PARLab Parallel Boot Camp



Computational Patterns and Autotuning

Jim Demmel EECS and Mathematics University of California, Berkeley

Outline



- Productive parallel computing depends on recognizing and exploiting useful patterns
 - Computational (7 Motifs) and Structural
- Simplest case: use "best" existing highly tuned implementation
 - Best: Fastest? Most accurate? Fewest keystrokes?
- Optimizing (some of) the 7 Motifs
 - To minimize time or energy, minimize communication (moving data)
 - Between levels of the memory hierarchy
 - Between processors over a network
 - Autotuning to explore large design spaces
 - Too hard (tedious) to write by hand, let machine do it
- SEJITS how to deliver autotuning to more programmers
- For more details, see
 - CS267: www.cs.berkeley.edu/~demmel/cs267_Spr13
 - 10-hour short course: issnla2010.ba.cnr.it/Courses.htm
 - Papers at bebop.cs.berkeley.edu, parlab.eecs.berkeley.edu

"7 Motifs" of High Performance Computing



- Phil Colella (LBL) identified 7 kernels of which most simulation and data-analysis programs are composed:
 - 1. Dense Linear Algebra
 - Ex: Solve Ax=b or $Ax = \lambda x$ where A is a dense matrix
 - 2. Sparse Linear Algebra
 - Ex: Solve Ax=b or Ax = λx where A is a sparse matrix (mostly zero)
 - 3. Operations on Structured Grids
 - Ex: $A_{new}(i,j) = 4*A(i,j) A(i-1,j) A(i+1,j) A(i,j-1) A(i,j+1)$
 - 4. Operations on Unstructured Grids
 - Ex: Similar, but list of neighbors varies from entry to entry
 - 5. Spectral Methods
 - Ex: Fast Fourier Transform (FFT)
 - 6. Particle Methods
 - Ex: Compute electrostatic forces on n particles
 - 7. Monte Carlo
 - Ex: Many independent simulations using different inputs

"7 Motifs" of High Performance Computing



- Phil Colella (LBL) identified 7 kernels of which most simulation and data-analysis programs are composed:
 - 1. Dense Linear Algebra
 - Ex: Solve Ax=b or Ax = λx where A is a dense matrix
 - 2. Sparse Linear Algebra
 - Ex: Solve Ax=b or Ax = λx where A is a sparse matrix (mostly zero)
 - 3. Operations on Structured Grids
 - Ex: $A_{new}(i,j) = 4*A(i,j) A(i-1,j) A(i+1,j) A(i,j-1) A(i,j+1)$
 - 4. Operations on Unstructured Grids
 - Ex: Similar, but list of neighbors varies from entry to entry
 - 5. Spectral Methods
 - Ex: Fast Fourier Transform (FFT)
 - 6. Particle Methods
 - Ex: Compute electrostatic forces on n particles
 - 7. Monte Carlo
 - Ex: Many independent simulations using different inputs

Organizing Linear Algebra Motifs in books and on-line



www.netlib.org/lapack



www.netlib.org/templates

gams.nist.gov

ScaLAPACK Users' Guide

L. S. Blackford and Chick A. Clearly - F. D'Arezedic J. Denmed - J. Difficial - I. Dingama - S. Parimaning G. Henry - A Petter - K Stanley - D. Walker - R. C. Whaley



www.netlib.org/scalapack



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

Why Minimize Communication? (1/2

- Running time of an algorithm is sum of 3 terms:
 - # flops * time_per_flop
 - # words moved / bandwidth
 - # messages * latency

communication

- Time_per_flop << 1/ bandwidth << latency
 - Gaps growing exponentially with time [FOSC]

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

• Minimize communication to save time

Why Minimize Communication? (2/2)4





President Obama cites Communication-Avoiding Algorithms in the FY 2012 Department of Energy Budget Request to Congress:

"New Algorithm Improves Performance and Accuracy on Extreme-Scale Computing Systems. On modern computer architectures, communication between processors takes longer than the performance of a floating point arithmetic operation by a given **processor.** ASCR researchers have developed a new method, derived from commonly used linear algebra methods, to **minimize** communications between processors and the memory hierarchy, by reformulating the communication patterns specified within the algorithm. This method has been implemented in the TRILINOS framework, a highly-regarded suite of software, which provides functionality for researchers around the world to solve large scale, complex multi-physics problems."

FY 2010 Congressional Budget, Volume 4, FY2010 Accomplishments, Advanced Scientific

CA-GMRES (Hoemmen, Mohiyuddin, Yelick, JD) "Tall-Skinny" QR (Grigori, Hoemmen, Langou, JD) Obstacle to avoiding communication: Low "computational intensity"



- Let f = #arithmetic operations in an algorithm
- Let m = #words of data needed
- Def: q = f/m = computational intensity
- If q small, say q=1, so one op/word, then algorithm can't run faster than memory speed
- But if q large, so many ops/word, then algorithm can (potentially) fetch data, do many ops while in fast memory, only limited by (faster!) speed of arithmetic
- We seek algorithms with high q
 - Still need to be clever to take advantage of high q



DENSE LINEAR ALGEBRA MOTIF

Brief history of (Dense) Linear Algebra software (1/6)

- 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 on vectors
 - Ex: $y = \alpha \cdot x + y$ ("AXPY"), dot product, etc
 - Goals
 - Common pattern to ease programming, efficiency, robustness
 - Used in libraries like LINPACK (for linear systems)
 - Source of the name "LINPACK Benchmark" (not the code!)
 - Why BLAS 1 ? 1 loop, do O(n¹) ops on O(n¹) data
 - Computational intensity = q = 2n/3n = 2/3 for AXPY
 - Very low!
 - BLAS1, and so LINPACK, limited by memory speed
 - Need something faster ...

