
CS 267

Dense Linear Algebra: History and Structure, Parallel Matrix Multiplication

James Demmel

www.cs.berkeley.edu/~demmel/cs267_Spr15

02/26/2015

CS267 Lecture 12

1

Quick review of earlier lecture

- What do you call
 - A program written in PyGAS, a Global Address Space language based on Python...
 - That uses a Monte Carlo simulation algorithm to approximate π ...
 - That has a race condition, so that it gives you a different funny answer every time you run it?

Monte - π - thon

02/26/2015

CS267 Lecture 12

2

Outline

- History and motivation
 - What is dense linear algebra?
 - Why minimize communication?
 - Lower bound on communication
- Structure of the Dense Linear Algebra motif
 - What does A\b do?
- Parallel Matrix-matrix multiplication
 - Attaining the lower bound
- Other Parallel Algorithms (next lecture)

02/26/2015

CS267 Lecture 12

3

Outline

- History and motivation
 - What is dense linear algebra?
 - Why minimize communication?
 - Lower bound on communication
- Structure of the Dense Linear Algebra motif
 - What does A\b do?
- Parallel Matrix-matrix multiplication
 - Attaining the lower bound
- Other Parallel Algorithms (next lecture)

02/26/2015

CS267 Lecture 12

4

Motifs

The Motifs (formerly “Dwarfs”) from “The Berkeley View” (Asanovic et al.)
Motifs form key computational patterns

	Embed	SPEC	DB	Games	ML	HPC	Health	Image	Speech	Music	Browser
Finite State Mach.	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
Circuits	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
Graph Algorithms	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
Structured Grid	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
Dense Matrix	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
Sparse matrix	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
Spectral (FFT)	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
Dynamic Prog	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
N-Body	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
Backtrack/ B&B	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
Graphical Models	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
Unstructured Grid	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red

What is dense linear algebra?

- Not just matmul!
- Linear Systems: $Ax=b$
- Least Squares: choose x to minimize $\|Ax-b\|_2$
 - Overdetermined or underdetermined
 - Unconstrained, constrained, weighted
- Eigenvalues and vectors of Symmetric Matrices
 - Standard ($Ax = \lambda x$), Generalized ($Ax=ABx$)
- Eigenvalues and vectors of Unsymmetric matrices
 - Eigenvalues, Schur form, eigenvectors, invariant subspaces
 - Standard, Generalized
- Singular Values and vectors (SVD)
 - Standard, Generalized
- Different matrix structures
 - Real, complex; Symmetric, Hermitian, positive definite; dense, triangular, banded ...
- Level of detail
 - Simple Driver (“ $x=Ab$ ”)
 - Expert Drivers with error bounds, extra-precision, other options
 - Lower level routines (“apply certain kind of orthogonal transformation”, matmul...)

02/26/2015

CS267 Lecture 13

6

A brief history of (Dense) Linear Algebra software (1/7)

- In the beginning was the do-loop...
 - Libraries like EISPACK (for eigenvalue problems)
- Then the BLAS (1) were invented (1973-1977)
 - Standard library of 15 operations (mostly) on vectors
 - “AXPY” ($y = \alpha \cdot x + y$), dot product, scale ($x = \alpha \cdot x$), etc
 - Up to 4 versions of each (S/D/C/Z), 46 routines, 3300 LOC
 - Goals
 - Common “pattern” to ease programming, readability
 - Robustness, via careful coding (avoiding over/underflow)
 - Portability + Efficiency via machine specific implementations
 - Why BLAS 1 ? They do $O(n^1)$ ops on $O(n^1)$ data
 - Used in libraries like LINPACK (for linear systems)
 - Source of the name “LINPACK Benchmark” (not the code!)

02/26/2015

CS267 Lecture 12

7

Current Records for Solving Dense Systems (11/2014)

- Linpack Benchmark
- Fastest machine overall (www.top500.org)
 - Tianhe-2 (Guangzhou, China)
 - 33.9 Petaflops out of 54.9 Petaflops peak ($n=10M$)
 - 3.1M cores, of which 2.7M are accelerator cores
 - Intel Xeon E5-2692 (Ivy Bridge) and Xeon Phi 31S1P
 - 1 Pbyte memory
 - 17.8 MWatts of power, 1.9 Gflops/Watt
- Historical data (www.netlib.org/performance)
 - Palm Pilot III
 - 1.69 Kiloflops
 - $n = 100$

02/26/2015

CS267 Lecture 12

8

A brief history of (Dense) Linear Algebra software (5/7)

- Is LAPACK parallel?
 - Only if the BLAS are parallel (possible in shared memory)
- ScaLAPACK – “Scalable LAPACK” (1995 – now)
 - For distributed memory – uses MPI
 - More complex data structures, algorithms than LAPACK
 - Only (small) subset of LAPACK’s functionality available
 - Details later (class projects!)
 - All at www.netlib.org/scalapack


02/26/2015

CS267 Lecture 12

13

Success Stories for Sca/LAPACK (6/7)

- Widely used
 - Adopted by Mathworks, Cray, Fujitsu, HP, IBM, IMSL, Intel, NAG, NEC, SGI, ...
 - 7.5M webhits/year @ Netlib (incl. CLAPACK, LAPACK95)
- New Science discovered through the solution of dense matrix systems
 - Nature article on the flat universe used ScaLAPACK
 - Other articles in Physics Review B that also use it
 - 1998 Gordon Bell Prize
 - www.nersc.gov/news/reports/newNERSResults050703.pdf



Cosmic Microwave Background Analysis, BOOMERanG collaboration, MADCAP code (Apr. 27, 2000).

