

COEN 45, MATLAB Programming  
Spring Quarter, 2011

Lab assignment #4  
The Collatz problem  
Feb. 1, 2, 3

The subject of this assignment is the Collatz problem, a long-standing mystery of number theory. If you solve the mystery, you will be famous! See “Collatz conjecture” on wikipedia.org for more information, including a really cool fractal plot that I would like to try to figure out how to program.

The Collatz problem is based on this very simple algorithm:

1. Start with any integer  $n > 2$
2. Does  $n$  equal 1? If so, stop.
3. If  $n$  is even, divide it by two. Otherwise, triple it and add 1.
4. Go to step 2.

Example sequence:

9, 28, 14, 7, 22, 11, 34, 17, 52, 26, 13, 40, 20, 10, 5, 16, 8, 4, 2, 1

The mystery is whether every possible starting number results in convergence to 1 in a finite number of steps. As far as I know, no exception to this rule has been found, but neither has anyone proved its truth.

Normally, round-off error does not present significant difficulty, but it could in this problem. That’s because we assume integer arithmetic in the Collatz problem whereas MATLAB does all computations in double floating-point form. The largest integer that can be represented exactly in this form is  $2^{53}$ . If your script produces a larger number, MATLAB will continue, but there will be round-off error that will very likely spoil the results. So when you code the algorithm you must watch for numbers that exceed  $2^{53}$  and stop if you get one.

To do:

- Write a function `collatz` that returns a Collatz sequence in a vector `c`. Download the template from the class web site and finish it.
- Check out your function by exercising it at the command line. When you are satisfied, get the script `lab4.m` from the class web site and finish it. After you get your plots, use the magnifying glass on the plot to see if you find any interesting patterns.
- Have some fun with the script. See how high you can go. Note: if you try to compute the Collatz results for all odd integers up to a very large number like  $10^9$  you will probably run out of memory before you get overflow. But you can try very large numbers one at a time. `lab4.m` as

written takes about 17 sec on my dual AMD64 at home. If it takes too long in the lab, you may have to cut back on the number of trials. *Warning:* if you run the CPU flat out on your laptop for 20 minutes, which you can probably do with this script, it could overheat!

- Demonstrate your work to the lab assistant. Submit a signed printout of your script file along with a printed copy of your plots with your name and student ID in the title.

You can find an enormous amount of math associated with the Collatz problem. In fact, an entire conference was held on just this problem:  
<http://www.math.grin.edu/~chamber1/conf.html>