

Week 1: Friday, Aug 24

Logistics

1. Lecture 1 notes are posted. In general, lecture notes will be posted under the “lectures” page.
2. My official office hours are Wednesday 1:30–2:30 pm, Thursday 3:00–4:00 pm, and Friday 10:00–11:00 pm. You can also drop by my office, assuming the door is open.

The lazy man’s approach to performance

An algorithm like matrix multiplication seems simple, but there is a lot under the hood of a tuned implementation, much of which has to do with the organization of memory. We often get the best “bang for our buck” by taking the time to formulate our algorithms in block terms, so that we can spend most of our computation inside someone else’s well-tuned matrix multiply routine (or something similar). There are several implementations of the Basic Linear Algebra Subroutines (BLAS), including some implementations provided by hardware vendors and some automatically generated by tools like ATLAS. The best BLAS library varies from platform to platform, but by using a good BLAS library and writing routines that spend a lot of time in *level 3* BLAS operations (operations that perform $O(n^3)$ computation on $O(n^2)$ data and can thus potentially get good cache re-use), we can hope to build linear algebra codes that get good performance across many platforms.

This is also a good reason to use MATLAB: it uses pretty good BLAS libraries, and so you can often get surprisingly good performance from it for the types of linear algebraic computations we will pursue.

Blocked matrix multiply

Supposing we have very limited memory (e.g. only registers) available, let’s consider where loads and stores might occur in the inner product version of square matrix-matrix multiply:

```

% To compute C = C + A*B
for i = 1:n
    for j = 1:n
        % load C(i,j)
        for k = 1:n
            % load A(i,j) and B(k,j)
            C(i,j) = C(i,j) + A(i,k)*B(k,j);
        end
        % store C(i,j)
    end
end
end

```

We can count that there are $2n^2 + 2n^3$ loads and stores in total – that is, we do roughly the same number of memory transfer operations as we do arithmetic operations. If memory transfers are much slower than arithmetic, this is trouble.

Now suppose that $n = Nb$, where b is a block size and N is the number of blocks. Suppose also that b is small enough that we can fit one block of A , B , and C in memory. Then we may write

```

for i = 1:N
    for j = 1:N
        % Make index ranges associated with blocks
        I = ((i-1)*N+1):i*N;
        J = ((j-1)*N+1):j*N;

        % load block C(I,J) into fast memory

        for k = 1:N
            % load blocks A(I,K) and B(K,J) into fast memory
            K = ((k-1)*N+1):k*N;
            C(I,J) = C(I,J) + A(I,K)*B(K,J);
        end

        % store block C(I,J)
    end
end
end

```

Following the previous analysis, we find there are $2N^2 + 2N^3$ loads and stores of blocks with this algorithm; since each block has size $b \times b$, this gives

$(2N^2 + 2N^3)b^2 = 2n^2 + 2n^3/b$ element transfers into fast memory. That is, the number of transfers into fast memory is about $1/b$ times the total number of element transfers.

We could consider using different block shapes, but it turns out that roughly square blocks are roughly optimal. It also turns out that this is only the start of a very complicated story. The last time I taught a high-performance computing class, we worked on a tuned matrix-matrix multiplication routine as one of the first examples. My version roughly did the following:

1. Split the matrices into 56×56 blocks.
2. For each product of 56×56 matrices, copy the submatrices into a contiguous block of aligned memory.
3. Further subdivide the 56×56 submatrices into 8×8 submatrices, which are multiplied using a simple fixed-size basic matrix multiply with a few annotations so that the compiler can do lots of optimizations. If the overall matrix size is not a multiple of 8, pad it with zeros.
4. Add the product submatrix into the appropriate part of C .

By using the Intel compiler and choosing the optimizer flags with some care, I got good performance in the tiny $8 \times 8 \times 8$ matrix multiply kernel that did all the computation – it took advantage of vector instructions, unrolled loops to get a good instruction mix, etc. The blocking then gave me reasonably effective cache re-use. It took me a couple days to get this to my liking, and my code was still not as fast as the matrix multiplication routine that I got from a well-tuned package.

Matrix-vector multiply revisited

Recall from last time that I asked you to look at three organizations of matrix-vector multiply in MATLAB: multiplication with the built-in routine, with a column-oriented algorithm, and with a row-oriented algorithm. I coded up the timing myself:

```
for n= 4095:4097
```

```
fprintf('-- %d --\n', n);
A= rand(n);
b= rand(n,1);

tic;
for trials = 1:20
    c1 = A*b;
end
fprintf('Standard multiply: %g ms\n', 50*toc);

tic;
for trials = 1:20
    c2 = zeros(n,1);
    for j = 1:n
        c2 = c2 + A(:,j)*b(j);
    end
end
fprintf('Column-oriented:    %g ms\n', 50*toc);

tic;
for trials = 1:20
    c3 = zeros(n,1);
    for j = 1:n
        c3(j) = A(j,:)*b;
    end
end
fprintf('Row-oriented:        %g ms\n', 50*toc);

end
```

