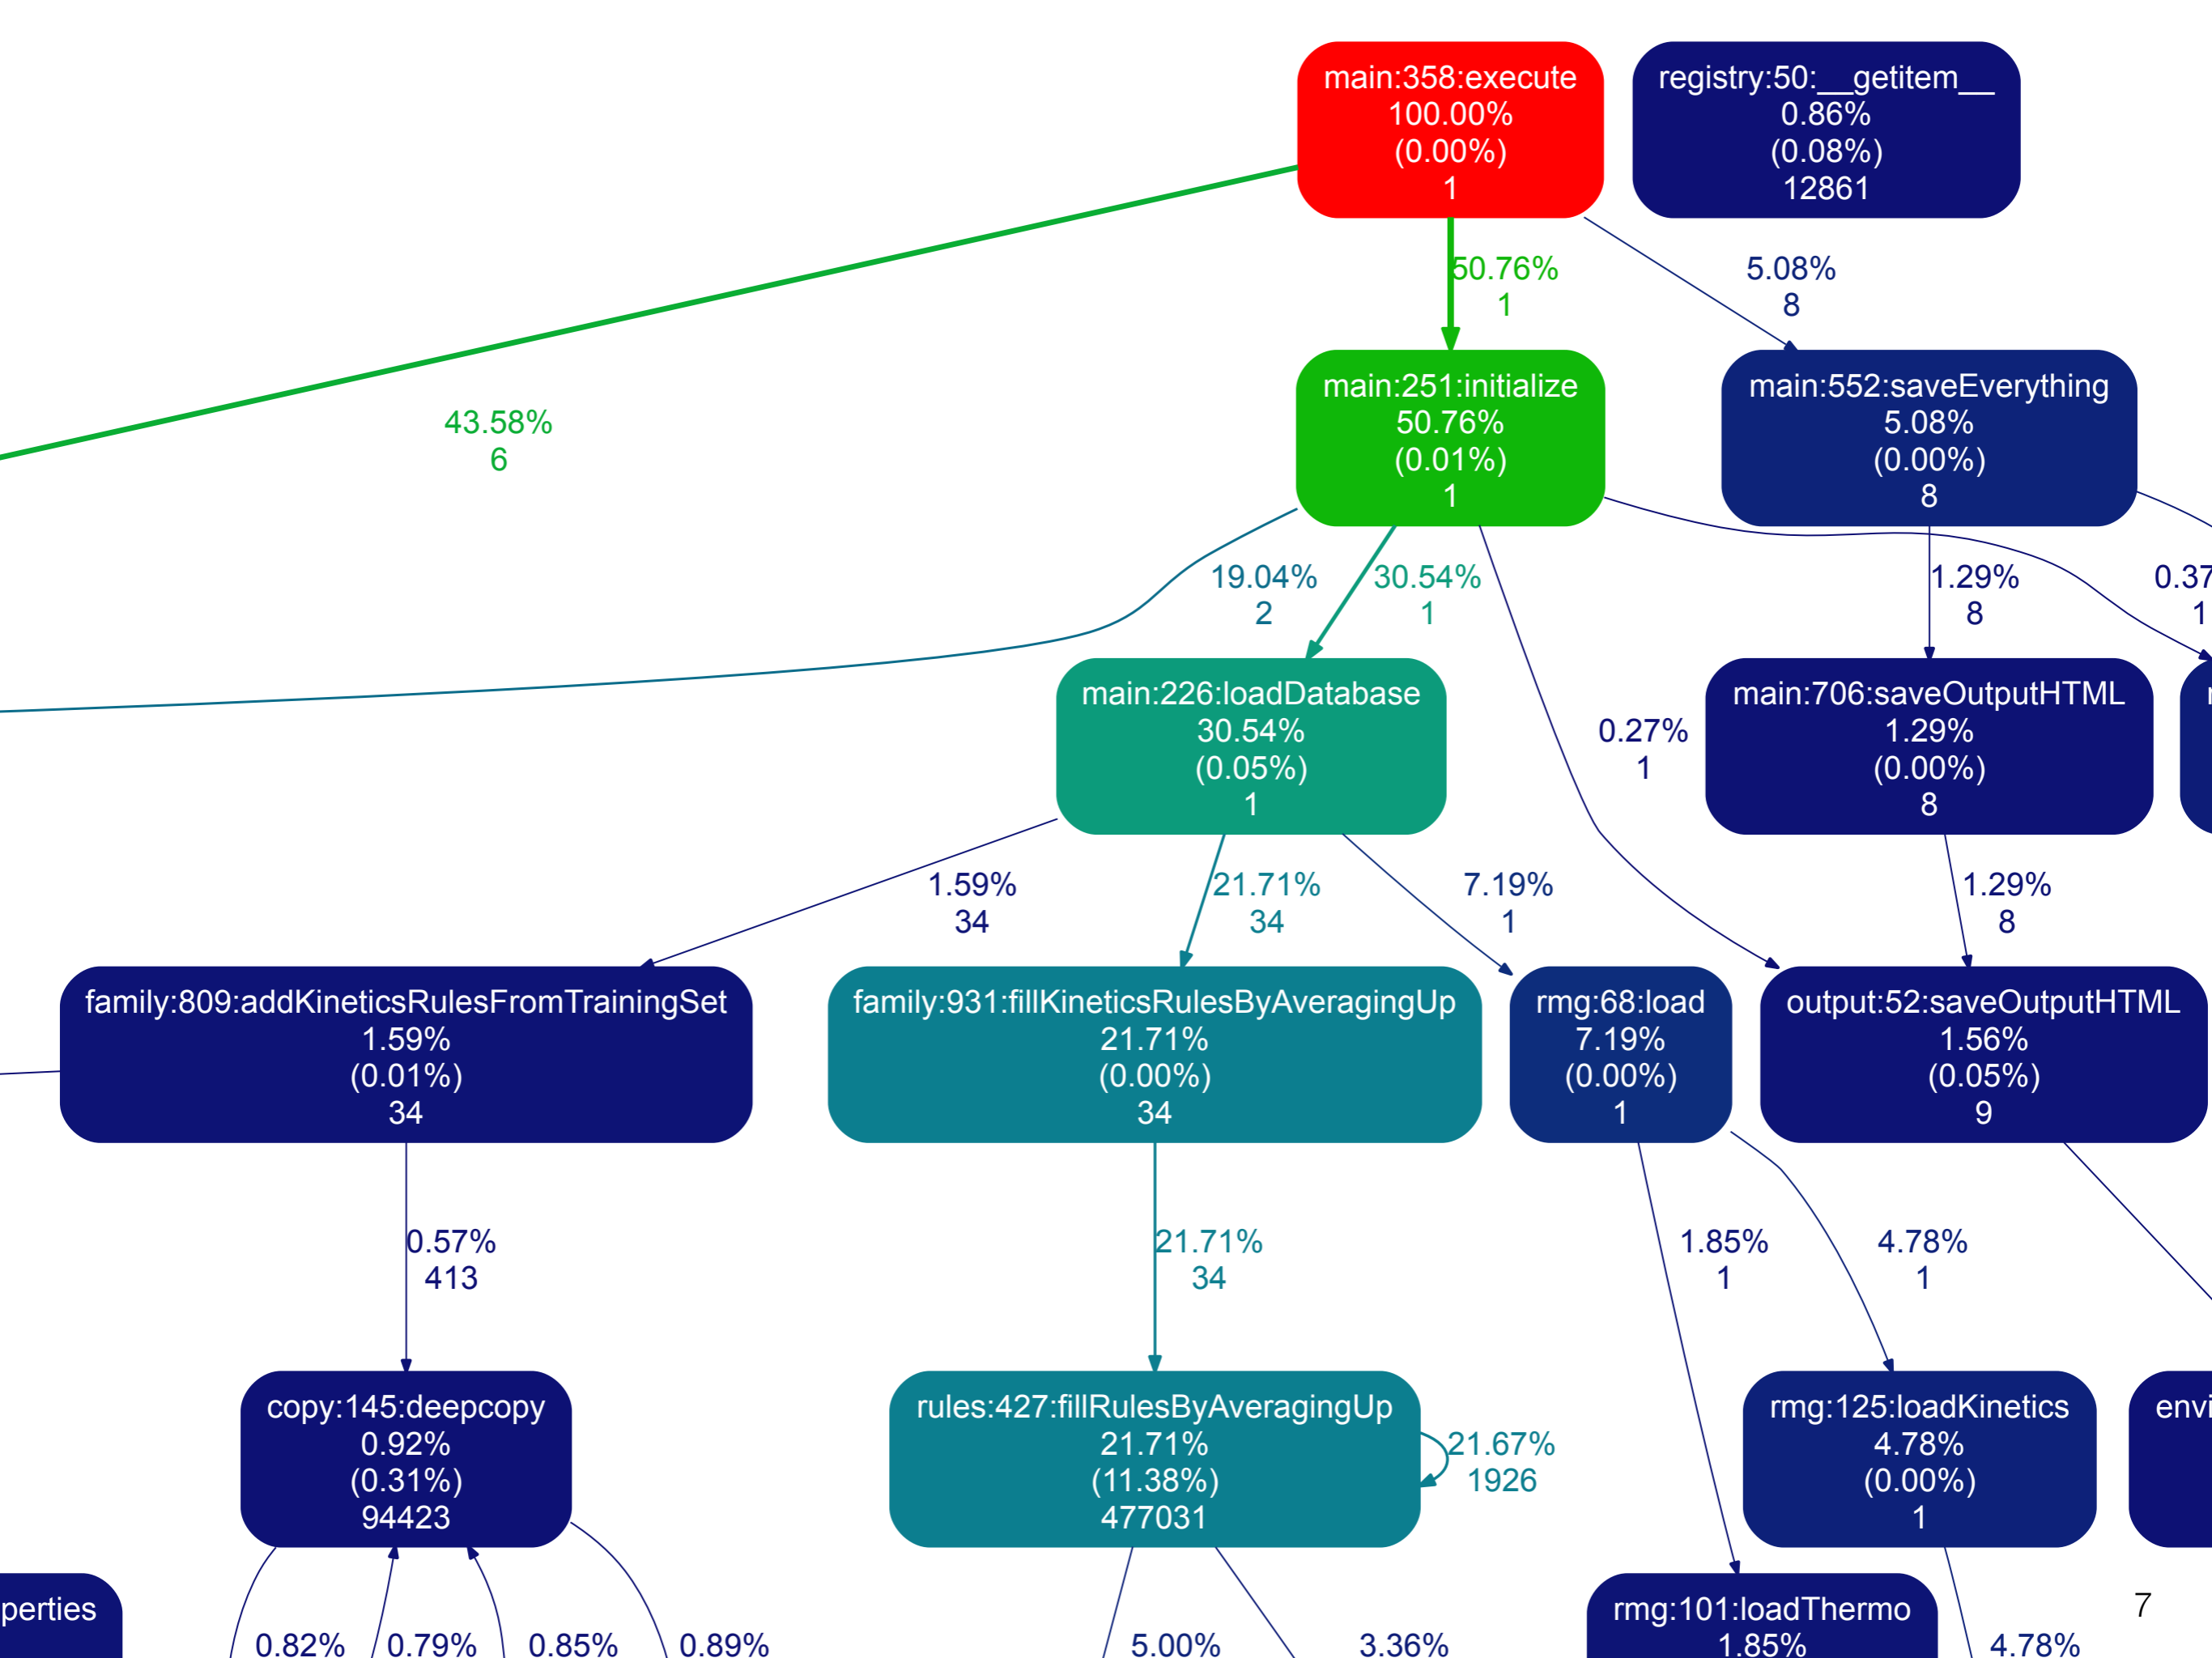
```
Wed Jan 22 14:42:49 2014 /Users/rwest/Code/RMG-Py/examples/rmg/1,3-hexadiene/RMG.profile
```