02/26/2015

CS267 Lecture 12

14

A brief future look at (Dense) Linear Algebra software (7/7)


- PLASMA, DPLASMA and MAGMA (now)
 - Ongoing extensions to Multicore/GPU/Heterogeneous
 - Can one software infrastructure accommodate all algorithms and platforms of current (future) interest?
 - How much code generation and tuning can we automate?
 - Details later (Class projects!) (icl.cs.utk.edu/plasma, [magma](http://icl.cs.utk.edu/magma))
- Other related projects
 - Elemental (libelemental.org)
 - Distributed memory dense linear algebra
 - “Balance ease of use and high performance”
 - FLAME (z.cs.utexas.edu/wiki/flame/wiki/FrontPage)
 - Formal Linear Algebra Method Environment
 - Attempt to automate code generation across multiple platforms
 - BLAST Forum (www.netlib.org/blas/blast-forum)
 - Attempt to extend BLAS, add new functions, extra-precision, ...

Back to basics:

Why avoiding communication is important (1/3)

Algorithms have two costs:

1. Arithmetic (FLOPS)
2. Communication: moving data between
 - levels of a memory hierarchy (sequential case)
 - processors over a network (parallel case).



02/26/2015

CS267 Lecture 12

16

Why avoiding communication is important (2/3)

- Running time of an algorithm is sum of 3 terms:
 - # flops * time_per_flop
 - # words moved / bandwidth } **communication**
 - # messages * latency
- Time_per_flop \ll 1/ bandwidth \ll latency
- Gaps growing exponentially with time

Time_per_flop	Annual improvements		
	DRAM	Bandwidth	Latency
59%	26%	23%	15%
	Network		5%

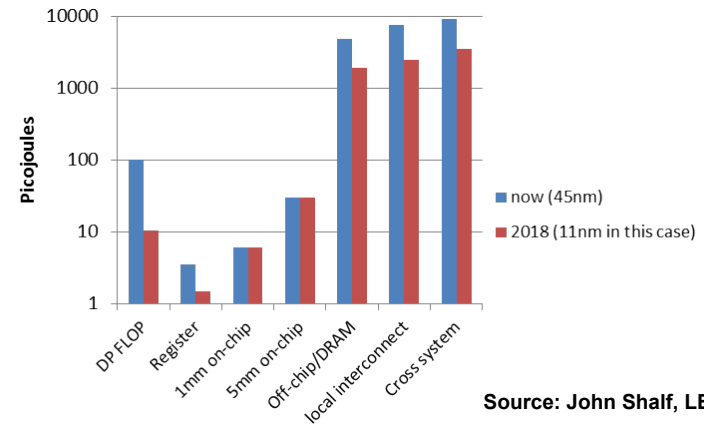
- Minimize communication to save time**

02/26/2015

CS267 Lecture 12

17

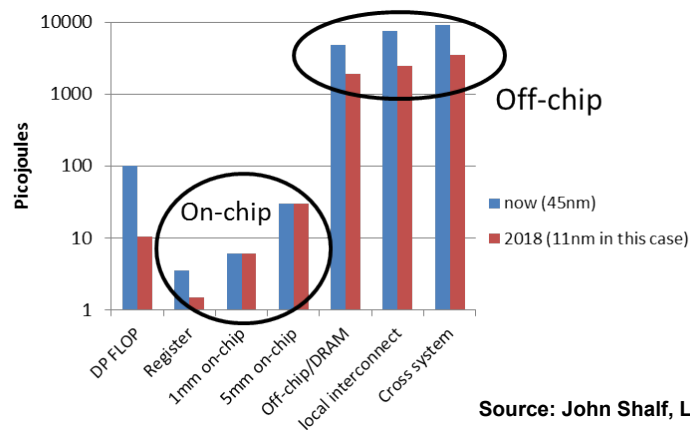
Why Minimize Communication? (3/3)



Source: John Shalf, LBL

Why Minimize Communication? (3/3)

Minimize communication to save energy



Source: John Shalf, LBL

Goal: Organize Linear Algebra to Avoid Communication

- Between all memory hierarchy levels
 - L1 \leftrightarrow L2 \leftrightarrow DRAM \leftrightarrow network, etc
- Not just *hiding* communication (overlap with arithmetic)
 - Speedup \leq 2x
- Arbitrary speedups/energy savings possible
- Later: Same goal for other computational patterns
 - Lots of open problems

02/26/2015

CS267 Lecture 12

20

Review: Blocked Matrix Multiply

- Blocked Matmul $C = A \cdot B$ breaks A, B and C into blocks with dimensions that depend on cache size