Brief history of (Dense) Linear Algebra software (2/6)

- So the BLAS-2 were invented (1984-1986)
 - Standard library of 25 operations (mostly) on matrix/vector pairs
 - Ex: $y = \alpha \cdot A \cdot x + \beta \cdot y$ ("GEMV"), $A = A + \alpha \cdot x \cdot y^{T}$ ("GER"), $y = T^{-1} \cdot x$ ("TRSV")
 - Why BLAS 2 ? 2 nested loops, do O(n²) ops on O(n²) data
 - But q = computational intensity still just ~ (2n²)/(n²) = 2
 - Was OK for vector machines, but not for machine with caches, since q still just a small constant

Brief history of (Dense) Linear Algebra software (3/6)

 \mathcal{M}

- The next step: BLAS-3 (1987-1988)
 - Standard library of 9 operations (mostly) on matrix/matrix pairs
 - Ex: $C = \alpha \cdot A \cdot B + \beta \cdot C$ ("GEMM"), $C = \alpha \cdot A \cdot A^T + \beta \cdot C$ ("SYRK"), $C = T^{-1} \cdot B$ ("TRSM")
 - Why BLAS 3 ? 3 nested loops, do O(n³) ops on O(n²) data
 - So computational intensity $q=(2n^3)/(4n^2) = n/2 big$ at last!
 - Tuning opportunities machines with caches, other mem. hierarchy levels
- How much faster can BLAS 3 go?

Matrix-multiply, optimized several ways



Speed of n-by-n matrix multiply on Sun Ultra-1/170, peak = 330 MFlops

Faster Matmul C=A*B by "Blocking"

• Replace usual 3 nested loops ...

for i=1 to n for j=1 to n for k=1 to n $C(i,j) = C(i,j) + A(i,k)^*B(k,j)$



Lower bounds on Communication for Matmu

- Assume sequential n³ algorithm for C=A*B
 - i.e. not Strassen-like
- Assume A, B and C fit in slow memory, but not in fast memory of size M
- Thm: Lower bound on #words_moved to/from slow memory, no matter the order n³ operations are done,
 = Ω (n³ / M^{1/2}) [Hong & Kung (1981)]
- Attained by "blocked" algorithm
 - Some other algorithms attain it too
 - Widely implemented in libraries (eg Intel MKL)

How hard is hand-tuning, anyway?



- Results of 22 student teams trying to tune matrix-multiply, in CS267 Spr09
- Students given "blocked" code to start with
 - Still hard to get close to vendor tuned performance (ACML)
- For more discussion, see <u>www.cs.berkeley.edu/~volkov/cs267.sp09/hw1/results/</u>
- Naïve matmul: just 2% of peak

How hard is hand-tuning, anyway



What part of the Matmul Search Space Looks Like



A 2-D slice of a 3-D register-tile search space. The dark blue region was pruned. (Platform: Sun Ultra-IIi, 333 MHz, 667 Mflop/s peak, Sun cc v5.0 compiler)

Autotuning DGEMM with ATLAS (n = 500)

Source: Jack Dongarra

MFLOPS



- ATLAS is faster than all other portable BLAS implementations and it is comparable with machine-specific libraries provided by the vendor.
- ATLAS written by C. Whaley, inspired by PHiPAC, by Asanovic, Bilmes, Chin, D.

Brief history of (Dense) Linear Algebra software (4/6)



- LAPACK "Linear Algebra PACKage" uses BLAS-3 (1989 now)
 - Ex: Obvious way to express Gaussian Elimination (GE) is adding multiples of each row to other rows – BLAS-1
 - Need to reorganize GE (and everything else) to use BLAS-3 instead
 - Contents of current LAPACK (summary)
 - Algorithms we can turn into (nearly) 100% BLAS 3 for large n
 - Linear Systems: solve Ax=b for x
 - Least Squares: choose x to minimize $\sqrt{\Sigma_i r_i^2}$ where r=Ax-b
 - Algorithms that are only up to ~50% BLAS 3, rest BLAS 1 & 2
 - "Eigenproblems": Find λ and x where Ax = λ x
 - Singular Value Decomposition (SVD): $A^TAx = \sigma^2 x$
 - Error bounds for everything
 - Lots of variants depending on A's structure (banded, A=A^T, etc)
 - Widely used (list later)
 - All at <u>www.netlib.org/lapack</u>

Brief history of (Dense) Linear Algebra software (5/6)

- 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 subset of LAPACK's functionality available
 - Work in progress (contributions welcome!)
 - All at <u>www.netlib.org/scalapack</u>

Success Stories for Sca/LAPACK

- Widely used
 - Adopted by Mathworks, Cray, Fujitsu, HP, IBM, IMSL, Intel, NAG, NEC, SGI, ...
 - >172M web hits(in 2013, 157M in 2012)
 @ Netlib (incl. CLAPACK, LAPACK95)
- New science discovered through the solution of dense matrix systems
 - Nature article on the flat universe used
 ScaLAPACK
 - 1998 Gordon Bell Prize

<u>www.nersc.gov/news/reports/newNERS</u> <u>Cresults050703.pdf</u> Cosmic Microwave Background Analysis, BOOMERanG collaboration, MADCAP code (Apr. 27, 2000).

• Currently funded to improve, update, maintain Sca/LAPACK







Lower bound for all "n³-like" linear algebra

• Let M = "fast" memory size (per processor)

#words_moved (per processor) = Ω (#flops (per processor) / M^{1/2})

- Parallel case: assume either load or memory balanced
 - Holds for
 - Matmul

Lower bound for all "n³-like" linear algebra

• Let M = "fast" memory size (per processor)

#words_moved (per processor) = Ω (#flops (per processor) / M^{1/2})

#messages_sent ≥ #words_moved / largest_message_size

- Parallel case: assume either load or memory balanced
 - Holds for
 - Matmul, BLAS, LU, QR, eig, SVD, tensor contractions, ...
 - Some whole programs (sequences of these operations, no matter how individual ops are interleaved, eg A^k)
 - Dense and sparse matrices (where #flops $<< n^3$)
 - Sequential and parallel algorithms
 - Some graph-theoretic algorithms (eg Floyd-Warshall)

Lower bound for all "n³-like" linear algebra

• Let M = "fast" memory size (per processor)

#words_moved (per processor) = Ω (#flops (per processor) / M^{1/2})

#messages_sent (per processor) = Ω (#flops (per processor) / M^{3/2})

- Parallel case: assume either load or memory balanced
 - Holds for
 SIAM SIAG/LA Best Paper 2012
 - Matmul, BLAS, LU, QR, eig, SVD, tensor contractions, ...
 - Some whole programs (sequences of these operations, no matter how individual ops are interleaved, eg A^k)
 - Dense and sparse matrices (where #flops << n³)
 - Sequential and parallel algorithms
 - Some graph-theoretic algorithms (eg Floyd-Warshall)

Can we attain these lower bounds?

