CS 140 : Matrix multiplication

- Warmup: Matrix times vector: communication volume
- Matrix multiplication I: parallel issues
- Matrix multiplication II: cache issues

Thanks to Jim Demmel and Kathy Yelick (UCB) for some of these slides

Matrix-Matrix Multiplication ("DGEMM")

```
{implements C = C + A^*B}
for i = 1 to n
for j = 1 to n
for k = 1 to n
C(i,j) = C(i,j) + A(i,k) * B(k,j)
```

Work: 2*n³ flops Memory: 3*n² words



Parallel matrix multiply, C = C + A*B

- Basic sequential algorithm:
 - C(i,j) += A(i,1) * B(1,j) + A(i,2) * B(1,j) + ... + A(i,n) * B(n,j)
 - work = $t_1 = 2n^3$ floating point ops
- Highly parallel: $t_p = 2n^3 / p$ is easy, for p up to at least n^2
- The issue is communication cost, as affected by:
 - Data layout
 - Structure & schedule of communication
- Where's the data?

Communication volume model

- Network of p processors
 - Each with local memory
 - Message-passing
- Communication volume (v)
 - Total size (words) of all messages passed during computation
 - Broadcasting one word costs volume p (actually, p-1)
- No explicit accounting for communication time
 - Thus, we can't model parallel efficiency or speedup; for that, we'd use the *latency-bandwidth model* (see extra slides)

Parallel Matrix Multiply with 1D Column Layout

• Assume matrices are n x n and n is divisible by p



(A reasonable assumption for analysis, not for code)

- Let A(k) be the n-by-n/p block column that processor k owns
 - similarly B(k) and C(k)

C(k) += A * B(k)

• Now let B(i,k) be a subblock of B(k) with n/p rows

C(k) += A(0) * B(0,k) + A(1) * B(1,k) + ... + A(p-1) * B(p-1,k)

Matmul for 1D layout on a Processor Ring

- Proc k communicates only with procs k-1 and k+1
- Different pairs of processors can communicate simultaneously
- Round-Robin "Merry-Go-Round" algorithm

```
Copy A(myproc) into MGR (MGR = "Merry-Go-Round")

C(myproc) = C(myproc) + MGR*B(myproc, myproc)

for j = 1 to p-1

send MGR to processor myproc+1 mod p (but see deadlock below)

receive MGR from processor myproc-1 mod p (but see below)

C(myproc) = C(myproc) + MGR * B( myproc-j mod p, myproc)
```

- Avoiding deadlock:
 - even procs send then recv, odd procs recv then send
 - or, use nonblocking sends and be careful with buffers
- Comm volume of one inner loop iteration = n^2

Matmul for 1D layout on a Processor Ring

- One iteration: v = n²
- All p-1 iterations: $v = (p-1) * n^2 \sim pn^2$
- Optimal for 1D data layout:
 - Perfect speedup for arithmetic
 - A(myproc) must meet each C(myproc)
- "Nice" communication pattern can probably overlap independent communications in the ring.
- In latency/bandwidth model (see extra slides), parallel efficiency e = 1 - O(p/n)

MatMul with 2D Layout

- Consider processors in 2D grid (physical or logical)
- Processors can communicate with 4 nearest neighbors
 - Alternative pattern: broadcast along rows and columns

p(0,0)	p(0,1)	p(0,2)		p(0,0)	p(0,1)	p(0,2)		p(0,0)	p(0,1)	p(0,2)
p(1,0)	p(1,1)	p(1,2)	=	p(1,0)	p(1,1)	p(1,2)	*	p(1,0)	p(1,1)	p(1,2)
p(2,0)	p(2,1)	p(2,2)		p(2,0)	p(2,1)	p(2,2)		p(2,0)	p(2,1)	p(2,2)

• Assume p is square s x s grid

Cannon's Algorithm: 2-D merry-go-round

... $C(i,j) = C(i,j) + \sum_{k} A(i,k)^*B(k,j)$... assume s = sqrt(p) is an integer forall i=0 to s-1 ... "skew" A left-circular-shift row i of A by i ... so that A(i,j) overwritten by A(i,(j+i)mod s) forall i=0 to s-1 ... "skew" B up-circular-shift B column i of B by i ... so that B(i,j) overwritten by B((i+j)mod s), j) for k=0 to s-1 ... sequential forall i=0 to s-1 and j=0 to s-1 ... all processors in parallel $C(i,j) = C(i,j) + A(i,j)^*B(i,j)$ left-circular-shift each row of A by 1 up-circular-shift each row of B by 1