... Break $A^{n \times n}$, $B^{n \times n}$, $C^{n \times n}$ into $b \times b$ blocks labeled $A(i,j)$, etc

... b chosen so 3 $b \times b$ blocks fit in cache

for $i = 1$ to n/b , for $j=1$ to n/b , for $k=1$ to n/b

$C(i,j) = C(i,j) + A(i,k) \cdot B(k,j)$... $b \times b$ matmul, $4b^2$ reads/writes

- When $b=1$, get “naïve” algorithm, want b larger ...
- $(n/b)^3 \cdot 4b^2 = 4n^3/b$ reads/writes altogether
- Minimized when $3b^2 = \text{cache size} = M$, yielding $O(n^3/M^{1/2})$ reads/writes
- What if we had more levels of memory? (L1, L2, cache etc)?
 - Would need 3 more nested loops per level
 - Recursive (cache-oblivious algorithm) also possible

02/26/2015

CS267 Lecture 12

21

Communication Lower Bounds: Prior Work on Matmul

- Assume n^3 algorithm (i.e. not Strassen-like)
- Sequential case, with fast memory of size M
 - Lower bound on #words moved to/from slow memory = $\Omega(n^3 / M^{1/2})$ [Hong, Kung, 81]
 - Attained using blocked or cache-oblivious algorithms
- Parallel case on P processors:
 - Let M be memory per processor; assume load balanced
 - Lower bound on #words moved = $\Omega(n^3 / (p \cdot M^{1/2}))$ [Irony, Tiskin, Toledo, 04]
 - If $M = 3n^2/p$ (one copy of each matrix), then lower bound = $\Omega(n^2 / p^{1/2})$
 - Attained by SUMMA, Cannon's algorithm

02/26/2015

CS267 Lecture 12

22

New lower bound for all “direct” linear algebra

Let M = “fast” memory size per processor
= cache size (sequential case) or $O(n^2/p)$ (parallel case)
#flops = number of flops done per processor

#words_moved per processor = $\Omega(\text{\#flops} / M^{1/2})$

#messages_sent per processor = $\Omega(\text{\#flops} / M^{3/2})$

- Holds for
 - Matmul, BLAS, LU, QR, eig, SVD, tensor contractions, ...
 - Some whole programs (sequences of these operations, no matter how they are interleaved, eg computing A^k)
 - Dense *and* sparse matrices (where #flops $\ll n^3$)
 - Sequential *and* parallel algorithms
 - Some graph-theoretic algorithms (eg Floyd-Warshall)
- Generalizations later (Strassen-like algorithms, loops accessing arrays)

02/26/2015

CS267 Lecture 12

23

New lower bound for all “direct” linear algebra

Let M = “fast” memory size per processor
= cache size (sequential case) or $O(n^2/p)$ (parallel case)
#flops = number of flops done per processor

#words_moved per processor = $\Omega(\text{\#flops} / M^{1/2})$

#messages_sent per processor = $\Omega(\text{\#flops} / M^{3/2})$

- Sequential case, dense $n \times n$ matrices, so $O(n^3)$ flops
 - #words_moved = $\Omega(n^3 / M^{1/2})$
 - #messages_sent = $\Omega(n^3 / M^{3/2})$
- Parallel case, dense $n \times n$ matrices
 - Load balanced, so $O(n^3/p)$ flops processor
 - One copy of data, load balanced, so $M = O(n^2/p)$ per processor
 - #words_moved = $\Omega(n^2 / p^{1/2})$
 - #messages_sent = $\Omega(p^{1/2})$ **SIAM Linear Algebra Prize, 2012**

02/26/2015

CS267 Lecture 12

24

Can we attain these lower bounds?

- Do conventional dense algorithms as implemented in LAPACK and ScaLAPACK attain these bounds?
 - Mostly not yet, work in progress
- If not, are there other algorithms that do?
 - Yes
- Goals for algorithms:
 - Minimize #words_moved
 - Minimize #messages_sent
 - Need new data structures
 - Minimize for multiple memory hierarchy levels
 - Cache-oblivious algorithms would be simplest
 - Fewest flops when matrix fits in fastest memory
 - Cache-oblivious algorithms don't always attain this
- Attainable for nearly all dense linear algebra
 - Just a few prototype implementations so far (class projects!)
 - Only a few sparse algorithms so far (eg Cholesky)

02/26/2015

CS267 Lecture 12

25

Outline

- History and motivation
 - What is dense linear algebra?
 - Why minimize communication?
 - Lower bound on communication
- Structure of the Dense Linear Algebra motif
 - What does A\b do?
- Parallel Matrix-matrix multiplication
 - Attaining the lower bound
 - Proof of the lower bound (if time)
- Other Parallel Algorithms (next lecture)

02/26/2015

CS267 Lecture 12

26

What could go into the linear algebra motif(s)?

For all linear algebra problems

For all matrix/problem structures

For all data types

For all architectures and networks

For all programming interfaces

Produce best algorithm(s) w.r.t.
performance and/or accuracy
(including error bounds, etc)

Need to prioritize, automate!