- Do conventional dense algorithms as implemented in LAPACK and ScaLAPACK attain these bounds?
 Mostly not
- If not, are there other algorithms that do?
 - Yes, for much of dense linear algebra
 - New algorithms, with new numerical properties, new ways to encode answers, new data structures
 - Not just loop transformations (need those too!)
- Only a few sparse algorithms so far
- Lots of work in progress





2.5D Matrix Multiply Timing Breakdown

c = 16 copies

Matrix multiplication on 16,384 nodes of BG/P



TSQR: QR of a Tall, Skinny matrix

$$W = \begin{pmatrix} W_0 \\ W_1 \\ W_2 \\ W_3 \end{pmatrix}$$
$$\begin{pmatrix} R_{00} \\ R_{10} \\ R_{20} \\ R_{30} \end{pmatrix} = \begin{pmatrix} Q_{01} R_{01} \\ Q_{11} R_{11} \end{pmatrix}$$
$$\begin{pmatrix} R_{01} \\ R_{01} \\ R_{11} \end{pmatrix} = (Q_{02} R_{02})$$

TSQR: QR of a Tall, Skinny matrix





$$\begin{pmatrix} \mathsf{R}_{00} \\ \mathsf{R}_{10} \\ \mathsf{R}_{20} \\ \mathsf{R}_{30} \end{pmatrix} = \begin{pmatrix} \mathsf{Q}_{01} & \mathsf{R}_{01} \\ \mathsf{Q}_{11} & \mathsf{R}_{11} \end{pmatrix} = \begin{pmatrix} \mathsf{Q}_{01} \\ \mathsf{Q}_{11} \end{pmatrix} \cdot \begin{pmatrix} \mathsf{R}_{01} \\ \mathsf{R}_{11} \end{pmatrix}$$
$$\begin{pmatrix} \frac{\mathsf{R}_{01}}{\mathsf{R}_{11}} \\ \mathsf{R}_{11} \end{pmatrix} = \begin{pmatrix} \mathsf{Q}_{02} & \mathsf{R}_{02} \end{pmatrix}$$

Output = { $Q_{00}, Q_{10}, Q_{20}, Q_{30}, Q_{01}, Q_{11}, Q_{02}, R_{02}$ }

TSQR: An Architecture-Dependent Algorithm

Parallel:
$$W = \begin{bmatrix} W_0 \\ W_1 \\ W_2 \\ W_3 \end{bmatrix} \xrightarrow{\rightarrow} R_{10} \xrightarrow{\rightarrow} R_{01} \xrightarrow{\rightarrow} R_{02}$$

$$\xrightarrow{\rightarrow} R_{20} \xrightarrow{\rightarrow} R_{11} \xrightarrow{\rightarrow} R_{02}$$

Sequential:
$$W = \begin{bmatrix} W_0 \\ W_1 \\ W_2 \\ W_3 \end{bmatrix} \xrightarrow{\rightarrow} R_{00} \xrightarrow{\rightarrow} R_{01} \xrightarrow{\rightarrow} R_{02} \xrightarrow{\rightarrow} R_{03}$$

Dual Core:
$$W = \begin{bmatrix} W_0 \\ W_1 \\ W_2 \\ W_3 \end{bmatrix} \xrightarrow{\rightarrow} R_{01} \xrightarrow{\rightarrow} R_{01} \xrightarrow{\rightarrow} R_{02} \xrightarrow{\rightarrow} R_{03}$$

Multicore / Multisocket / Multirack / Multisite / Out-of-core: ? Can choose reduction tree dynamically

TSQR Performance Results

Parallel

- Intel Clovertown
 - Up to 8x speedup (8 core, dual socket, 10M x 10)
- Pentium III cluster, Dolphin Interconnect, MPICH
 - Up to 6.7x speedup (16 procs, 100K x 200)
- BlueGene/L
 - Up to **4x** speedup (32 procs, 1M x 50)
- Tesla C 2050 / Fermi
 - Up to **13x** (110,592 x 100)
- Grid 4x on 4 cities vs 1 city (Dongarra, Langou et al)
- Cloud (Gleich, Benson) cost ~ 2 MapReduce operations, orders of magnitude faster
- Sequential
 - · "Infinite speedup" for out-of-Core on PowerPC laptop
 - As little as 2x slowdown vs (predicted) infinite DRAM
 - LAPACK with virtual memory never finished
- Joint work with Grigori, Hoemmen, Langou, Anderson, Ballard, Keutzer, others

Brief history/future of (Dense) Linear Algebra software (6/6)

- Communication-Avoiding for everything (open problems...)
 - Perfect strong scaling for time and energy
 - Extensions to Strassen-like algorithms
- Extensions for multicore
 - PLASMA Parallel Linear Algebra for Scalable Multicore Architectures
 - Dynamically schedule tasks into which algorithm is decomposed, to minimize synchronization, keep all processors busy
 - Release 2.5.1 at icl.cs.utk.edu/plasma/
- Extensions for heterogeneous architectures, eg CPU + GPU
 - "Benchmarking GPUs to tune Dense Linear Algebra"
 - Best Student Paper Prize at SC08 (Vasily Volkov)
 - Paper, slides and code at <u>www.cs.berkeley.edu/~volkov</u>
 - Lower, matching upper bounds (tech report at bebop.cs.berkeley.edu)
 - MAGMA Matrix Algebra on GPU and Multicore Architectures
 - Release 1.4 at icl.cs.utk.edu/magma/
- How much code generation can we automate?
 - MAGMA , and FLAME (<u>www.cs.utexas.edu/users/flame/</u>)

Beyond Dense Linear Algebra

- Recall matmul: #words_moved = Ω (#flops/M^{1/2}), attained by "block algorithm" with M^{1/2} x M^{1/2} blocks
 Where do all the 1/2 's come from?
- "Thm": Lower bound extends to any algorithm that
 - "smells like" nested loops over i1, i2, i3, ... ,ik
 - Accesses arrays like A(i1,i2+3*i3-2*i4,i3+4*i5,...), B(pntr(i2),...)
 - Lower bound becomes Ω (#loop_iteration/M^S)
 - Can write down optimal algorithms in many cases of interest
 - Linear Algebra, N-body, data-base-join, ...
 - Open problem: Does an algorithm attaining lower bound always exist?