Cannon's Matrix Multiplication



Cannon's Matrix Multiplication Algorithm

C(1,2) = A(1,0) * B(0,2) + A(1,1) * B(1,2) + A(1,2) * B(2,2)

Initial Step to Skew Matrices in Cannon

Initial blocked input

A(0,0)A(0,1)A(0,2)A(1,0)A(1,1)A(1,2)A(2,0)A(2,1)A(2,2)

B(0,0)	B(0,1)	B(0,2)
B(1,0)	B(1,1)	B(1,2)
B(2,0)	B(2,1)	B(2,2)

• After skewing before initial block multiplies

A(0,0)A(0,1)A(0,2)A(1,1)A(1,2)A(1,0)A(2,2)A(2,0)A(2,1)

B(0,0)	B(1,1)	B(2,2)
B(1,0)	B(2,1)	B(0,2)
B(2,0)	B(0,1)	B(1,2)

Skewing Steps in Cannon

First step	A(0,0)	A(0,1)	A(0,2)
	A(1,1)	A(1,2)	A(1,0)
	A(2,2)	A(2,0)	A(2,1)
Second	A(0,1)	A(0,2)	A(0,0)
	A(1,2)	A(1,0)	A(1,1)
	A(2,0)	A(2,1)	A(2,2)
Third	A(0,2)	A(0,0)	A(0,1)
	A(1,0)	A(1,1)	A(1,2)
	A(2,1)	A(2,2)	A(2,0)

	B(0,0)	B(1,1)	B(2,2)
	B(1,0)	B(2,1)	B(0,2)
	B(2,0)	B(0,1)	B(1,2)
I			
	B(1,0)	B(2,1)	B(0,2)
	B(2,0)	B(0,1)	B(1,2)
	B(0,0)	B(1,1)	B(2,2)
	B(2,0)	B(0,1)	B(1,2)
	B(0,0)	B(1,1)	B(2,2)
	B(1.0)	B(2.1)	B(0.2)

Communication Volume of Cannon's Algorithm

forall i=0 to s-1... recall s = sqrt(p)left-circular-shift row i of A by i... v = n² / s for each iforall i=0 to s-1... v = n² / s for each iup-circular-shift B column i of B by i... v = n² / s for each ifor k=0 to s-1for all i=0 to s-1 and j=0 to s-1C(i,j) = C(i,j) + A(i,j)*B(i,j)... v = n² for each kup-circular-shift each row of A by 1... v = n² for each k

- ° Total comm v = 2*n² + 2*s*n² ~ 2*sqrt(p)*n²
- ° Computational intensity $q = t_1 / v \sim n / sqrt(p)$
- In latency/bandwidth model (see extra slides), parallel efficiency e = 1 - O(sqrt(p) / n)

Cannon is beautiful, but maybe too beautiful ...

- Drawbacks to Cannon:
 - Hard to generalize for p not a perfect square, A & B not square, dimensions not divisible by s = sqrt(p), different memory layouts, etc.
 - Memory hog needs extra copies of local matrices.
- Algorithm used instead in practice is **SUMMA**:
 - uses row and column broadcasts, not merry-go-round
 - see extra slides below for details
- Comparing *computational intensity* = work / comm volume:
 - 1-D MGR has computational intensity q = O(n / p)
 - Cannon has computational intensity q = O(n / sqrt(p))
 - SUMMA has computational intensity q = O(n / sqrt(p)log p)

Sequential Matrix Multiplication

Simple mathematics, but getting good performance is complicated by memory hierarchy --- cache issues.

Naïve 3-loop matrix multiply

```
{implements C = C + A^*B}
for i = 1 to n
for j = 1 to n
for k = 1 to n
C(i,j) = C(i,j) + A(i,k) * B(k,j)
```

Work: 2*n³ flops Memory: 3*n² words



3-Loop Matrix Multiply [Alpern et al., 1992]



O(N³) performance would have constant cycles/flop Performance looks much closer to O(N⁵)

Slide source: Larry Carter, UCSD

3-Loop Matrix Multiply [Alpern et al., 1992]



Slide source: Larry Carter, UCSD

Avoiding data movement: Reuse and locality



- Large memories are slow, fast memories are small
- Parallel processors, collectively, have large, fast cache
 - the slow accesses to "remote" data we call "communication"
- Algorithm should do most work on local data

