

Adaptive Linear Solvers and Eigensolvers

Jack Dongarra

University of Tennessee Oak Ridge National Laboratory University of Manchester



1



• Common Operations

$$Ax = b; \quad \min_{x} ||Ax - b||; \quad Ax = \lambda x$$

- A major source of large dense linear systems is problems involving the solution of boundary integral equations.
 - The price one pays for replacing three dimensions with two is that what started as a sparse problem in $O(n^3)$ variables is replaced by a dense problem in $O(n^2)$.
- Dense systems of linear equations are found in numerous other applications, including:
 - airplane wing design;
 - radar cross-section studies;
 - flow around ships and other off-shore constructions;
 - diffusion of solid bodies in a liquid;
 - noise reduction; and
- 8/2/19 diffusion of light through small particles 2



Existing Math Software - Dense LA

DIRECT SOLVERS	License	Support	1	Гуре	I	Languag	ge		Mode		Dense	S	parse D	irect	Spa Itera	arse ative	Sp Eiger	arse ivalue	Last release date
			Real	Complex	F77/ F95	С	C++	Shared	Accel.	Dist		SPD	SI	Gen	SPD	Gen	Sym	Gen	
Chameleon	CeCILL-C	yes	X	X		X		X	С	Μ	X								2018-09-15
DPLASMA	BSD	yes	X	X		X		X	С	М	X								2014-04-14
Eigen	MPL2	yes	X	X			X	X			X	X		X	X	Х			2018-07-23
Elemental	New BSD	yes	X	X			X			М	X	X	X	X					2017-02-06
ELPA	LGPL	yes	X	X	F90	X		X		М	X								2018-06-01
FLENS	<u>BSD</u>	yes	X	X			X	X			X								2014-05-11
LAPACK	<u>BSD</u>	yes	X	X	X	X		X			X								2017-11-12
LAPACK95	<u>BSD</u>	yes	X	X	X			X			X								2000-11-30
libflame	New BSD	yes	X	X	Х	X		X			X								2014-03-18
MAGMA	<u>BSD</u>	yes	X	X	Х	X		X	C/O/X		X				X	Х	X		2018-06-25
NAPACK	BSD	yes	X		X			X			X				X		X		?
PLAPACK	LGPL	yes	X	X	X	X				М	X								2007-06-12
PLASMA	<u>BSD</u>	yes	X	X	X	X		X			X								2018-09-04
ScaLAPACK	BSD	yes	X	X	X	X				M/P	X								2018-08-20
Trilinos/Pliris	<u>BSD</u>	yes	X	X		X	X			М	X								2015-05-07
ViennaCL	MIT	yes	X				X	X	C/O/X		X				X	х	X	X	2016-01-20

http://www.netlib.org/utk/people/JackDongarra/la-sw.html

LINPACK, EISPACK, LAPACK, ScaLAPACK
 > PLASMA, MAGMA
 ³

8/2/19



- We are interested in developing Dense Linear Algebra Solvers
- Retool LAPACK and ScaLAPACK for multicore and hybrid architectures



50 Years Evolving SW and Alg

J. H. Wilkinson C. Reinsch Linear Algebra

Tracking Hardware Developments Software/Algorithms follow hardware evolution in time

EISPACK (1970's) Rely on (Translation of Algol to F66) - Fortran, but row oriented LINPACK (1980's) Rely on - Level-1 BLAS operations (Vector operations) - Column oriented LAPACK (1990's) Rely on (Blocking, cache friendly) - Level-3 BLAS operations ScaLAPACK (2000's) Rely on # (Distributed Memory) - PBLAS Mess Passing PLASMA (2010's) Rely on - DAG/scheduler New Algorithms (many-core friendly) - block data layout - some extra kernels SLATE (2020's) Rely on C++ - Tasking DAG scheduling - Tiling, but tiles can come from anywhere - Batched Dispatch

What do you mean by performance?

- What is a xflop/s?
 - > xflop/s is a rate of execution, some number of floating point operations per second.
 - > Whenever this term is used it will refer to 64 bit floating point operations and the operations will be either addition or multiplication.

> Tflop/s refers to trillions (10¹²) of floating point operations per second and

> Pflop/s refers to 10¹⁵ floating point operations per second.

What is the theoretical peak performance?

- > The theoretical peak is based not on an actual performance from a benchmark run, but on a paper computation to determine the theoretical peak rate of execution of floating point operations for the machine.
- > The theoretical peak performance is determined by counting the number of floating-point additions and multiplications (in full precision) that can be completed during a period of time, usually the cycle time of the machine.
- For example, an Intel Skylake processor at 2.1 GHz can complete 32 floating point operations per cycle per core or a theoretical peak performance of 67.2 GFlop/s per core or 1.61 Tflop/s for the socket of 24 cores.