SPARSE LINEAR ALGEBRA MOTIF

Sparse Matrix Computations

- Similar problems to dense matrices
 - Ax=b, Least squares, Ax = λx , SVD, ...
- But different algorithms!
 - Exploit structure: only store, work on nonzeros
 - Direct methods
 - LU, Cholesky for Ax=b, QR for Least squares
 - See <u>crd.lbl.gov/~xiaoye/SuperLU/SparseDirectSurvey.pdf</u> for a survey of available serial and parallel sparse solvers
 - See <u>crd.lbl.gov/~xiaoye/SuperLU/index.html</u> for LU codes
 - Iterative methods for Ax=b, least squares, eig, SVD
 - Use simplest operation: Sparse-Matrix-Vector-Multiply (SpMV)
 - Krylov Subspace Methods: find "best" solution in space spanned by vectors generated by SpMVs



- See <u>www.netlib.org/templates</u> for Ax=b
- See <u>www.cs.ucdavis.edu/~bai/ET/contents.html</u> for Ax=λx and SVD

Sparse Outline

- Approaches to Automatic Performance Tuning
- Results for sparse matrix kernels
 - Sparse Matrix Vector Multiplication (SpMV)
 - Sequential and Multicore results
- OSKI = Optimized Sparse Kernel Interface
 pOSKI = parallel OSKI
- Tuning Entire Sparse Solvers
 - Avoiding Communication
- What is a sparse matrix?



Approaches to Automatic Performance Tuning



- Goal: Let machine do hard work of writing fast code
- Why is tuning dense BLAS "easy"?
 - Can do the tuning off-line: once per architecture, algorithm
 - Can take as much time as necessary (hours, a week...)
 - At run-time, algorithm choice may depend only on few parameters (matrix dimensions)
- Can't always do tuning off-line
 - Algorithm and implementation may strongly depend on data only known at run-time
 - Ex: Sparse matrix nonzero pattern determines both best data structure and implementation of Sparse-matrix-vector-multiplication (SpMV)
 - Part of search for best algorithm must be done (very quickly!) at run-time
- Tuning FFTs and signal processing
 - Seems off-line, but maybe not, because of code size
 - www.spiral.net, www.fftw.org



Source: Accelerator Cavity Design Problem (Ko via Husbands)

Linear Programming Matrix

A Sparse Matrix You Use Every Day





SpMV with Compressed Sparse Row (CSR) Storage



Matrix-vector multiply kernel: $y(i) \leftarrow y(i) + A(i,j) * x(j)$

```
for each row i
for k=ptr[i] to ptr[i+1] do
    y[i] = y[i] + val[k]*x[ind[k]]
```

Only 2 flops per 2 mem_refs: Limited by getting data from memory Example: The Difficulty of Tuning



- n = 21200
- nnz = 1.5 M
- kernel: SpMV

 Source: NASA structural analysis problem Example: The Difficulty of Tuning



Matrix 02-raefsky3

- n = 21200
- nnz = 1.5 M
- kernel: SpMV
- Source: NASA structural analysis problem
- 8x8 dense substructure: exploit this to limit #mem_refs

Speedups on Itanium 2: The Need for Search





Register Profile: Itanium 2

 SpMV BCSR Profil 	e [ref=294.5 Mflop/s;	900 MHz	ltanium 2,	Intel (C v7.0]
--------------------------------------	-----------------------	---------	------------	---------	---------

	12	1.75	1.52	.99	1.33	1.51	1.64	1.79	1.83	1.89	1.75	1.85	1.72	
	11	1.72	1.64	1.12	1.23	1.45	1.60	1.71	1.80	1.88	1.91	1.88	1.97	
	10	1.73	1.47	1.14	1.23	1.38	1.54	1.69	1.67	1.86	1.89	1.88	1.93	
	9	1.54	1.74	1.24	1.00	1.27	1.42	1.55	1.61	1.71	1.73	1.75	1.90	
Ξ	8	3.89	2.40	1.44	1.16	1.16	1.32	1.44	1.47	1.68	1.75	1.77	1.84	
< size	7	3.98	2.04	1.65	1.22	1.04	1.20	1.30	1.44	1.52	1.63	1.65	1.74	
/ block	6	3.79	1.77	1.72	1.44	1.19	1.14	1.23	1.31	1.41	1.52	1.58	1.65	
õ	5	3.20	1.74	1.99	1.52	1.34	1.19	.97	1.17	1.27	1.36	1.42	1.50	
	4	3.32	4.07	1.74	2.37	1.52	1.38	1.19	1.14	.92	1.19	1.22	1.29	
	3	2.55	3.35	.61	1.74	1.97	1.71	1.52	1.34	1.19	1.08	1.03	.88	
	2	1.89	2.54	2.76	2.73	1.62	1.70	1.85	2.40	1.70	1.54	1.27	1.17	
	1	1.00	1.35	1.39	1.44	1.43	1.47	1.48	1.49	1.34	1.42	1.41	1.43	
		1	2	3	4	5	6	7	8	9	10	11	12	
	column block size (c)													

1190 Mflop/s

190 Mflop/s

Power3 - 17% 'rofile [ref=163.9 Mflop/s; 375 MHz Power3, IBM xlc v5]