Simplified model of hierarchical memory

- Assume just 2 levels in the hierarchy, fast and slow
- All data initially in slow memory
 - m = number of memory elements (words) moved between fast and slow memory
 - t_m = time per slow memory operation
 - f = number of arithmetic operations
 - t_f = time per arithmetic operation << t_m
 - q = f / m (computational intensity) flops per slow element access
- Minimum possible time = f* t_f when all data in fast memory
- Actual time
 - $f * t_f + m * t_m = f * t_f * (1 + t_m/t_f * 1/q)$
- Larger q means time closer to minimum $f * t_f$

"Naïve" Matrix Multiply

```
{implements C = C + A^*B}
for i = 1 to n
{read row i of A into fast memory}
for j = 1 to n
{read C(i,j) into fast memory}
{read column j of B into fast memory}
for k = 1 to n
C(i,j) = C(i,j) + A(i,k) * B(k,j)
{write C(i,j) back to slow memory}
```



"Naïve" Matrix Multiply

How many references to slow memory?

m = n³ read each column of B n times + n² read each row of A once + $2n^2$ read and write each element of C once = n³ + $3n^2$

So q = f/m =
$$2n^3 / (n^3 + 3n^2)$$

~= 2 for large n



Blocked Matrix Multiply

Consider A,B,C to be N by N matrices of b by b subblocks where b=n / N is called the block size

```
for i = 1 to N

for j = 1 to N

{read block C(i,j) into fast memory}

for k = 1 to N

{read block A(i,k) into fast memory}

{read block B(k,j) into fast memory}

C(i,j) = C(i,j) + A(i,k) * B(k,j) {do a matrix multiply on blocks}

{write block C(i,j) back to slow memory}
```



Blocked Matrix Multiply

m is amount memory traffic between slow and fast memory matrix has nxn elements, and NxN blocks each of size bxb f is number of floating point operations, $2n^3$ for this problem q = f / m measures data reuse, or computational intensity

Computational intensity $q = f / m = 2n^3 / ((2N + 2) * n^2)$ ~= n / N = b for large n

We can improve performance by increasing the blocksize b (but only until 3b² gets as big as the fast memory size)

Can be much faster than matrix-vector multiply (q=2)

Multi-Level Blocked Matrix Multiply

- More levels of memory hierarchy => more levels of blocking!
- Version 1: One level of blocking for each level of memory (L1 cache, L2 cache, L3 cache, DRAM, disk, ...)
 - Version 2: Recursive blocking, O(log n) levels deep

In the "Uniform Memory Hierarchy" cost model, the 3-loop algorithm is $O(N^5)$ time, but the blocked algorithms are $O(N^3)$

BLAS: Basic Linear Algebra Subroutines

- Industry standard interface
- Vendors, others supply optimized implementations
- History
 - BLAS1 (1970s):
 - vector operations: dot product, saxpy (y= α *x+y), etc
 - m=2*n, f=2*n, q ~1 or less
 - BLAS2 (mid 1980s)
 - matrix-vector operations: matrix vector multiply, etc
 - m=n², f=2*n², q~2, less overhead
 - somewhat faster than BLAS1
 - BLAS3 (late 1980s)
 - matrix-matrix operations: matrix matrix multiply, etc
 - m >= n^2, f=O(n^3), so q can possibly be as large as n
 - BLAS3 is potentially much faster than BLAS2
- Good algorithms use BLAS3 when possible (LAPACK)
 - See www.netlib.org/blas, www.netlib.org/lapack

BLAS speeds on an IBM RS6000/590



ScaLAPACK Parallel Library

ScalaPACK SOFTWARE HIERARCHY



Extra Slides:

Parallel matrix multiplication in the latency-bandwidth cost model

Latency Bandwidth Model

- Network of p processors, each with local memory
 - Message-passing
- Latency (α)
 - Cost of communication per message
- Inverse bandwidth (β)
 - Cost of communication per unit of data
- Parallel time (t_p)
 - Computation time plus communication time
- Parallel efficiency:
 - $e(p) = t_1 / (p * t_p)$
 - perfect speedup $\rightarrow e(p) = 1$

Matrix Multiply with 1D Column Layout

• Assume matrices are n x n and n is divisible by p



May be a reasonable assumption for analysis, not for code

- A(k) is the n-by-n/p block column that processor k owns (similarly B(i) and C(i))
- B(i,k) is an n/p-by-n/p sublock of B(k)
 - in rows i*n/p through (i+1)*n/p
- Formula: $C(k) = C(k) + A^*B(k) = C(k) + \sum_{i=0:p} A(i) * B(i,k)$