02/26/2015

CS267 Lecture 12

27

For all linear algebra problems: Ex: LAPACK Table of Contents

- Linear Systems
- Least Squares
 - Overdetermined, underdetermined
 - Unconstrained, constrained, weighted
- Eigenvalues and vectors of Symmetric Matrices
 - Standard ($Ax = \lambda x$), Generalized ($Ax = \lambda Bx$)
- Eigenvalues and vectors of Unsymmetric matrices
 - Eigenvalues, Schur form, eigenvectors, invariant subspaces
 - Standard, Generalized
- Singular Values and vectors (SVD)
 - Standard, Generalized
- Level of detail
 - Simple Driver
 - Expert Drivers with error bounds, extra-precision, other options
 - Lower level routines ("apply certain kind of orthogonal transformation")

02/26/2015

CS267 Lecture 12

28

What does A\b do? What could it do? Ex: LAPACK Table of Contents

- BD – bidiagonal
- GB – general banded
- GE – general
- GG – general , pair
- GT – tridiagonal
- HB – Hermitian banded
- HE – Hermitian
- HG – upper Hessenberg, pair
- HP – Hermitian, packed
- HS – upper Hessenberg
- OR – (real) orthogonal
- OP – (real) orthogonal, packed
- PB – positive definite, banded
- PO – positive definite
- PP – positive definite, packed
- PT – positive definite, tridiagonal
- SB – symmetric, banded
- SP – symmetric, packed
- ST – symmetric, tridiagonal
- SY – symmetric
- TB – triangular, banded
- TG – triangular, pair
- TP – triangular, packed
- TR – triangular
- TZ – trapezoidal
- UN – unitary
- UP – unitary packed

02/26/2015

CS267 Lecture 12

29

What does A\b do? What could it do? Ex: LAPACK Table of Contents

- BD – bidiagonal
- GB – general banded
- **GE – general**
- GG – general , pair
- GT – tridiagonal
- HB – Hermitian banded
- HE – Hermitian
- HG – upper Hessenberg, pair
- HP – Hermitian, packed
- HS – upper Hessenberg
- OR – (real) orthogonal
- OP – (real) orthogonal, packed
- PB – positive definite, banded
- PO – positive definite
- PP – positive definite, packed
- PT – positive definite, tridiagonal
- SB – symmetric, banded
- SP – symmetric, packed
- ST – symmetric, tridiagonal
- SY – symmetric
- TB – triangular, banded
- TG – triangular, pair
- TP – triangular, packed
- TR – triangular
- TZ – trapezoidal
- UN – unitary
- UP – unitary packed

02/26/2015

CS267 Lecture 12

30

What does A\b do? What could it do? Ex: LAPACK Table of Contents

- BD – bidiagonal
- GB – general banded
- GE – general
- GG – general, pair
- GT – tridiagonal
- **HB – Hermitian** banded
- **HE – Hermitian**
- HG – upper Hessenberg, pair
- **HP – Hermitian**, packed
- HS – upper Hessenberg
- OR – (real) orthogonal
- OP – (real) orthogonal, packed
- PB – positive definite, banded
- PO – positive definite
- PP – positive definite, packed
- PT – positive definite, tridiagonal
- **SB – symmetric**, banded
- **SP – symmetric**, packed
- **ST – symmetric**, tridiagonal
- **SY – symmetric**
- TB – triangular, banded
- TG – triangular, pair
- TP – triangular, packed
- TR – triangular
- TZ – trapezoidal
- UN – unitary
- UP – unitary packed

02/26/2015

CS267 Lecture 12

31

What does A\b do? What could it do? Ex: LAPACK Table of Contents

- BD – bidiagonal
- GB – general banded
- GE – general
- GG – general, pair
- GT – tridiagonal
- HB – Hermitian banded
- HE – Hermitian
- HG – upper Hessenberg, pair
- HP – Hermitian, packed
- HS – upper Hessenberg
- OR – (real) orthogonal
- OP – (real) orthogonal, packed
- **PB – positive definite**, banded
- **PO – positive definite**
- **PP – positive definite**, packed
- **PT – positive definite**, tridiagonal
- SB – symmetric, banded
- SP – symmetric, packed
- ST – symmetric, tridiagonal
- SY – symmetric
- TB – triangular, banded
- TG – triangular, pair
- TP – triangular, packed
- TR – triangular
- TZ – trapezoidal
- UN – unitary
- UP – unitary packed

02/26/2015

CS267 Lecture 12

32

What does A\b do? What could it do? Ex: LAPACK Table of Contents


- BD – bidiagonal
- GB – general banded
- GE – general
- GG – general, pair
- GT – tridiagonal
- HB – Hermitian banded
- HE – Hermitian
- HG – upper Hessenberg, pair
- HP – Hermitian, packed
- HS – upper Hessenberg
- OR – (real) orthogonal
- OP – (real) orthogonal, packed
- PB – positive definite, banded
- PO – positive definite
- PP – positive definite, packed
- PT – positive definite, tridiagonal
- SB – symmetric, banded
- SP – symmetric, packed
- ST – symmetric, tridiagonal
- SY – symmetric
- **TB – triangular, banded**
- **TG – triangular, pair**
- **TP – triangular, packed**
- **TR – triangular**
- **TZ – trapezoidal**
- UN – unitary
- UP – unitary packed