252 Mflop/s

122 Mflop/s

820 Mflop/s Power4 - 16% ifile [ref=594.9 Mflop/s; 1.3 GHz Power4, IBM xlc v6]

	12	1.28	1.23	1.28	1.06	1.26	1.24	1.29	1.16	1.24	1.17	1.30	1.33		- 799
	11	1.34	1.29	1.26	1.14	1.12	1.26	1.27	1.28	1.26	1.25	1.26	1.28		- 779
														-	- 759
	IU	1.33	1.29	1.21	1.29	1.16	1.14	1.27	1.38	1.27	1.25	1.29	1.29	-	-739
	9	1.32	1.29	1.20	1.20	1.16	1.16	1.21	1.31	1.27	1.27	1.25	1.30	-	-719
	_													-	-699
Ξ	8	1.32	1.27	1.19	1.19	1.25	1.12	1.26	1.15	1.27	1.10	1.13	1.23	-	-679
size	7	1.35	1.30	1.30	1.26	1.17	1.10	1.11	1.24	1.25	1.26	1.27	1.29	-	- 659
ž	_													-	-639
입	6	1.32	1.26	1.34	1.09	1.16	1.19	1.18	1.10	1.20	1.18	1.25	1.28	-	-619
õ	5	1.27	1.24	1.15	1.27	1.15	.91	.99	.94	.80	.81	1.06	1.03	-	- 599
														-	- 579
	4	1.29	1.26	.98	1.24	1.23	1.12	.77	.99	.95	.89	.81	.86	-	- 559
	3	1.23	1.29	1.27	1 18	1 13	1.23	1 18	99	93	82	80	90	-	<mark>-</mark> 539
	_	1.20	1.20	1.21	1.10	1.10	1.20	1.10	.00		.02			-	-519
	2	1.13	1.18	1.25	1.20	1.17	1.00	.92	.87	1.17	1.17	1.15	.99	-	- 499
	1	1.00	1 10	1 15	1 1 5	1 18	1 16	1 14	1.06	1.05	1.02	88	91	-	<mark>-</mark> 479
		1	2	2	4	5	6	7	0	9	10	11	12		459
		I	۷	J	4	colu	mn blo	, ck size	e (c)	э	10	11	14	45	69 Mflop/s

Itanium 1 - 8%

ofile [ref=161.2 Mflop/s; 800 MHz Itanium, Intel C v7]

	12	92	77	99	116	1.30	1.35	1.39	1.39	1.45	1 24	1.22	1.32
	11	1 5 2				1.24	1.05	1 20	1.20	1.44	1.24	1.24	1.20
		1.52	.00	.90	1.11	1.24	1.55	1.58	1.58	1.44	1.24	1.24	1.29
	10	1.34	.74	.88	1.04	1.16	1.30	1.36	1.35	1.38	1.22	1.23	1.32
	9	1.25	.78	.82	1.01	1.10	1.19	1.34	1.35	1.39	1.21	1.20	1.28
Ξ	8	1.47	.72	.77	.92	1.05	1.16	1.22	1.31	1.37	1.22	1.15	1.27
< size	7	1.38	.80	.76	.82	.98	1.11	1.16	1.26	1.34	1.25	1.18	1.32
V DIOCI	6	1.25	.84	.78	.79	.87	.99	1.09	1.18	1.23	1.29	1.14	1.41
õ	5	1.10	1.17	.75	.71	.80	.88	.97	1.06	1.09	1.15	1.14	1.29
	4	1.55	1.30	.80	.72	.71	.77	.80	.94	1.00	1.08	1.11	1.16
	3	1.54	1.04	1.15	.80	.71	.66	.76	.77	.81	.87	.95	.98
	2	1.48	1.48	1.02	1.27	1.05	.83	.70	.67	.66	.68	.77	.77
	1	1.00	1.07	1.05	1.12	.89	.95	1.07	1.21	1.02	.94	.82	.73
		1	2	3	4	5	6	7	8	9	10	11	12
						colu	mn blo	ck size	e (c)				

247	Mflop/	S

	247
	237
	227
	217
_	207
	197
	187
	177
	167
	157
_	147
	137
	127
	117
	107
107 I	Vlflop/s

Itanium 2 - 33%

lon/s [,]	900	MH ₂	Itanii

1.2 Gflop/s

12	1.75	1.52	.99	1.33	1.51	1.64	1.79	1.83	1.89	1.75	1.85	1.72
11	1.72	1.64	1.12	1.23	1.45	1.60	1.71	1.80	1.88	1.91	1.88	1.97
10	1.73	1.47	1.14	1.23	1.38	1.54	1.69	1.67	1.86	1.89	1.88	1.93
9	1.54	1.74	1.24	1.00	1.27	1.42	1.55	1.61	1.71	1.73	1.75	1.90
Ξ ⁸	3.89	2.40	1.44	1.16	1.16	1.32	1.44	1.47	1.68	1.75	1.77	1.84
2 size	3.98	2.04	1.65	1.22	1.04	1.20	1.30	1.44	1.52	1.63	1.65	1.74
block	3.79	1.77	1.72	1.44	1.19	1.14	1.23	1.31	1.41	1.52	1.58	1.65
ē 5	3.20	1.74	1.99	1.52	1.34	1.19	.97	1.17	1.27	1.36	1.42	1.50
4	3.32	4.07	1.74	2.37	1.52	1.38	1.19	1.14	.92	1.19	1.22	1.29
3	2.55	3.35	.61	1.74	1.97	1.71	1.52	1.34	1.19	1.08	1.03	.88
2	1.89	2.54	2.76	2.73	1.62	1.70	1.85	2.40	1.70	1.54	1.27	1.17
1	1.00	1.35	1.39	1.44	1.43	1.47	1.48	1.49	1.34	1.42	1.41	1.43
	1	2	3	4	5	6	7	8	9	10	11	12

190 Mflop/s

Ultra 2i - 11% ofile [ref=35.8 Mflop/s; 333 MHz Sun Ultra 2i, Sun C v6.0]

