

Supercomputing in Plain English, Spring 2009

Exercise #2: Tiling

In this exercise, we'll use the same conventions and commands as in Exercise #1. You should refer back to the Exercise #1 description for details on various Unix commands.

In the exercise, you'll **benchmark** matrix-matrix multiplication algorithms, choosing various matrix sizes and various ways of performing the multiplication. **Benchmark** means to run timing tests.

Specifically, you'll benchmark the following methods:

- a naïve algorithm that performs the matrix-matrix multiplication the way you would normally do it by hand, which you'll run on various matrix sizes;
- a tiling algorithm, for which you'll be selecting not only various matrix sizes, but also various tile sizes;
- if you're running the Fortran90 version, the Fortran90 intrinsic routine MATMUL, for various matrix sizes.

1. Log in to Sooner.
2. Copy the Tiling directory:

```
cp -r ~hneeman/SIPE2009_exercises/Tiling/ ~/SIPE2009_exercises/
```

3. Choose which language you want to use (C or Fortran90), and `cd` into the appropriate directory:

```
cd ~/SIPE2009_exercises/Tiling/C/
```

OR:

```
cd ~/SIPE2009_exercises/Tiling/Fortran90/
```

4. Edit the batch script `matmatmult.bsub` so that it contains your username and your e-mail address.
5. Compile:

```
make
```

6. Edit the input file `matmatmult_input.txt` so that it contains your preferred problem size and type:
 - a. The first three numbers are, respectively:
 - i. the number of rows in the product matrix;
 - ii. the number of columns in the product matrix;
 - iii. the number of columns of the first matrix to multiply by, which is also the number of rows of the second matrix to multiply by.
 - b. The next number is the preferred matrix-matrix multiplication algorithm:
 - i. choose 1 for the naïve algorithm;
 - ii. choose 2 for the tiling algorithm;
 - iii. for Fortran90 only, choose 3 for the Fortran90 intrinsic routine `MATMUL`.
 - c. If you've chosen the tiling algorithm, then the last three numbers are the dimensions of the tiles, which correspond to the dimensions of the matrices. (For other algorithms, the last three numbers are ignored.)

7. Submit the batch job:

```
bsub < matmatmult.bsub
```

8. Once the batch job completes, examine the stdout file to see the timing for your run.
9. Run many more runs, of each of the available algorithms, each for several different problem sizes:
 - a. Start by finding the largest problem that reports a runtime of zero seconds (that is, immeasurably small runtime).
 - b. Work your way up in problem size to as big as you want, but the biggest problem you should do should be at most about 14 GB total. (Beyond that, your timings will become ridiculously long, because you'll spend most of your runtime swapping in and out of disk, which would be really really bad).
 - c. For the tiling algorithm, also try many different tile sizes, from very small up to the same size as the matrices.
10. When you've completed enough runs to satisfy yourself, use your favorite graphing program (for example, Microsoft Excel) to create a graph (or graphs) of your various runs, so that you can compare the various methods visually.
11. If you're feeling especially adventuresome, you can edit your `makefile` to incorporate the `-x8` option in the `CFLAGS` or `FFLAGS` macros, then repeat the same experiments. (Be careful not to run problems that are too large, which will be easier to do in double precision, because you'll consume twice as many bytes for the same matrix dimensions.) Compare these new double precision runs against the original single precision runs.