

Comparing performance of different implementations of matrix multiplication

Last modified: October 20, 2015

The following function implements multiplication of two $n \times n$ matrices:

$$C_{ij} = \sum_{k=1}^n A_{ik} B_{kj}$$

```
1  * Naive implementation of matrix multiplication c += ab,
2  * where a is m x n, b is n x p, and c is m x p, in column-major order.
3  *
4  * The physical sizes of a, b, and c are lda x n, ldb x p, and ldc x p,
5  * but only the first m/n/m rows are used, respectively.
6  *
7  *   c_{ij} += \sum_{k=0}^{n-1} a_{ik} b_{kj}
8  *
9  *   a_{ij} <-> a[i + j*lda] <- column-major order
10 */
11 void matmul_naive (const double *a, const double *b, double *c, const int m,
12     const int n, const int p, const int lda, const int ldb, const int ldc)
13 {
14     for (int i = 0; i < m; i++)
15     {
16         for (int j = 0; j < p; j++)
17         {
18             double sum = 0.0;
19
20             for (int k = 0; k < n; k++)
21             {
22                 sum += a[i + lda * k] * b[k + ldb * j];
23             }
24             c[i + ldc * j] += sum;
25         }
26     }
27 }
```

The graph in Fig. 1 compares the CPU performance, in MFLOPS, vs. matrix size n achieved during matrix multiplication by two different implementations - the naive one shown above and the state of the art implementation in OpenBLAS library.

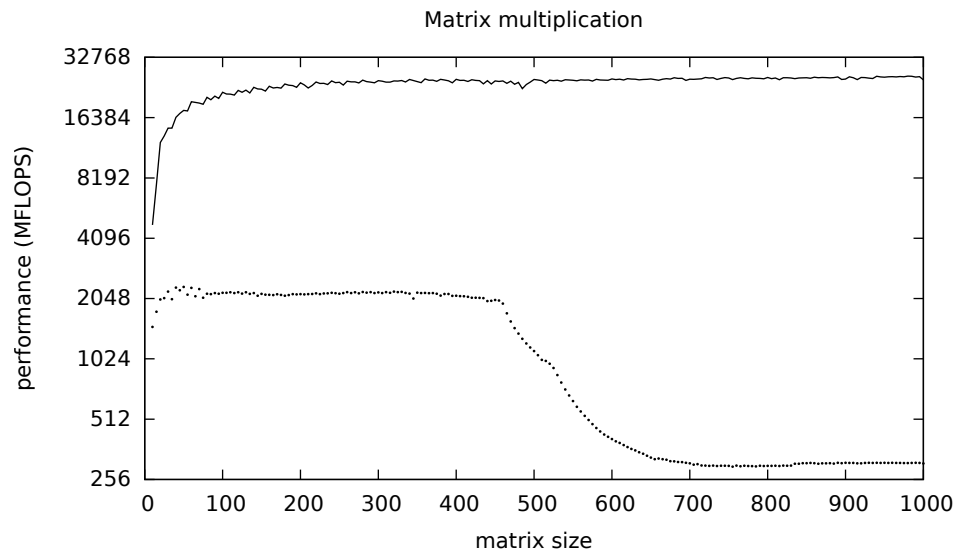


Figure 1: Performance of two implementations of matrix multiplication: the code above (dotted line) and dgemm function from the OpenBLAS library (solid line). Notice the logarithmic vertical axis.

```

1 set term pdf mono lw .5
2 set o "performance.pdf"
3
4 unset key
5 set title "Matrix multiplication"
6 set xlabel "matrix size"
7 set ylabel "performance (MFLOPS) "
8 set logscale y 2
9
10 p [:] [:] "naive.res" u 1:2 w d, "dgemm.res" u 1:2 w l

```

Figure 2: Gnuplot script that produced Figure 1.