Matmul for 1D layout on a Processor Ring

- Proc k communicates only with procs k-1 and k+1
- Different pairs of processors can communicate simultaneously
- Round-Robin "Merry-Go-Round" algorithm

```
Copy A(myproc) into MGR (MGR = "Merry-Go-Round")

C(myproc) = C(myproc) + MGR*B(myproc, myproc)

for j = 1 to p-1

send MGR to processor myproc+1 mod p (but see deadlock below)

receive MGR from processor myproc-1 mod p (but see below)

C(myproc) = C(myproc) + MGR * B( myproc-j mod p, myproc)
```

- Avoiding deadlock:
 - even procs send then recv, odd procs recv then send
 - or, use nonblocking sends

• Time of inner loop =
$$2^{(\alpha + \beta n^2/p)} + 2^{n^{(n/p)^2}}$$

Matmul for 1D layout on a Processor Ring

- Time of inner loop = $2^{(\alpha + \beta^{n^2/p})} + 2^{n^*(n/p)^2}$
- Total Time = 2*n* (n/p)² + (p-1) * Time of inner loop

~
$$2^{n_3/p}$$
 + $2^{p_1} \alpha$ + $2^{r_2} \beta^{r_2}$

- Optimal for 1D layout on Ring or Bus, even with broadcast:
 - Perfect speedup for arithmetic

۲

- A(myproc) must move to each other processor, costs at least (p-1)*cost of sending n*(n/p) words
- Parallel Efficiency = $2^n^3 / (p * \text{Total Time})$ = $1/(1 + \alpha * p^2/(2^n^3) + \beta * p/(2^n))$ = 1/(1 + O(p/n))= 1 - O(p/n)
- Grows to 1 as n/p increases (or α and β shrink)

MatMul with 2D Layout

- Consider processors in 2D grid (physical or logical)
- Processors can communicate with 4 nearest neighbors
 - Alternative pattern: broadcast along rows and columns

p(0,0)	p(0,1)	p(0,2)		p(0,0)	p(0,1)	p(0,2)		p(0,0)	p(0,1)	p(0,2)
p(1,0)	p(1,1)	p(1,2)	=	p(1,0)	p(1,1)	p(1,2)	*	p(1,0)	p(1,1)	p(1,2)
p(2,0)	p(2,1)	p(2,2)		p(2,0)	p(2,1)	p(2,2)		p(2,0)	p(2,1)	p(2,2)

• Assume p is square s x s grid

Cannon's Algorithm: 2-D merry-go-round

... $C(i,j) = C(i,j) + \sum_{k} A(i,k)^*B(k,j)$... assume s = sqrt(p) is an integer forall i=0 to s-1 ... "skew" A left-circular-shift row i of A by i ... so that A(i,j) overwritten by A(i,(j+i)mod s) forall i=0 to s-1 ... "skew" B up-circular-shift B column i of B by i ... so that B(i,j) overwritten by B((i+j)mod s), j) for k=0 to s-1 ... sequential forall i=0 to s-1 and j=0 to s-1 ... all processors in parallel $C(i,j) = C(i,j) + A(i,j)^*B(i,j)$ left-circular-shift each row of A by 1 up-circular-shift each row of B by 1

Cannon's Matrix Multiplication



Cannon's Matrix Multiplication Algorithm

C(1,2) = A(1,0) * B(0,2) + A(1,1) * B(1,2) + A(1,2) * B(2,2)

Initial Step to Skew Matrices in Cannon

• Initial blocked input

A(0,0)A(0,1)A(0,2)A(1,0)A(1,1)A(1,2)A(2,0)A(2,1)A(2,2)

B(0,0)	B(0,1)	B(0,2)
B(1,0)	B(1,1)	B(1,2)
B(2,0)	B(2,1)	B(2,2)

After skewing before initial block multiplies

A(0,0)	A(0,1)	A(0,2)
A(1,1)	A(1,2)	A(1,0)
A(2,2)	A(2,0)	A(2,1)

C	DIOCK	multi	plies
	B(0,0)	B(1,1)	B(2,2)
	B(1,0)	B(2,1)	B(0,2)
	B(2,0)	B(0,1)	B(1,2)

Skewing Steps in Cannon

•	First step	A(0,0)	A(0,1)	A(0,2)
		A(1,1)	A(1,2)	A(1,0)
		A(2,2)	A(2,0)	A(2,1)
•	Second	A(0,1)	A(0,2)	A(0,0)
		A(1,2)	A(1,0)	A(1,1)
	Third	A(2,0)	A(2,1)	A(2,2)
•	THIL	A(0,2)	A(0,0)	A(0,1)
		A(1,0)	A(1,1)	A(1,2)
		A(2,1)	A(2,2)	A(2,0)