56 1.73 1.56 1.72	1.71 1.56	1.62 1	1.62	1.82	1.76	1.72	1.87	1.63	12
64 1.75 1.63 1.78	1.71 1.64	1.57 1	1.85	1.85	1.70	1.69	1.70	1.64	11
56 1.66 1.58 1.63	1.67 1.56	1.51 1	1.84	1.84	1.68	1.64	1.61	1.61	10
56 1.70 1.57 1.63	1.68 1.56	1.93 1	1.82	1.78	1.70	1.65	1.67	1.61	9
67 1.67 1.62 1.69	1.93 <mark>1.67</mark>	1.91 1	1.74	1.81	1.79	1.86	1.99	1.58	Э ⁸
84 <mark>1.61 1.60 1.66</mark>	1.86 1.84	1.88 1	1.82	1.76	1.63	1.61	1.57	1.54	7 Size
90 1.88 1.87 <mark>1.63</mark>	2.03 1.90	1.92 2	1.77	1.74	1.85	1.67	1.81	1.51	
94 1.90 1.90 1.82	1.93 1.94	1.92 1	1.80	1.32	1.74	1.68	1.69	1.47	ē 5
81 1.77 1.79 1.69	1.75 1.81	1.82 1	1.60	1.90	1.57	1.60	1.55	1.41	4
64 1.65 1.66 1.69	1.86 <mark>1.64</mark>	1.62 1	1.65	1.60	1.77	1.60	1.46	1.35	3
80 1.87 1.83 1.82	1.54 1.80	1.74 1	1.50	1.52	1.48	1.50	1.38	1.23	2
46 1.48 1.47 1.48	1.45 1.46	1.43 1	1.43	1.41	1.32	1.29	1.23	1.00	1
3 10 11 12	8 9 (c)	7 ck size (6 mn blo	5	4	3	2	1	
90 1.88 1.87 94 1.90 1.90 81 1.77 1.79 64 1.65 1.66 80 1.87 1.83 46 1.48 1.47 9 10 11	2.03 1.90 1.93 1.94 1.75 1.81 1.86 1.64 1.54 1.80 1.45 9 (c)	1.92 2 1.92 1 1.82 1 1.62 1 1.74 1 1.43 1 7 ck size (1.77 1.80 1.60 1.65 1.50 1.43 6 mn blo	1.74 1.32 1.90 1.60 1.52 1.41 5 colu	1.85 1.74 1.57 1.77 1.48 1.32 4	1.67 1.68 1.60 1.60 1.50 1.29 3	1.81 1.69 1.55 1.46 1.38 1.23 2	1.51 1.47 1.41 1.35 1.23 1.00 1	1017 A01 5 4 3 2 1

72 Mflop/s

	71.8
	69.8
	67.8
	65.8
	63.8
	61.8
	59.8
	57.8
	55.8
	53.8
	51.8
	49.8
	47.8
	45.8
_	43.8
	41.8
	39.8
	37.8
	35.8
35 N	flop/s

108 Mflop/s

107.1

102.2

97.2

92.2

87.2

82.2

-77.2

72.2

67.2

62.2

57.2

52.2

47.2

42.2

Pentium III - 21% =42.1 Mflop/s; 500 MHz Pentium III, Intel C v7.0]

12	1.69	1.96	2.22	2.22	2.06	2.26	2.05	2.28	1.88	2.13	1.88	2.22
11	1.69	1.98	2.04	2.17	1.99	2.18	1.89	2.15	1.88	2.17	1.84	2.15
10	1.67	1.87	2.01	2.21	2.07	2.23	1.89	2.27	1.92	2.27	1.83	2.27
9	1.67	2.02	1.99	2.19	1.97	2.15	1.96	2.16	1.82	2.19	1.96	2.18
Ξ ⁸	1.69	1.81	2.18	2.02	1.97	2.21	2.11	2.01	1.86	2.13	1.96	2.09
< size	1.76	2.04	1.78	1.97	1.85	1.94	1.85	2.00	1.91	2.03	1.74	2.04
v bloc	2.17	2.36	2.11	2.14	1.92	2.34	1.99	2.18	1.91	2.17	2.14	2.19
ē ₅	2.15	2.39	2.12	2.10	2.03	1.92	2.21	2.16	2.24	1.97	2.24	2.18
4	2.04	2.28	2.16	2.13	2.17	2.31	2.13	2.44	2.34	2.49	2.36	2.35
3	1.80	2.09	2.31	2.16	2.31	2.25	2.39	2.11	2.42	2.38	2.45	2.23
2	1.44	1.89	2.10	2.18	2.09	2.23	2.25	2.38	2.35	2.54	2.33	2.18
1	1.00	1.53	1.60	1.66	1.71	1.75	1.76	1.73	1.73	1.80	1.80	1.79
	1	2	3	4	5 colu	6 mn blo	7 ck size	8 (c)	9	10	11	12

Ultra 3 - 5%	Profile [ref=50.3 Mflop/s; 900 MHz Sun Ultra 3, Sun C v6.0]	90 Mflop/s

12	1.57	1.59	1.61	1.66	1.66	1.63	1.65	1.77	1.76	1.76	1.76	1.78	
11	1.48	1.59	1.55	1.53	1.66	1.72	1.74	1.74	1.66	1.75	1.74	1.77	
10	1.55	1.56	1.63	1.59	1.67	1.71	1.73	1.78	1.75	1.76	1.74	1.77	82.4 80.4
9	1.54	1.59	1.60	1.61	1.63	1.70	1.73	1.77	1.74	1.74	1.75	1.77	78.4
5 ⁸	1.65	1.58	1.58	1.64	1.65	1.60	1.70	1.76	1.74	1.74	1.74	1.77	76.4 74.4
SIZE	1.54	1.52	1.59	1.59	1.62	1.67	1.69	1.76	1.71	1.73	1.71	1.74	72.4 70.4
6	1.53	1.54	1.55	1.59	1.65	1.69	1.70	1.76	1.73	1.70	1.64	1.72	68.4
ē 5	1.47	1.59	1.57	1.57	1.65	1.59	1.68	1.74	1.72	1.72	1.71	1.72	64.4 64.4
4	1.39	1.53	1.55	1.58	1.55	1.54	1.61	1.69	1.65	1.66	1.59	1.65	62.4 60.4
3	1.34	1.48	1.51	1.53	1.53	1.58	1.62	1.69	1.68	1.68	1.66	1.66	58.4
2	1.16	1.38	1.46	1.55	1.45	1.55	1.55	1.60	1.57	1.62	1.60	1.61	56.4 54.4
1	1.00	1.21	1.31	1.35	1.39	1.42	1.43	1.44	1.46	1.42	1.47	1.47	52.4
1 2 3 4 5 6 7 8 9 10 11 12 column block size (c)										50 Mflop/s			