Peak Performance - Per Core





Commodity Processors ...

Over provisioned for floating point operations

Today it's all about data movement



Each Core: 32 Flops per core / cycle

With 2.6 GHz Each Core Peak DP 83.2 Gflop/s Each Socket Peak 665.6 Gflop/s

Memory Access Latencies in Clock Cycles 167 cycles to move a word from memory to a register In 167 cycles single core: 5344 DP Flops, socket: >40K Flops Main memory 167 L3 Cache Full Random access 38 L3 Cache In Page Random access 18 L3 Cache sequential access 14 L2 Cache Full Random access 11 L2 Cache In Page Random access 11 L2 Cache seguential access 11 L1 Cache In Full Random access Need Cache Friendly Algorithms L1 Cache In Page Random access Matrix Multiply and Data Reuse L1 Cache sequential access 0 50 100 150 200

Memory transfer

• One level of memory model on my laptop:



The model IS simplified (see next slide) but it provides an upper bound on performance as well. I.e., we will never go faster than what the model predicts. (And, of course, we can go slower ...)

8/2/19

FMA: fused multiply-add



Note: It is reasonable to expect the one loop codes shown here to perform as well as their Level 1 BLAS counterpart (on multicore with an OpenMP pragma for example).

The true gain these days with using the BLAS is (1) Level 3 BLAS, and (2) portability.

• Take two double precision vectors x and y of size n=375,000.



- Data size:
 - (375,000 double) * (8 Bytes / double) = 3 MBytes per vector

(Two vectors fit in cache (6 MBytes). OK.)

- Time to move the vectors from memory to cache:
 (6 MBytes) / (25.6 GBytes/sec) = 0.23 ms
- Time to perform computation of DOT:
 - (2n flops) / (56 Gflop/sec) = 0.013 ms

Vector Operations

total_time \geq max (time_comm , time_comp) = max (0.23ms , 0.01ms) = 0.23ms

Performance = (2 x 375,000 flops)/.23ms = 3.2 Gflop/s

Performance for DOT ≤ 3.2 Gflop/s Peak is 56 Gflop/s

We say that the operation is communication bounded. No reuse of data.

Level 1, 2 and 3 BLAS



• Double precision matrix A and vectors x and y of size n=860.



• Data size:

- (860² + 2*860 double) * (8 Bytes / double) ~ 6 MBytes
 Matrix and two vectors fit in cache (6 MBytes).

- Time to move the data from memory to cache:
 - (6 MBytes) / (25.6 GBytes/sec) = 0.23 ms
- Time to perform computation of GEMV:

- (2n² flops) / (56 Gflop/sec) = 0.026 ms

Matrix - Vector Operations

total_time \geq max (time_comm , time_comp) = max (0.23ms , 0.026ms) = 0.23ms

Performance = (2 x 860² flops)/.23ms = 6.4 Gflop/s

Performance for GEMV ≤ 6.4 Gflop/s

Performance for DOT ≤ 3.2 Gflop/s

Peak is 56 Gflop/s

We say that the operation is communication bounded. Very little reuse of data.

• Take two double precision vectors x and y of size n=500.



• Data size:

- (500² double) * (8 Bytes / double) = 2 MBytes per matrix
(Three matrices fit in cache (6 MBytes). OK.)

- Time to move the matrices in cache:
 - (6 MBytes) / (25.6 GBytes/sec) = 0.23 ms
- Time to perform computation in GEMM:
 - (2n³ flops) / (56 Gflop/sec) = 4.5 ms

Matrix Matrix Operations

```
total_time \geq max ( time_comm , time_comp )
```

= max(0.23ms , 4.46ms) = 4.46ms

For this example, communication time is less than 6% of the computation time.

Performance = $(2 \times 500^{3} \text{ flops})/4.5 \text{ms} = 55.5 \text{ Gflop/s}$

There is a lots of data reuse in a GEMM; 2/3n per data element. Has good temporal locality.

If we assume total_time ≈ time_comm +time_comp, we get Performance for GEMM ≈ 55.5 Gflop/sec

Performance for DOT ≤ 3.2 Gflop/s Performance for GEMV ≤ 6.4 Gflop/s

(Out of 56 Gflop/sec possible, so that would be 99% peak performance efficiency.)