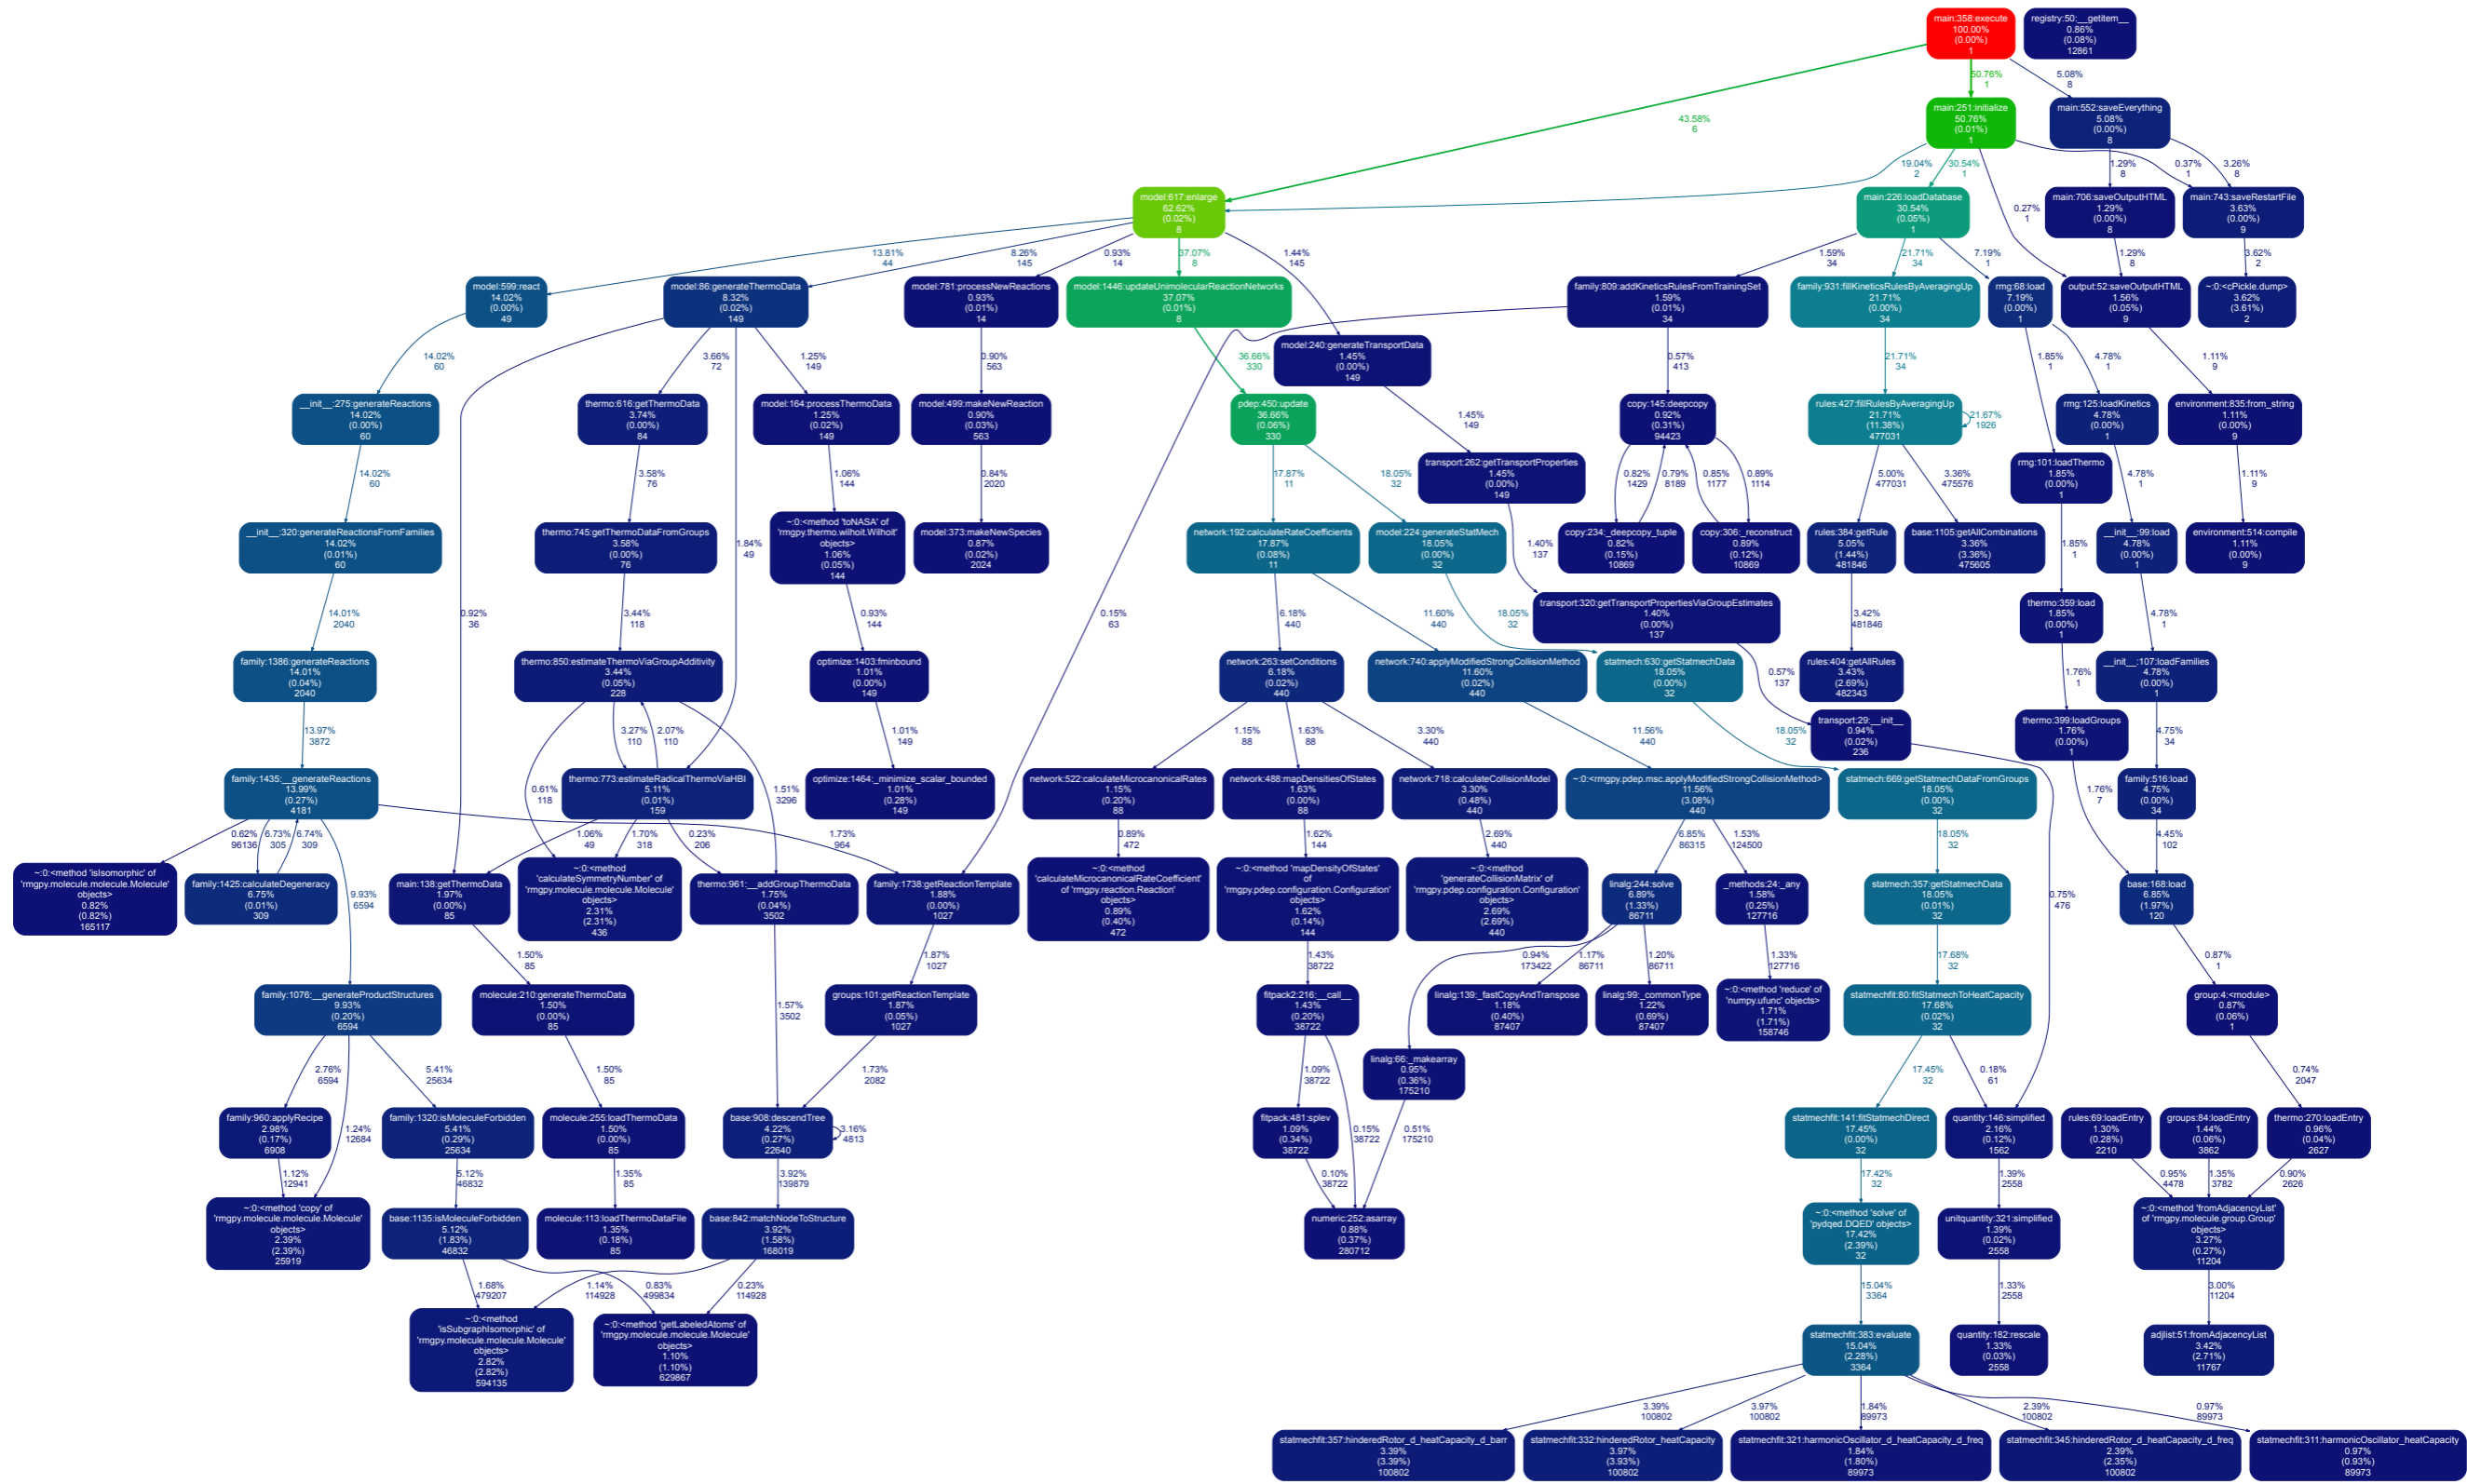
```
25196711 function calls (24439766 primitive calls) in 46.146 seconds
```

```
Ordered by: internal time
```

```
List reduced from 2468 to 25 due to restriction <25>
```