Pentium III-M - 15% [lop/s; 800 MHz Pentium III-M, Intel C v7.0]

1.98 2.01 2.02

> 3 4 5

.84

2

1.65 1.76 1.78

12	1.70	1.73	2.05	2.06	2.06	2.07	2.04	2.07	1.97	2.06	1.95	2.08
11	1.69	1.77	2.02	2.07	2.03	2.07	1.96	2.07	1.95	2.07	1.92	2.07
10	1.71	1.71	2.05	2.07	2.06	2.08	1.97	2.08	2.00	2.08	1.91	2.09
9	1.73	1.69	1.92	2.06	2.03	2.08	2.05	2.06	1.98	2.08	2.06	2.08
Ξ ⁸	1.75	1.62	2.02	2.05	2.06	2.07	2.07	2.08	2.05	2.08	2.07	2.09
size	1.88	1.66	1.84	2.06	2.00	2.07	2.02	2.08	2.07	2.05	2.02	2.08
9 block	1.93	1.87	1.99	2.05	2.06	2.06	2.07	2.07	2.07	2.08	2.08	2.08
₫ 5	1.89	1.73	1.94	2.03	2.05	2.07	2.07	2.07	2.07	2.07	2.08	2.08
4	1.85	1.96	2.00	1.98	2.04	2.05	2.06	2.05	2.06	2.07	2.08	2.07

2.03

2.01

6 7 8 9

column block size (c)

1.99

2.05

2.02

1.82

2.05 2.05

1.79 1.89

2.02 2.03 2.04 2.04 2.04

1.76

10

2.06 2.06

11

1.88 1.83

12



3 1.74

2 1.52

1 1.00 1.53 1.58

1

42 Mflop/s

Another example of tuning challenges





- More complicated non-zero structure in general
- N = 16614
- NNZ = 1.1M

Zoom in to top corner





- More complicated non-zero structure in general
- N = 16614
- NNZ = 1.1M

3x3 blocks look natural, but...





- More complicated non-zero structure in general
 - Example: 3x3 blocking
 - Logical grid of 3x3 cells
- But would lead to lots of "fill-in"

Extra Work Can Improve Efficiency



- More complicated non-zero structure in general
- Example: 3x3 blocking
 - Logical grid of 3x3 cells
 - Fill-in explicit zeros
 - Unroll 3x3 block multiplies

Selecting Register Block Size r x c



• Off-line benchmark

- Precompute Mflops(r,c) using dense A for each r x c
- Once per machine/architecture
- Run-time "search"
 - Sample A to estimate Fill(r,c) for each r x c
 - − Control cost = O(s·nnz) by controlling fraction $s \in [0,1]$ sampled
 - Control s automatically by computing statistical confidence intervals, by monitoring variance

Run-time heuristic model

– Choose r, c to minimize time ~ Fill(r,c) / Mflops(r,c)

Cost of tuning

- Lower bound: convert matrix in 5 to 40 unblocked SpMVs
- Heuristic: 1 to 11 SpMVs

Tuning only useful when we do many SpMVs

Common case, eg in sparse solvers



עוטוו טו ווומטוווופ אפמו





כנוטוו טו ווומכווווב הבמי

Summary of Other Sequential Performance Optimizations



- Register blocking (RB): up to 4x over CSR
- Variable block splitting: 2.1x over CSR, 1.8x over RB
- Diagonals: 2x over CSR
- Reordering to create dense structure + splitting: 2x over CSR
- Symmetry: 2.8x over CSR, 2.6x over RB
- Cache blocking: 2.8x over CSR
- Multiple vectors (SpMM): 7x over CSR
- And combinations...
- Sparse triangular solve
 - Hybrid sparse/dense data structure: 1.8x over CSR
- Higher-level kernels
 - A·A^T·x, A^T·A·x: 4x over CSR, 1.8x over RB
 - A²·x: 2x over CSR, 1.5x over RB
 - $[A \cdot x, A^2 \cdot x, A^3 \cdot x, ..., A^k \cdot x]$ more to say later





nz = 2287944

Can we reorder the rows and columns to create dense blocks, to accelerate SpMV?

x 10

Source: Accelerator Cavity Design Problem (Ko via Husbands)



Post-RCM (Breadth-first-search) Reordering

Moving nonzeros nearer the diagonal should create dense block, but let's zoom in and see...

100x100 Submatrix Along Diagonal



Here is the top 100x100 submatrix before RCM

"Microscopic" Effect of RCM Reordering



Here is the top 100x100 submatrix after RCM: red entries move to the blue locations. More dense blocks, but could be better, so let's try reordering again, using TSP (Travelling Saleman Problem)

"Microscopic" Effect of Combined RCM+TSP Reordering



Here is the top 100x100 submatrix after RCM and TSP: red entries move to the blue locations. Lots of dense blocks, as desired!

Speedups (using symmetry too):

Itanium 2: 1.7x Pentium 4: 2.1x Power 4: 2.1x Ultra 3: 3.3x Multicore SMPs Used for Tuning SpMV

Multicore SMPs Used for Tuning SpMV

Intel Xeon E5345 (Clovertown)

AMD Opteron 2356 (Barcelona)

- Cache based
- 8 Threads
- 75 GFlops
- 21/10 GB/s R/W BW
- Sun T2+ T5140 (Victoria Falls)
 - Cache based
 - 128 Threads (CMT)
 - NUMA
 - 19 GFlops
 - 42/21 GB/s R/W BW

- Cache based
- 8 Threads
- NUMA
- 74 GFlops
- 21 GB/s R/W BW

IBM QS20 Cell Blade

- Local-Store based
- 16 Threads
- NUMA
- 29 Gflops (SPEs only)
- 51 GB/s R/W BW

Set of 14 test matrices

• All bigger than the caches of our SMPs

2K x 2K Dense matrix stored in sparse format

Dense Well Structured (sorted by nonzeros/row) FEM / FEM / Wind FEM / FEM / QCD Protein Economics Epidemiology Spheres Cantilever Tunnel Harbor Ship **Poorly Structured** hodgepodge Circuit webbase Accelerator