02/26/2015


CS267 Lecture 12

33

Organizing Linear Algebra – in books



www.netlib.org/lapack



gams.nist.gov



www.netlib.org/scalapack



www.netlib.org/templates



www.cs.utk.edu/~dongarra/etemplates

Outline


- History and motivation
 - What is dense linear algebra?
 - Why minimize communication?
 - Lower bound on communication
- Structure of the Dense Linear Algebra motif
 - What does A\b do?
- Parallel Matrix-matrix multiplication
 - Attaining the lower bound
- Other Parallel Algorithms (next lecture)

02/26/2015


CS267 Lecture 12

35


Different Parallel Data Layouts for Matrices (not all!)



1) 1D Column Blocked Layout




2) 1D Column Cyclic Layout




3) 1D Column Block Cyclic Layout

4) Row versions of the previous layouts



5) 2D Row and Column Blocked Layout



Generalizes others

6) 2D Row and Column Block Cyclic Layout

02/26/2015

CS267 Lecture 12

36

Parallel Matrix-Vector Product

- Compute $y = y + A*x$, where A is a dense matrix
- Layout:
 - **1D row blocked**

- $A(i)$ refers to the n by n/p block row that processor i owns,


- $x(i)$ and $y(i)$ similarly refer to segments of x, y owned by i

Algorithm:

- **Foreach processor i**
- **Broadcast $x(i)$**
- **Compute $y(i) = A(i)*x$**

- Algorithm uses the formula

$$y(i) = y(i) + A(i)*x = y(i) + \sum_j A(i,j)*x(j)$$




02/26/2015

CS267 Lecture 12

37

Matrix-Vector Product $y = y + A*x$

- A **column layout** of the matrix eliminates the broadcast of x
 - But adds a reduction to update the destination y
- A **2D blocked layout** uses a broadcast and reduction, both on a subset of processors
 - \sqrt{p} for square processor grid



02/26/2015

CS267 Lecture 12

38

Parallel Matrix Multiply

- Computing $C=C+A*B$
- Using basic algorithm: $2*n^3$ Flops
- Variables are:

- Data layout: 1D? 2D? Other?
- Topology of machine: Ring? Torus?
- Scheduling communication

- Use of performance models for algorithm design

$$\text{Message Time} = \text{"latency"} + \text{\#words} * \text{time-per-word} \\ = \alpha + n*\beta$$

- Efficiency (in any model):

- serial time / (p * parallel time)
- perfect (linear) speedup \leftrightarrow efficiency = 1


02/26/2015

CS267 Lecture 12

39

Matrix Multiply with 1D Column Layout

- Assume matrices are $n \times n$ and n is divisible by p



May be a reasonable assumption for analysis, not for code

- $A(i)$ refers to the n by n/p block column that processor i owns (similarly for $B(i)$ and $C(i)$)
- $B(i,j)$ is the n/p by n/p subblock of $B(i)$
 - in rows $j*n/p$ through $(j+1)*n/p - 1$

- Algorithm uses the formula

$$C(i) = C(i) + A*B(i) = C(i) + \sum_j A(j)*B(j,i)$$

02/26/2015

CS267 Lecture 12

40

Matrix Multiply: 1D Layout on Bus or Ring

- Algorithm uses the formula
$$C(i) = C(i) + A * B(i) = C(i) + \sum_j A(j) * B(j,i)$$
- First consider a bus-connected machine without broadcast: only one pair of processors can communicate at a time (ethernet)
- Second consider a machine with processors on a ring: all processors may communicate with nearest neighbors simultaneously

02/26/2015

CS267 Lecture 12

41

MatMul: 1D layout on Bus without Broadcast

Naïve algorithm:

```
C(myproc) = C(myproc) + A(myproc)*B(myproc,myproc)
for i = 0 to p-1
  for j = 0 to p-1 except i
    if (myproc == i) send A(i) to processor j
    if (myproc == j)
      receive A(i) from processor i
      C(myproc) = C(myproc) + A(i)*B(i,myproc)
  barrier
```

Cost of inner loop:

computation: $2 * n * (n/p)^2 = 2 * n^3 / p^2$
communication: $\alpha + \beta * n^2 / p$

02/26/2015

CS267 Lecture 12

42

Naïve MatMul (continued)

Cost of inner loop:

computation: $2 * n * (n/p)^2 = 2 * n^3 / p^2$
communication: $\alpha + \beta * n^2 / p$... approximately

Only 1 pair of processors (i and j) are active on any iteration,
and of those, only i is doing computation

=> the algorithm is almost entirely serial

Running time:

$= (p * (p-1) + 1) * \text{computation} + p * (p-1) * \text{communication}$
 $\approx 2 * n^3 + p^2 * \alpha + p * n^2 * \beta$

This is worse than the serial time and grows with p.