| ncalls | tottime | percall | cumtime | percall | filename:lineno(function) |
|-----------|---------|---------|---------|---------|--|
| 477031/34 | 5.251 | 0.000 | 10.022 | 0.295 | rules.py:427(fillRulesByAveragingUp) |
| 100802 | 1.812 | 0.000 | 1.831 | 0.000 | statmechfit.py:332(hinderedRotor_heatCapacity) |
| 2 | 1.668 | 0.834 | 1.671 | 0.835 | {cPickle.dump} |
| 100802 | 1.563 | 0.000 | 1.563 | 0.000 | statmechfit.py:357(hinderedRotor_d_heatCapacity_d_barr) |
| 475605 | 1.549 | 0.000 | 1.549 | 0.000 | base.py:1105(getAllCombinations) |
| 440 | 1.422 | 0.003 | 5.334 | 0.012 | {rmgpy.pdep.msc.applyModifiedStrongCollisionMethod} |
| 594135 | 1.300 | 0.000 | 1.300 | 0.000 | {method 'isSubgraphIsomorphic' of 'rmgpy.molecule.molecule.Molecule' objects} |
| 11767 | 1.252 | 0.000 | 1.577 | 0.000 | adjlist.py:51(fromAdjacencyList) |
| 482343 | 1.242 | 0.000 | 1.582 | 0.000 | rules.py:404(getAllRules) |
| 440 | 1.240 | 0.003 | 1.240 | 0.003 | {method 'generateCollisionMatrix' of 'rmgpy.pdep.configuration.Configuration' objects} |
| 25919 | 1.104 | 0.000 | 1.104 | 0.000 | {method 'copy' of 'rmgpy.molecule.molecule.Molecule' objects} |
| 32 | 1.101 | 0.034 | 8.042 | 0.251 | {method 'solve' of 'Invdged DOED' objects} |