1 core Intel Haswell i7-4850HQ, 2.3 GHz, Memory: DDR3L-1600MHz 6 MB shared L3 cache, and each core has a private 256 KB L2 and 64 KB L1. The theoretical peak per core double precision is 56 Gflop/s per core. Compiled with gcc and using Veclib

Level 1, 2 and 3 BLAS

18 cores Intel Xeon Gold 6140 (Skylake), 2.3 GHz, Peak DP = 1325 Gflop/s



Issues

- Reuse based on matrices that fit into cache.
- What if you have matrices bigger than cache?

Issues

- Reuse based on matrices that fit into cache.
- What if you have matrices bigger than cache?
- Break matrices into blocks or tiles that will fit.



8/2/19

LU Factorization in LINPACK (1970's)



- Factor one column at a time
 - i_amax and _scal
- Update each column of trailing matrix, one column at a time
 - _ахру
- Level 1 BLAS

22

- Bulk synchronous
 - Single main thread
 - Parallel work in BLAS
 - "Fork-and-join" model



The Standard LU Factorization LAPACK 1980's HPC of the Day: Cache Based SMP



- Factor panel of *nb* columns
 - getf2, unblocked BLAS-2 code
- Level 3 BLAS update block-row of U
 - trsm
- Level 3 BLAS update trailing matrix
 - gemm
 - Aimed at machines with cache hierarchy
- Bulk synchronous



Most flops in gemm update

- 2/3 n³ term
- Easily parallelized using multi-threaded BLAS
- Done in any reasonable software
- Other operations lower order
 - · Potentially expensive if not parallelized



Ĉ Last Generations of DLA Software

Software/Algor	Software/Algorithms follow hardware evolution in time										
LINPACK (70's) (Vector operations)		Rely on - Level-1 BLAS operations									
LAPACK (80's) (Blocking, cache friendly)		Rely on - Level-3 BLAS operations									
ScaLAPACK (90's) (Distributed Memory)		Rely on - PBLAS Mess Passing									
	2D Block Cyclic Layout										

					~,				.,		•												
			Ma	ıtrix p	point	of vi	ew							Proc	æ	ssoi	, bo	int of	F١	view			
	_									ī í	٢	-	-			_	-		۱ſ	-	-	_	
	0	2	4	0	2	4	0	2	4			0	0	0		2	2	2		4	4	4	
	1	3	5	1	3	5	1	3	5			0	0	0		2	2	2		4	4	4	
	0	2	4	0	2	4	0	2	4			0	0	0		2	2	2		4	4	4	
	1	3	5	П	3	5	h	3	5			0	0	0		2	2	2		4	4	4	
	<u> </u>	Ľ	<u> </u>	Ŀ	<u> </u>	-	Ŀ	-		.		0	0	0		2	2	2		4	4	4	
	0	2	4	0	2	4	0	2	4		l	_				_				_		_	
	1	3	5	1	3	5	1	3	5			1	1	1		3	3	3		5	5	5	
T										1		÷	-	-		-	0	-		-	-		
	0	2	4	0	2	4	<u> </u>	2	4			1	1	1		3	3	3		5	5	5	
	1	3	5	1	3	5	1	3	5			1	1	1		3	3	3		5	5	5	
Ī	0	2	4	0	2	4	0	2	4	1		1	1	1		3	3	3		5	5	5	
	-			<u> </u>		<u> </u>	Ľ.	_	<u> </u>														_

8/2/19

ScaLAPACK

Scalable Linear Algebra PACKage

- Distributed memory
- Message Passing
 - Clusters of SMPs
 - Supercomputers
- Dense linear algebra
- Modules
 - PBLAS: Parallel BLAS

Parallelism in ScaLAPACK

- Similar to LAPACK
- Bulk-synchronous processing

 separate message passing & compute
- Most flops in gemm update
 - 2/3 n³ term
 - Can use sequential BLAS,
 p x q = # cores
 = # MPI processes,
 - num_threads = 1
 - Or multi-threaded BLAS,



Today's HPC Environment for Numerical Libraries

- Highly parallel
 - Distributed memory
 - MPI + Open-MP programming model
- Heterogeneous
 - Commodity processors + GPU accelerators
- Simple loop level parallelism too limiting in terms of performance
- Communication between parts very expensive compared to floating point ops
- Comparison of operation counts may not reflect time to solution
- Floating point hardware at 64, 32, and 16 bit levels







\land ≠ 🏈

Туре	Size	Range	$u = 2^{-t}$
half	16 bits	10 ^{±5}	$2^{-11}\approx 4.9\times 10^{-4}$
single double	32 bits 64 bits	10 ^{±38} 10 ^{±308}	$\begin{array}{l} 2^{-24}\approx 6.0\times 10^{-8} \\ 2^{-53}\approx 1.1\times 10^{-16} \end{array}$
quadruple	128 bits	10 ^{±4932}	$2^{-113} \approx 9.6 \times 10^{-35}$