02/26/2015

CS267 Lecture 12

43

Matmul for 1D layout on a Processor Ring

- Pairs of adjacent processors can communicate simultaneously

```
Copy A(myproc) into Tmp
C(myproc) = C(myproc) + Tmp*B(myproc , myproc)
for j = 1 to p-1
  Send Tmp to processor myproc+1 mod p
  Receive Tmp from processor myproc-1 mod p
  C(myproc) = C(myproc) + Tmp*B( myproc-j mod p , myproc)
```

- Same idea as for gravity in simple sharks and fish algorithm
 - May want double buffering in practice for overlap
 - Ignoring deadlock details in code
- Time of inner loop = $2 * (\alpha + \beta * n^2 / p) + 2 * n * (n/p)^2$

02/26/2015

CS267 Lecture 12

44

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 = $2*n*(n/p)^2 + (p-1) * \text{Time of inner loop}$
- $\approx 2*n^3/p + 2*p*\alpha + 2*\beta*n^2$
- (Nearly) Optimal for 1D layout on Ring or Bus, even with Broadcast:
 - Perfect speedup for arithmetic
 - A(myproc) must move to each other processor, costs at least $(p-1)*\text{cost of sending } n*(n/p) \text{ words}$
- Parallel Efficiency = $2*n^3 / (p * \text{Total Time})$
 $= 1/(1 + \alpha * p^2/(2*n^3) + \beta * p/(2*n))$
 $= 1/(1 + O(p/n))$
- Grows to 1 as n/p increases (or α and β shrink)
- But far from communication lower bound

02/26/2015

CS267 Lecture 12

45

Need to try 2D Matrix layout

1) 1D Column Blocked Layout

2) 1D Column Cyclic Layout

3) 1D Column Block Cyclic Layout

4) Row versions of the previous layouts

5) 2D Row and Column Blocked Layout

6) 2D Row and Column Block Cyclic Layout

Generalizes others

02/26/2015

CS267 Lecture 12

46

Summary of Parallel Matrix Multiply

- SUMMA
 - Scalable Universal Matrix Multiply Algorithm
 - Attains communication lower bounds (within $\log p$)
- Cannon
 - Historically first, attains lower bounds
 - More assumptions
 - A and B square
 - P a perfect square
- 2.5D SUMMA
 - Uses more memory to communicate even less
- Parallel Strassen
 - Attains different, even lower bounds

02/26/2015

CS267 Lecture 12

47

SUMMA Algorithm

- SUMMA = Scalable Universal Matrix Multiply
- 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

02/26/2015

CS267 Lecture 12

48

SUMMA uses Outer Product form of MatMul

- $C = A*B$ means $C(i,j) = \sum_k A(i,k)*B(k,j)$
- Column-wise outer product:

$$C = A*B$$

$$= \sum_k A(:,k)*B(k,:)$$

$$= \sum_k (\text{k-th col of } A) * (\text{k-th row of } B)$$
- Block column-wise outer product
(block size = 4 for illustration)

$$C = A*B$$

$$= A(:,1:4)*B(1:4,:) + A(:,5:8)*B(5:8,:) + \dots$$


$$= \sum_k (\text{k-th block of 4 cols of } A) * (\text{k-th block of 4 rows of } B)$$

02/26/2015

CS267 Lecture 12

49

SUMMA – $n \times n$ matmul on $P^{1/2} \times P^{1/2}$ grid




- $C[i, j]$ is $n/P^{1/2} \times n/P^{1/2}$ submatrix of C on processor P_{ij}
- $A[i,k]$ is $n/P^{1/2} \times b$ submatrix of A
- $B[k,j]$ is $b \times n/P^{1/2}$ submatrix of B
- $C[i,j] = C[i,j] + \sum_k A[i,k]*B[k,j]$
 - summation over submatrices
- Need not be square processor grid

02/26/2015

CS267 Lecture 12

50

SUMMA– $n \times n$ matmul on $P^{1/2} \times P^{1/2}$ grid



For $k=0$ to $n/b-1$
 for all $i = 1$ to $P^{1/2}$
 owner of $A[i,k]$ broadcasts it to whole processor row (using binary tree)
 for all $j = 1$ to $P^{1/2}$
 owner of $B[k,j]$ broadcasts it to whole processor column (using bin. tree)
 Receive $A[i,k]$ into $Acol$
 Receive $B[k,j]$ into $Brow$
 $C_{myproc} = C_{myproc} + Acol * Brow$

02/26/2015

CS267 Lecture 12

51

SUMMA Costs

For $k=0$ to $n/b-1$
 for all $i = 1$ to $P^{1/2}$
 owner of $A[i,k]$ broadcasts it to whole processor row (using binary tree)
 ... #words = $\log P^{1/2} * b * n/P^{1/2}$, #messages = $\log P^{1/2}$
 for all $j = 1$ to $P^{1/2}$
 owner of $B[k,j]$ broadcasts it to whole processor column (using bin. tree)
 ... same #words and #messages
 Receive $A[i,k]$ into $Acol$
 Receive $B[k,j]$ into $Brow$
 $C_{myproc} = C_{myproc} + Acol * Brow$... #flops = $2n^2*b/P$