This is what I saw the last time I ran the class, in 2009:

```
-- 4095 --
Standard multiply: 21.2793 ms
Column-oriented:   96.799 ms
Row-oriented:      185.404 ms
-- 4096 --
Standard multiply: 22.7508 ms
Column-oriented:   97.1834 ms
```

```
Row-oriented:      379.368 ms
-- 4097 --
Standard multiply: 22.4621 ms
Column-oriented:   97.5955 ms
Row-oriented:      185.886 ms
```

Notice that the row-oriented algorithm is about twice as slow as the column-oriented algorithm for $n = 4095$, and that *it is twice as slow for $n = 4096$ as for $n = 4095$!* The reason again is the cache architecture. On my machine, the cache is 8-way set associative, so that each memory location can be mapped to only one of eight possible lines in the cache, and each row of A uses exactly the same set of eight lines when $n = 4096$. Therefore, we effectively get *no* cache re-use with $n = 4096$. With $n = 4095$ we don't have this conflict (or at least we suffer from it far less); and since each cache line holds two doubles, we effectively load two rows of A at a time.

This is what I saw this time:

```
-- 4095 --
Standard multiply: 8.58982 ms
Column-oriented:   33.963 ms
Row-oriented:      304.871 ms
-- 4096 --
Standard multiply: 8.02207 ms
Column-oriented:   32.9369 ms
Row-oriented:      310.911 ms
-- 4097 --
Standard multiply: 8.2646 ms
Column-oriented:   32.4177 ms
Row-oriented:      305.567 ms
```

I've changed machines between now and then, and the performance of the new machine is better when I do the smart thing (using the built-in library) and even when I do the column-oriented version. But the performance of the new machine is even worse than the old one when it comes to the cache-thrashing row-oriented version.

If by now you're feeling too lazy to want to untangle the details of how the cache architecture will affect matrix-matrix and matrix-vector multiplies, let alone anything more complicated, don't worry. That's basically the point.

Vector norms

We now switch gears for a little while and contemplate some analysis. When we compute, the entities in the computer have some error, both due to round-off and due to imprecise input values. We need to study *perturbation theory* in order to understand how error is introduced and propagated. Since it may be cumbersome to write separate error bounds for every component, it will be very convenient to express our analysis in terms of *norms*.

A norm is a scalar-valued function from a vector space into the real numbers with the following properties:

1. *Positive-definiteness*: For any vector x , $\|x\| \geq 0$; and $\|x\| = 0$ iff $x = 0$
2. *Triangle inequality*: For any vectors x and y , $\|x + y\| \leq \|x\| + \|y\|$
3. *Homogeneity*: For any scalar α and vector x , $\|\alpha x\| = |\alpha| \|x\|$

We will pay particular attention to three norms on \mathbb{R}^n and \mathbb{C}^n :

$$\begin{aligned}\|v\|_1 &= \sum_i |v_i| \\ \|v\|_\infty &= \max_i |v_i| \\ \|v\|_2 &= \sqrt{\sum_i |v_i|^2}\end{aligned}$$

Also, note that if $\|\cdot\|$ is a norm and M is any nonsingular square matrix, then $v \mapsto \|Mv\|$ is also a norm. The case where M is diagonal is particularly common in practice.

In finite-dimensional spaces, all norms are equivalent. That is, if $\|\cdot\|$ and $|||\cdot|||$ are two norms on the same space, then there are constants $C_1, C_2 > 0$ such that for all x ,

$$C_1 \|x\| \leq |||x||| \leq C_2 \|x\|.$$

However, C_1 and C_2 are not necessarily small, so it makes sense to choose norms somewhat judiciously. In particular, if different elements of x have different units (furlongs and fortnights), it usually pays to nondimensionalize the problem before doing numerics; and this can be interpreted as choosing a diagonally scaled norm.

Norms and inner products

Recall that an *inner product* on a vector space is a function of two vectors with the following properties:

1. *Positive-definiteness*: $\langle x, x \rangle \geq 0$; and $\langle x, x \rangle = 0$ iff $x = 0$
2. *Linearity in the first argument*: $\langle \alpha x + y, z \rangle = \alpha \langle x, z \rangle + \langle y, z \rangle$.
3. *Sesquilinearity*: $\langle v, w \rangle = \overline{\langle w, v \rangle}$.

In finite-dimensional vector spaces, any inner product can be identified with a positive-definite Hermitian matrix:

$$\langle v, w \rangle_H = w^* H v.$$

We will have more to say about this later in the class.

A vector space with an inner product automatically inherits the norm $\|v\| = \sqrt{\langle v, v \rangle}$. We are used to seeing this as the standard two-norm with the standard Euclidean inner product, but the same concept works with other inner products. Using an inner product and the associated norm, we can also define the notion of the angle between two vectors:

$$\cos(\theta_{v,w}) = \frac{\langle v, w \rangle}{\|v\| \|w\|}.$$

Problems to ponder

1. The *computational intensity* of an algorithm is the number of flops per memory access. What is the computational intensity of matrix-vector multiplication? Matrix-matrix multiplication?
2. Show that the norm induced by an inner product is indeed a norm.
3. If $\|\cdot\|$ is a norm induced by an inner product, show how to write $\langle v, w \rangle$ using only norms. Hint: try expanding $\|x \pm y\|^2$.
4. Show that $|\cos(\theta_{v,w})| \leq 1$, with equality iff v and w are proportional to each other.
5. What is $\nabla_x \|x\|^2$?