Tile Algorithms: Matrix Decomposition

Track dependencies — Directed acyclic graph (DAG)









Fork-join schedule on 4 cores with artificial synchronizations



Reorder without synchronizations



Critical path

Dataflow Based Design

- Objectives
 - High utilization of each core
 - Scaling to large number of cores
 - Synchronization reducing algorithms
- Methodology
 - Dynamic DAG scheduling using OpenMP
 - Explicit parallelism
 - Implicit communication
 - Fine granularity / block data layout
- Arbitrary DAG with dynamic scheduling



Cholesky; 45% improvement













Total: 18(3t+6)

Assume a t by t matrix tiling then Cholesky Factorization alone: 3t-2 Total: 25(7t-3)

32

Standard for Batched Computations

- Define standard API for batched BLAS and LAPACK in collaboration with Intel/Nvidia/other users
- Fixed size: most of BLAS and LAPACK released
- Variable size: most of BLAS released
- Variable size: LAPACK in the branch
- Native GPU algorithms (Cholesky, LU, QR) in the branch
- Tiled algorithm using batched routines on tile or LAPACK data layout in the branch
- Framework for Deep Neural Network kernels
- CPU, KNL and GPU routines
- FP16 routines in progress





Batched Computations

Non-batched computation

• **loop over the matrices one by one** and compute using multithread (note that, since matrices are of small sizes there is not enough work for all the cores). So we expect low performance as well as threads contention might also affect the performance







Batched Computations

Batched computation

- Distribute all the matrices over the available resources by assigning a matrix to each group of core/TB to operate on it independently
 - For very small matrices, assign a matrix/core (CPU) or per TB for GPU
 - For medium size a matrix go to a team of cores (CPU) or many TB's (GPU)
 - For large size switch to multithreads classical 1 matrix per round.







Batched Computations: How do we design and optimize



Machine Learning in Computational Science

Many fields are beginning to adopt machine learning to augment modeling and simulation methods

- Climate
- Biology
- Drug Design
- Epidemology
- Materials
- Cosmology
- High-Energy Physics







Deep Learning Needs Small Matrix Operations



EEE 754 Half Precision (16-bit) Floating Pt Standard

A lot of interest driven by "machine learning"



	AMD Rad	eon Instinct	
	Instinct MI6	Instinct MI8	Instinct MI25
Memory Type	16GB GDDR5	4GB HBM	"High Bandwidth Cache and Controller"
Memory Bandwidth	224GB/sec	512GB/sec	?
Single Precision	5.7 TELOPS	8 2 TELOPS	12.5 TFLOPS
(FP32)			
Half Precision (FP16)	5.7 TFLOPS	8.2 TFLOPS	25 TFLOPS
TDP	<150W	<175W	<300
Cooling	Passive	Passive (SFF)	Passive
GPU	Polaris 10	Fiji	Vega
Manufacturing Process	GloFo 14nm	TSMC 28nm	?

GPU PERF	ORMAN	CE COMP	ARISON
	P100	V100	Ratio
DL Training FP16	10 TFLOPS	120 TFLOPS	12x
DL Inferencing FP16	21 TFLOPS	120 TFLOPS	6x
FP64/FP32	5/10 TFLOPS	7.5/15 TFLOPS	1.5x
HBM2 Bandwidth	720 GB/s	900 GB/s	1.2x
STREAM Triad Perf	557 GB/s	855 GB/s	1.5x
NVLink Bandwidth	160 GB/s	300 GB/s	1.9x
L2 Cache	4 MB	6 MB	1.5x
L1 Caches	1.3 MB	10 MB	7.7x



• Today many precisions to deal with (IEEE Standard)

Туре	Size	Range	$u = 2^{-t}$
half	16 bits	10 ^{±5}	$2^{-11}\approx 4.9\times 10^{-4}$
single double	32 bits 64 bits	10 ^{±38} 10 ^{±308}	$\begin{array}{l} 2^{-24} \approx 6.0 \times 10^{-8} \\ 2^{-53} \approx 1.1 \times 10^{-16} \end{array}$
quadruple	128 bits	10 ^{±4932}	$2^{-113}\approx9.6\times10^{-35}$







- 64 bit floating point (FMA): 7.5 Tflop/s
- 32 bit floating point (FMA): 15 Tflop/s
- 16 bit floating point (FMA): 30 Tflop/s
- 16 bit floating point with Tensor core: 120 Tflop/s