How to profile RMG-Py for memory use

Just run it



statistics.xls

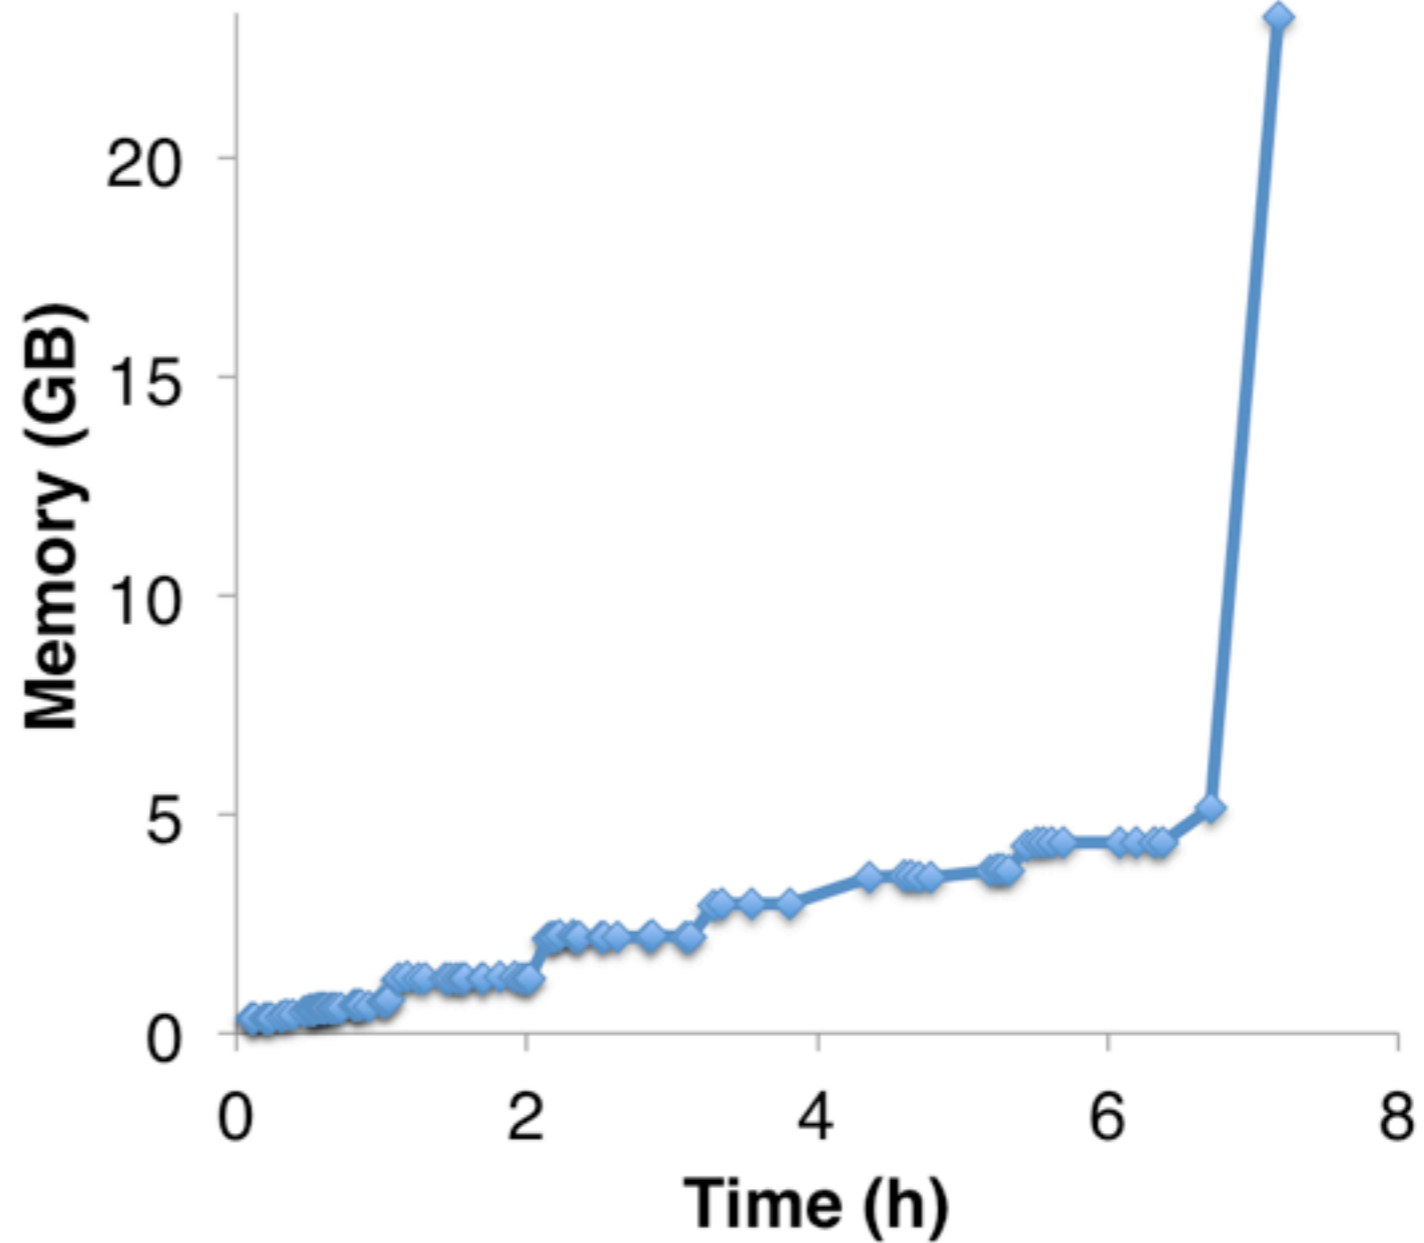
| | A | B | C | D | E | F | G | H |
|----|----------------|--------------|----------------|--------------|---------------|-------------|------------------------|---|
| 1 | Execution time | Core species | Core reactions | Edge species | Edge reaction | Memory used | Restart file size (MB) | |
| 2 | 139.520491 | 5 | 0 | 26 | 16 | 214.130688 | 0 | |
| 3 | 169.839562 | 6 | 1 | 80 | 70 | 242.659328 | 0 | |
| 4 | 178.095639 | 6 | 1 | 84 | 75 | 253.657088 | 0 | |
| 5 | 187.592839 | 6 | 1 | 88 | 79 | 271.437824 | 0 | |
| 6 | 201.334465 | 6 | 1 | 93 | 84 | 259.887104 | 0 | |
| 7 | 233.356402 | 7 | 2 | 150 | 169 | 263.737344 | 0 | |
| 8 | 265.418484 | 8 | 4 | 196 | 239 | 284.758016 | 0 | |
| 9 | 268.541452 | 8 | 4 | 196 | 240 | 261.169152 | 0 | |
| 10 | 279.033994 | 8 | 4 | 198 | 242 | 273.956864 | 0 | |
| 11 | 285.383976 | 9 | 7 | 200 | 260 | 274.026496 | 0 | |
| 12 | 328.440526 | 10 | 10 | 249 | 363 | 310.157312 | 0 | |
| 13 | 352.161246 | 11 | 16 | 283 | 415 | 280.141824 | 0 | |
| 14 | 355.457457 | 12 | 17 | 283 | 430 | 283.7504 | 0 | |
| 15 | 375.273565 | 13 | 21 | 324 | 492 | 284.073984 | 0 | |
| 16 | 421.504672 | 14 | 28 | 438 | 648 | 302.526464 | 0 | |
| 17 | | | | | | | | |