- Total #words = $\log P * n^2 / P^{1/2}$
 - Within factor of $\log P$ of lower bound
 - (more complicated implementation removes $\log P$ factor)
- Total #messages = $\log P * n/b$
 - Choose b close to maximum, $n/P^{1/2}$, to approach lower bound $P^{1/2}$
- Total #flops = $2n^3/P$

52

Performance of PBLAS

PDGEMM = PBLAS routine for matrix multiply

Observations:
 For fixed N, as P increases Mflops increases, but less than 100% efficiency
 For fixed P, as N increases, Mflops (efficiency) rises

Machine	Procs	Block Size	Speed in Mflops of PDGEMM		
			2000	4000	10000
Cray T3E	4=2x2	32	1055	1070	0
	16=4x4		3630	4005	4292
	64=8x8		13456	14287	16755
IBM SP2	4	50	755	0	0
	16		2514	2850	0
	64		6205	8709	10774
Intel XP/S MP Paragon	4	32	330	0	0
	16		1233	1281	0
	64		4496	4864	5257
Berkeley NOW	4	32	463	470	0
	32=4x8		2490	2822	3450
	64		4130	5457	6647

DGEMM = BLAS routine for matrix multiply

Maximum speed for PDGEMM = # Procs * speed of DGEMM

Observations (same as above):
 Efficiency always at least 48%
 For fixed N, as P increases, efficiency drops
 For fixed P, as N increases, efficiency increases

02/26/2015

Machine	Peak/proc	DGEMM Mflops	Procs	Efficiency = MFlops(PDGEMM)/(Procs*MFlops(DGEMM))		
				2000	4000	10000
Cray T3E	600	360	4	.73	.74	
			16	.63	.70	.75
			64	.58	.62	.73
IBM SP2	266	200	4	.94		
			16	.79	.89	
			64	.48	.68	.84
Intel XP/S MP Paragon	100	90	4	.92		
			16	.86	.89	
			64	.78	.84	.91
Berkeley NOW	334	129	4	.90	.91	
			32	.60	.68	.84
			64	.50	.66	.81

53

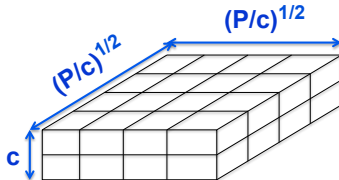
Can we do better?

- Lower bound assumed 1 copy of data: $M = O(n^2/P)$ per proc.
- What if matrix small enough to fit $c > 1$ copies, so $M = cn^2/P$?
 - #words_moved = $\Omega(\text{\#flops} / M^{1/2}) = \Omega(n^2 / (c^{1/2} P^{1/2}))$
 - #messages = $\Omega(\text{\#flops} / M^{3/2}) = \Omega(P^{1/2} / c^{3/2})$
- Can we attain new lower bound?
 - Special case: "3D Matmul": $c = P^{1/3}$
 - Bernsten 89, Agarwal, Chandra, Snir 90, Aggarwal 95
 - Processors arranged in $P^{1/3} \times P^{1/3} \times P^{1/3}$ grid
 - Processor (i,j,k) performs $C(i,j) = C(i,j) + A(i,k) * B(k,j)$, where each submatrix is $n/P^{1/3} \times n/P^{1/3}$
 - Not always that much memory available...

02/26/2015
CS267 Lecture 12
54

2.5D Matrix Multiplication

- Assume can fit cn^2/P data per processor, $c > 1$
- Processors form $(P/c)^{1/2} \times (P/c)^{1/2} \times c$ grid

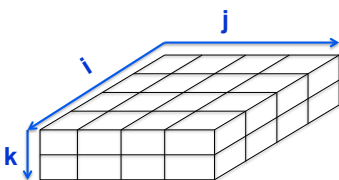


Example: P = 32, c = 2

02/26/2015 CS267 Lecture 12

2.5D Matrix Multiplication

- Assume can fit cn^2/P data per processor, $c > 1$
- Processors form $(P/c)^{1/2} \times (P/c)^{1/2} \times c$ grid




Initially P(i,j,0) owns A(i,j) and B(i,j) each of size $n(c/P)^{1/2} \times n(c/P)^{1/2}$


- P(i,j,0) broadcasts A(i,j) and B(i,j) to P(i,j,k)
- Processors at level k perform 1/c-th of SUMMA, i.e. 1/c-th of $\sum_m A(i,m) * B(m,j)$
- Sum-reduce partial sums $\sum_m A(i,m) * B(m,j)$ along k-axis so P(i,j,0) owns C(i,j)

2.5D Matmul on IBM BG/P, n=64K

- As P increases, available memory grows → c increases proportionally to P
 - #flops, #words_moved, #messages per proc all decrease proportionally to P
 - #words_moved = $\Omega(\text{\#flops} / M^{1/2}) = \Omega(n^2 / (c^{1/2} P^{1/2}))$
 - #messages = $\Omega(\text{\#flops} / M^{3/2}) = \Omega(P^{1/2} / c^{3/2})$
- Perfect strong scaling! But only up to $c = P^{1/3}$