B(0,0)	B(1,1)	B(2,2)
B(1,0)	B(2,1)	B(0,2)
B(2,0)	B(0,1)	B(1,2)
B(1,0)	B(2,1)	B(0,2)
B(2,0)	B(0,1)	B(1,2)
B(0,0)	B(1,1)	B(2,2)
B(2,0)	B(0,1)	B(1,2)
B(0,0)	B(1,1)	B(2,2)
B(1,0)	B(2,1)	B(0.2)

Cost of Cannon's Algorithm

```
forall i=0 to s-1 ... recall s = sqrt(p)

left-circular-shift row i of A by i ... cost = s*(\alpha + \beta*n2/p)

forall i=0 to s-1

up-circular-shift B column i of B by i ... cost = s*(\alpha + \beta*n2/p)

for k=0 to s-1

forall i=0 to s-1 and j=0 to s-1

C(i,j) = C(i,j) + A(i,j)*B(i,j) ... cost = 2*(n/s)^3 = 2*n3/p^{3/2}

left-circular-shift each row of A by 1 ... cost = \alpha + \beta*n2/p

up-circular-shift each row of B by 1 ... cost = \alpha + \beta*n2/p
```

```
° Total Time = 2*n^3/p + 4*s*\alpha + 4*\beta*n^2/s

° Parallel Efficiency = 2*n^3 / (p * Total Time)

= 1/(1 + \alpha * 2*(s/n)^3 + \beta * 2*(s/n))

= 1 - O(sqrt(p)/n)

° Grows to 1 as n/s = n/sqrt(p) = sqrt(data per processor) grows

° Better than 1D layout, which had Efficiency = 1 - O(p/n)
```

Extra Slides:

SUMMA parallel matrix multiplication algorithm

SUMMA Algorithm

- SUMMA = Scalable Universal Matrix Multiply
- Slightly less efficient than Cannon ... but simpler and easier to generalize
- Presentation from van de Geijn and Watts
 - www.netlib.org/lapack/lawns/lawn96.ps
 - Similar ideas appeared many times
- Used in practice in PBLAS = Parallel BLAS
 - www.netlib.org/lapack/lawns/lawn100.ps



- I, J represent all rows, columns owned by a processor
- k is a single row or column
 - or a block of b rows or columns
- $C(I,J) = C(I,J) + \Sigma_k A(I,k)^*B(k,J)$
- Assume a p_r by p_c processor grid (p_r = p_c = 4 above)
 Need not be square

```
SUMMA
                                        J ∕B(k,J)
                                           k
                        *
                                            =
                                                           C(I,J)
A(I,k)
   For k=0 to n-1 ... or n/b-1 where b is the block size
                    \dots = # cols in A(I,k) and # rows in B(k,J)
      for all I = 1 to p_r \dots in parallel
          owner of A(I,k) broadcasts it to whole processor row
      for all J = 1 to p_C ... in parallel
          owner of B(k,J) broadcasts it to whole processor column
      Receive A(I,k) into Acol
      Receive B(k,J) into Brow
      C(myproc, myproc) = C(myproc, myproc) + Acol * Brow
```

SUMMA performance

To simplify analysis only, assume s = sqrt(p)

```
For k=0 to n/b-1
   for all I = 1 to s ... s = sqrt(p)
       owner of A(I,k) broadcasts it to whole processor row
         ... time = log s *( \alpha + \beta * b*n/s), using a tree
   for all J = 1 to s
       owner of B(k,J) broadcasts it to whole processor column
         ... time = log s *( \alpha + \beta * b*n/s), using a tree
   Receive A(I,k) into Acol
   Receive B(k,J) into Brow
   C(myproc, myproc) = C(myproc, myproc) + Acol * Brow
         ... time = 2^{(n/s)^{2}}b
```

° Total time = $2*n^3/p + \alpha * \log p * n/b + \beta * \log p * n^2 /s$

SUMMA performance

- Total time = $2*n^3/p + \alpha * \log p * n/b + \beta * \log p * n^2 /s$
- Parallel Efficiency = $1/(1 + \alpha * \log p * p / (2*\beta*n^2) + \beta * \log p * s/(2*n))$
- ~Same β term as Cannon, except for log p factor log p grows slowly so this is ok
- Latency (α) term can be larger, depending on b When b=1, get α * log p * n As b grows to n/s, term shrinks to α * log p * s (log p times Cannon)
- Temporary storage grows like 2*b*n/s
- Can change b to tradeoff latency cost with memory