D = AB + C

VOLTA TENSOR OPERATION



Also supports FP16 accumulator mode for inferencing





Study of the Matrix Matrix multiplication kernel on Nvidia V100



Study of the Matrix Matrix multiplication kernel on Nvidia V100



dgemm achieve about 6.4 Tflop/s sgemm achieve about 14 Tflop/s hgemm achieve about 27 Tflop/s Tensor cores gemm reach about 85 Tflop/s

Matrix matrix multiplication GEMM



Study of the Matrix Matrix multiplication kernel on Nvidia V100



dgemm achieve about 6.4 Tflop/s sgemm achieve about 14 Tflop/s hgemm achieve about 27 Tflop/s Tensor cores gemm reach about 85 Tflop/s

Matrix matrix multiplication GEMM



Study of the rank k update used by the LU factorization algorithm on Nvidia V100



 In LU factorization need matrix multiple but operations is a rank-k update computing the Schur complement

	=		+	×	
--	---	--	---	---	--



Idea: use low precision to compute the expensive flops (LU $O(n^3)$) and then iteratively refine the solution in order to achieve the FP64 arithmetic

Iterative refinement for dense systems, Ax = b, can work this way. L U = lu(A) x = U\(L\b) r = b - Ax	lower precisio lower precisio FP64 precisio	0(n ³) 0(n ²) 0(n ²)
<pre>WHILE r not small enough 1. find a correction "z" to adjust x that satisfy Az=r solving Az=r could be done by either:</pre>	efinement using GMRes FP64 precision FP64 precision FP64 precision	O(n ²) O(n ²) O(n ¹) O(n ²)
END Higham and Carson showed can solve the inner problem with iterative method and no	infect the solution.	
Wilkinson, Moler, Stewart, & Higham provide error bound for SP fl pt results when using DP fl pt.	E. Carson & N. Higham, "Accelera Linear Systems by Iterative Refine Precisions SIAM J. Sci. Comput., 4	ing the Solution of ment in Three 0(2), A817–A847.

> Need the original matrix to compute residual (r) and matrix cannot be too badly conditioned

Improving Solution

- *z* is the correction or $(x_{i+1} x_i)$
- Computed in lower precision and then added to the approximate solution in higher precision $x_i + z$



$$x_{i+1} = x_i - \frac{f(x_i)}{f'(x_i)}$$
$$\underbrace{x_{i+1} - x_i}_{f'(x_i)} = -\frac{f(x_i)}{f'(x_i)}$$











$Flops = 2n^3/(3 time)$ meaning twice higher is twice faster

- solving Ax = b using FP64 LU
- solving Ax = b using FP32 LU and iterative refinement to achieve FP64 accuracy
- solving Ax = b using FP16 LU and iterative refinement to achieve FP64 accuracy
- solving Ax = b using FP16 Tensor Cores LU and iterative refinement to achieve FP64 accuracy



$Flops = 2n^3/(3 time)$ meaning twice higher is twice faster

- solving Ax = b using FP64 LU
- solving Ax = b using FP32 LU and iterative refinement to achieve FP64 accuracy
- solving Ax = b using FP16 LU and iterative refinement to achieve FP64 accuracy
- solving Ax = b using FP16 Tensor Cores LU and iterative refinement to achieve FP64 accuracy

Problem generated with an arithmetic distribution of the singular values $\sigma_i = 1 - (\frac{i-1}{n-1})(1 - \frac{1}{cond})$ and positive eigenvalues.

Critical Issues at Exascale for Algorithm and Software Design

- Synchronization-reducing algorithms
 - Break Fork-Join model
- Communication-reducing algorithms
 - Use methods which have lower bound on communication
- Mixed precision methods (half (16bit), single(32 bit), & double precision (64))
 - 2x 10x speed of ops and 2x 4x speed for data movement
- Autotuning Performance Debugging
 - Today's machines are very complicated, build "smarts" into software to adapt to the hardware
- Fault resilient algorithms
 - Implement algorithms that can recover from failures/bit flips
- Reproducibility of results
 - Today we can't guarantee this. We understand the issues, but some of our "colleagues" have a hard time with this.



Collaborators / Software / Support

- PLASMA <u>http://icl.cs.utk.edu/plasma/</u>
- MAGMA <u>http://icl.cs.utk.edu/magma/</u>
- SLATE
 - https://icl.utk.edu/slate/
 - https://bitbucket.org/icl/slate/src/default/
- PaRSEC (Parallel Runtime Scheduling & Execution Control)
- http://icl.cs.utk.edu/parsec/

> Collaborating partners University of Tennessee, Knoxville University of California, Berkeley University of Colorado, Denver

EXASCALE COMPUTING PROJEC

Looking for Grad Students and Post-Docs