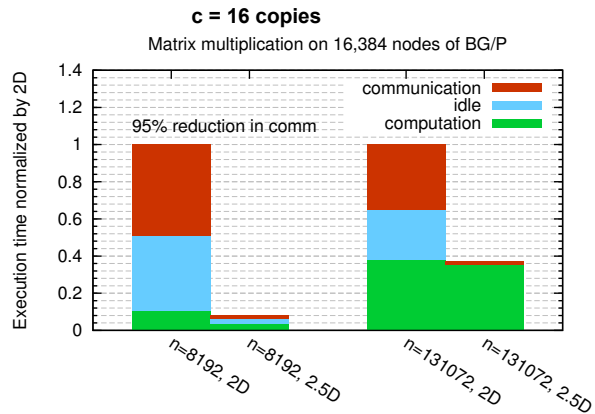
2.5D Matmul on IBM BG/P, 16K nodes / 64K cores



02/26/2015

CS267 Lecture 12

2.5D Matmul on IBM BG/P, 16K nodes / 64K cores



02/26/2015

Distinguished Paper Award, EuroPar'11
SC'11 paper by Solomonik, Bhatle, D.

Perfect Strong Scaling – in Time and Energy

- Every time you add a processor, you should use its memory M too
- Start with minimal number of procs: $PM = 3n^2$
- Increase P by a factor of c → total memory increases by a factor of c
- Notation for timing model:
 - $\gamma_T, \beta_T, \alpha_T$ = secs per flop, per word_moved, per message of size m
- $T(cP) = n^3/(cP) [\gamma_T + \beta_T/M^{1/2} + \alpha_T/(mM^{1/2})]$
= $T(P)/c$
- Notation for energy model:
 - $\gamma_E, \beta_E, \alpha_E$ = joules for same operations
 - δ_E = joules per word of memory used per sec
 - ϵ_E = joules per sec for leakage, etc.
- $E(cP) = cP \{ n^3/(cP) [\gamma_E + \beta_E/M^{1/2} + \alpha_E/(mM^{1/2})] + \delta_E MT(cP) + \epsilon_E T(cP) \}$
= $E(P)$
- c cannot increase forever: $c \leq P^{1/3}$ (3D algorithm)
 - Corresponds to lower bound on #messages hitting 1
- Perfect scaling extends to Strassen's matmul, direct N-body, ...
 - "Perfect Strong Scaling Using No Additional Energy"
 - "Strong Scaling of Matmul and Memory-Indep. Comm. Lower Bounds"
 - Both at bebop.cs.berkeley.edu

Classical Matmul

- Complexity of **classical Matmul**
- Flops: $O(n^3/p)$
- Communication lower bound on #words:

$$\Omega((n^3/p)/M^{1/2}) = \Omega(M(n/M^{1/2})^3/p)$$
- Communication lower bound on #messages:

$$\Omega((n^3/p)/M^{3/2}) = \Omega((n/M^{1/2})^3/p)$$
- All attainable as M increases past $O(n^2/p)$, up to a limit:
 - can increase M by factor up to $p^{1/3}$
 - #words as low as $\Omega(n/p^{2/3})$


02/27/2014

CS267 Lecture 12

61

Strong scaling of Matmul on Hopper (n=94080)

G. Ballard, D., O. Holtz, B. Lipshitz, O. Schwartz



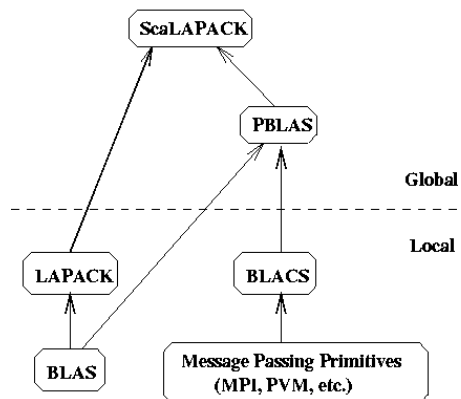
02/26/2015

“Communication-Avoiding Parallel Strassen”
bebop.cs.berkeley.edu, Supercomputing’12

62

ScaLAPACK Parallel Library

ScaLAPACK SOFTWARE HIERARCHY



02/26/2015

CS267 Lecture 12

63

Extensions of Lower Bound and Optimal Algorithms

- For each processor that does G flops with fast memory of size M
 - #words_moved = $\Omega(G/M^{1/2})$
- Extension: for any program that “smells like”
 - Nested loops ...
 - That access arrays ...
 - Where array subscripts are linear functions of loop indices
 - Ex: $A(i,j)$, $B(3*i-4*k+5*j, i-j, 2*k, \dots)$, ...
 - There is a constant s such that
 - #words_moved = $\Omega(G/M^{s-1})$
 - s comes from recent generalization of Loomis-Whitney ($s=3/2$)
 - Ex: linear algebra, n-body, database join, ...
 - Lots of open questions: deriving s, optimal algorithms ...

02/26/2015

CS267 Lecture 12

64