How to profile RMG-Py for memory use

Just run it



statistics.xls



How to profile RMG-Py for memory use

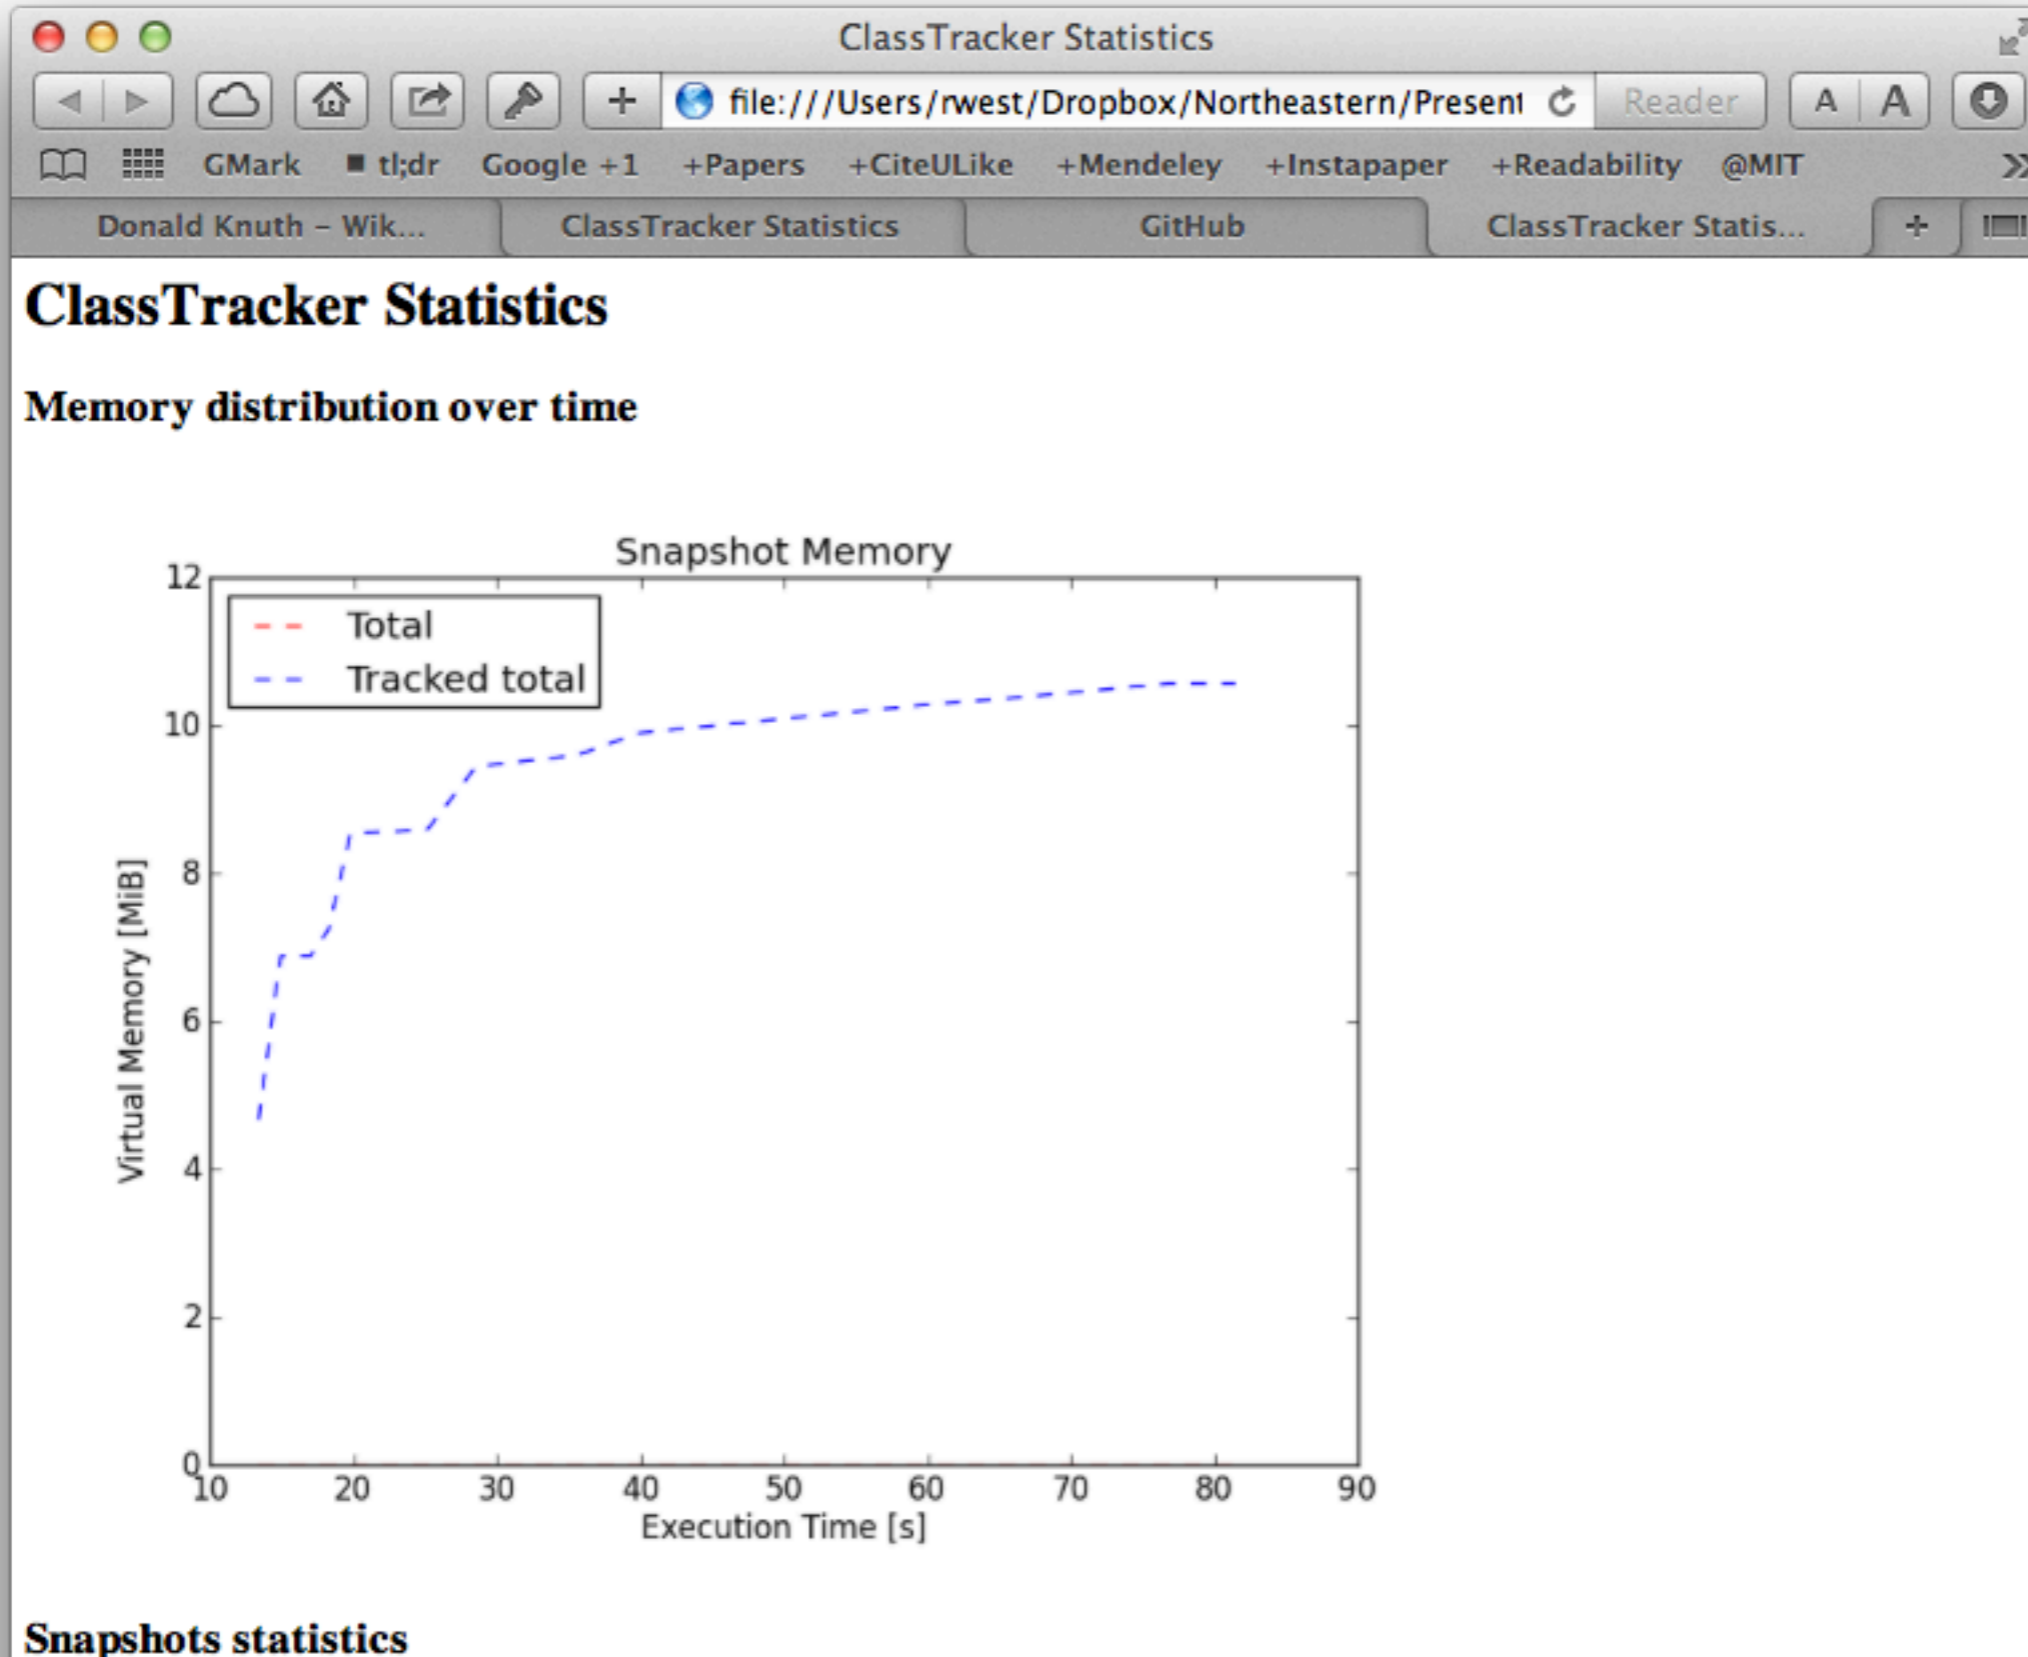
Take your pick of toolkits, libraries, etc.

