

University of Saskatchewan
Department of Mathematics and Statistics

Numerical Analysis I
(MATH 211)

Instructor: Dr. Raymond J. Spiteri

ASSIGNMENT 05

Due: 8:30 a.m. Tuesday, April 02, 2013

1. **[25 marks]** Simpson's 3/8-rule is a quadrature rule with 4 equally spaced interpolation nodes that include the interval endpoints a and b . Use the method of undetermined coefficients to derive this quadrature rule.

By considering the polynomials that Simpson's rule and Simpson's 3/8-rule can integrate exactly, explain why not too many people will have heard of Simpson's 3/8-rule.

2. **[25 marks]** In this question we will try to estimate

$$I = \int_0^1 x \sin\left(\frac{1}{x}\right) dx.$$

- (a) Use `ezplot` to plot $f(x) = x \sin(1/x)$ for $0 \leq x \leq 1$.
- (b) Why might this integrand be misbehaved? Why might this misbehaviour ultimately not be problematic?
- (c) Use the `int` command to find the exact value of I (Note: This will be in terms of some funny quantities.) Use the `vpa` command to give I to 25 decimal places.
- (d) What happens when you try to evaluate this integral directly using the `quad` function with the default tolerance of 10^{-6} ? Is the answer really accurate to the specified tolerance? Using your knowledge of the exact answer, how many decimal places of accuracy does the answer really have?
- (e) Normally (of course) the exact answer is unavailable! We can however still determine how accurate our estimate from part (d) (with the default tolerance) is by re-solving the problem with a tolerance that is lower than the default. Do this and confirm that the number of decimal places you claimed were accurate in part (d) by knowing the exact answer is the same now if we pretend not to know the exact answer.

3. [25 marks] Ramanujan proposed that the number of positive integers n that are either perfect squares or the sums of two perfect squares and lie between the positive integers a and b is given approximately by the integral

$$n \approx 0.764 \int_a^b \frac{dx}{\sqrt{\ln(x)}}.$$

We wish to solve this problem for $a = 1$, $b = 30$.

Explain where the integrand may be problematic.

What is the name of quadrature rule behind the MATLAB function `quadgk`?

By noting the location of the abscissae of this rule, explain why `quadgk` may be a suitable function to use to compute n for these values of a and b .

Use `quadgk` to numerically test the accuracy of Ramanujan's approximation.

4. [25 marks] Use (serial) MATLAB to compute the value of the double integral

$$I = \int_0^\pi \int_0^\pi (y^2 \sin x + x^2 \cos^2 y) dx dy$$

to 3 decimal places using the Monte Carlo method.

On `socrates.usask.ca`, compute the time it takes to compute I using the same number of function evaluations that you just used to answer the first part of this question with a `matlabpool` of 2, 4, and 8 workers. Are the times what you would expect? Explain why or why not.

Bonus: [10 marks] There will be small differences in the results. Explain.

Hints:

(a)

$$\int_0^\pi \int_0^\pi (y^2 \sin x + x^2 \cos^2 y) dx dy = \int_0^\pi y^2 dy \int_0^\pi \sin x dx + \int_0^\pi x^2 dx \int_0^\pi \cos^2 y dy.$$

(b) There are the main modifications to your (usual) MATLAB code that must be made to parallelize it in this case.

- i. A `matlabpool` of workers must be `opened` and set to the appropriate number of workers. Of course, it must also be `closed` at the end of the program.
- ii. The *outermost* `for` loop of your Monte Carlo code must be changed to `parfor`. Any inner `for` loops should remain unaltered.

(c) In order to execute the code in parallel, it must be submitted to a queue using an appropriately modified version of `submitMatlabScript.sh`.

It is submitted using the command

```
qsub -l nodes=1:ppn=N submitMatlabScript.sh
```

where `N` is the number of workers opened in the `matlabpool` command. Upon submission, you should see a message like

```
12345.socrates.usask.ca
```

where 12345 is the job number. Upon completion, you should see two files `submitMatlabScript.sh.e12345` and `submitMatlabScript.sh.o12345`. The first is the error log and should be empty (if all went well). The second is the output file, from which you can record the result as well as the execution time.

Your serial MATLAB code should have `exit` as its last statement for it to terminate properly when executed through the queue.

- (d) In order for the `matlab` command to be recognized, the appropriate modules must be loaded by typing the following commands at the command line:

```
module add matlab/R2012a matlab/mcr
module initadd matlab/R2012a matlab/mcr
```

These commands only need to be typed in once; all subsequent sessions should have the modules loaded automatically.

To check whether the modules have been loaded properly, you can type the command `module list`. The output should contain the following:

Currently Loaded Modulefiles:

```
1) rocks-openmpi  2) matlab/R2012a  3) matlab/mcr
```

- (e) In order for your result of the Monte Carlo method to be reproducible, the random number generator must be *seeded* with a number unique to the worker.

For a `parfor` loop counter `i`, one way to do this is to set `seed = i; rng(seed)`.

- (f) Don't leave this question until the last minute! There are a limited number of Parallel Computing Toolbox licenses, so you may not simply be able to log in whenever you wish and run your code.