Guppy

Meliae

Pympler

```
$ python rmg.py -m example/input.py
```



Snapshots statistics

\$ p

2 core 5 edge species 13 seconds snapshot at 00:00:13.42

Total virtual memory assigned to the program at that time was 2.58 GB, which includes 12.14 KB profiling overhead. The ClassTracker tracked 4.66 MB in total. The measurable objects including code objects but excluding overhead have a total size of 0 B.

| Class | Instance # | Total | Average size | Share |
|---|------------|---------|--------------|-------|
| rmgpy_rmg_model.CoreEdgeReactionModel | 1 | 4.66 MB | 4.66 MB | 0.18% |
| rmgpy_rmg_model.Species | 7 | 3.12 KB | 456 B | 0.00% |

Other

3 core 8 edge species 15 seconds snapshot at 00:00:14.94

Total virtual memory assigned to the program at that time was 2.61 GB, which includes 21.41 KB profiling overhead. The ClassTracker tracked 6.87 MB in total. The measurable objects including code objects but excluding overhead have a total size of 0 B.

| Class | Instance # | Total | Average size | Share |
|---|------------|---------|--------------|-------|
| rmgpy_rmg_model.CoreEdgeReactionModel | 1 | 6.87 MB | 6.87 MB | 0.26% |
| rmgpy_rmg_model.Species | 11 | 4.87 KB | 453 B | 0.00% |

Other

4 core 8 edge species 16 seconds snapshot at 00:00:16.02

Other resources

SlideShare.net

e.g. http://www.slideshare.net/g3_nittala/profiling-and-optimization-14511764

Google.com

e.g. “Python memory profiling”