Extreme Aspect Ratio (linear programming)

LΡ

SpMV Performance: Naive parallelization

- Out-of-the box SpMV performance on a suite of 14 matrices
- Scalability isn't great: Compare to # threads
 8

128 16

Naïve

SpMV Performance: NUMA and Software Prefetching

- NUMA-aware allocation is essential on NUMA SMPs.
- Explicit software prefetching can boost bandwidth and change cache replacement policies
- used exhaustive search

SpMV Performance: "Matrix Compression"

- Compression includes
 - register blocking
 - other formats
 - smaller indices
- Use heuristic rather than search

SpMV Performance: cache and TLB blocking

SpMV Performance: Architecture specific optimizations





SpMV Performance: max speedup







QCD

FEM-Ship

Tunnel

FEM-Har

FEM-Sphr FEM-Cant

Dense

Epidem

Circuit

Webbase

FEM

5

Median

- Fully auto-tuned SpMV performance across the suite of matrices
- Included SPE/local store optimized version
- Why do some optimizations work better on some architectures?



Optimized Sparse Kernel Interface - pOSKI

- Provides sparse kernels automatically tuned for matrix & machine
 - BLAS-style functionality: SpMV, $Ax \& A^Ty$
 - Hides complexity of run-time tuning
- Faster than previous implementations
 - Up to 7.8x over reference serial implementation on Sandy Bridge E
 - Up to 4.5x over OSKI on Sandy Bridge E
 - Up to 2.1x over MKL on Nehalem
- bebop.cs.berkeley.edu/poski
- Ongoing work by the Berkeley Benchmarking and Optimization (BeBop) group

Optimizations in pOSKI, so far

- Fully automatic heuristics for
 - Sparse matrix-vector multiply (Ax, A^Tx)
 - Register-level blocking, Thread-level blocking
 - SIMD, software prefetching, software pipelining, loop unrolling
 - NUMA-aware allocations
- "Plug-in" extensibility
 - Very advanced users may write their own heuristics, create new data structures/code variants and dynamically add them to the system
- Other kernels just in OSKI so far
 - $A^{T}Ax, A^{k}x$
 - A⁻¹x : Sparse triangular solver (SpTS)
- Other optimizations under development
 - Cache-level blocking, Reordering (RCM, TSP), variable block structure, index compressing, Symmetric storage, etc.





To user: Matrix handle for kernel calls

How pOSKI Tunes (Overview)



• At library build/install-time

- Generate code variants
 - Code generator (Python) generates code variants for various implementations
- Collect benchmark data
 - Measures and records speed of possible sparse data structure and code variants on target architecture
- Select best code variants & benchmark data
 - prefetching distance, SIMD implementation
- Installation process uses standard, portable GNU AutoTools

• At run-time

- Library "tunes" using heuristic models
 - Models analyze user's matrix & benchmark data to choose optimized data structure and code
 - User may re-collect benchmark data with user's sparse matrix (under development)
- Non-trivial tuning cost: up to ~40 mat-vecs
 - Library limits the time it spends tuning based on estimated workload
 - provided by user or inferred by library
 - User may reduce cost by saving tuning results for application on future runs with same or similar matrix (under development)

How to call pOSKI: Basic Usage

- May gradually migrate existing apps
 - Step 1: "Wrap" existing data structures
 - Step 2: Make BLAS-like kernel calls

int* ptr = ..., *ind = ...; double* val = ...; /* Matrix, in CSR format */
double* x = ..., *y = ...; /* Let x and y be two dense vectors */

/* Compute y = $\beta \cdot y + \alpha \cdot A \cdot x$, 500 times */
for(i = 0; i < 500; i++)
my_matmult(ptr, ind, val, α , x, β , y);

How to call pOSKI: Basic Usage

- May gradually migrate existing apps
 - Step 1: "Wrap" existing data structures
 - Step 2: Make BLAS-like kernel calls

int* ptr = ..., *ind = ...; double* val = ...; /* Matrix, in CSR format */
double* x = ..., *y = ...; /* Let x and y be two dense vectors */
/* Step 1: Create a default pOSKI thread object */
poski_threadarg_t *poski_thread = poski_InitThread();
/* Step 2: Create pOSKI wrappers around this data */
poski_mat_t A_tunable = poski_CreateMatCSR(ptr, ind, val, nrows, ncols,
 nnz, SHARE_INPUTMAT, poski_thread, NULL, ...);
poski_vec_t x_view = poski_CreateVec(x, ncols, UNIT_STRIDE, NULL);
poski_vec_t y_view = poski_CreateVec(y, nrows, UNIT_STRIDE, NULL);

```
/* Compute y = β·y + α·A·x, 500 times */
for( i = 0; i < 500; i++ )
    my_matmult( ptr, ind, val, α, x, β, y );</pre>
```

How to call pOSKI: Basic Usage

- May gradually migrate existing apps
 - Step 1: "Wrap" existing data structures
 - Step 2: Make BLAS-like kernel calls

int* ptr = ..., *ind = ...; double* val = ...; /* Matrix, in CSR format */
double* x = ..., *y = ...; /* Let x and y be two dense vectors */
/* Step 1: Create a default pOSKI thread object */
poski_threadarg_t *poski_thread = poski_InitThread();
/* Step 2: Create pOSKI wrappers around this data */
poski_mat_t A_tunable = poski_CreateMatCSR(ptr, ind, val, nrows, ncols,
 nnz, SHARE_INPUTMAT, poski_thread, NULL, ...);
poski_vec_t x_view = poski_CreateVec(x, ncols, UNIT_STRIDE, NULL);
poski_vec_t y_view = poski_CreateVec(y, nrows, UNIT_STRIDE, NULL);

/* Step 3: Compute $y = \beta \cdot y + \alpha \cdot A \cdot x$, 500 times */

for(i = 0; i < 500; i++)</pre>

poski_MatMult(A_tunable, OP_NORMAL, α , x_view, β , y_view);

How to call pOSKI: Tune with Explicit Hints



User calls "tune" routine (optional)
 May provide explicit tuning hints

```
poski_mat_t A_tunable = poski_CreateMatCSR( ... );
    /* ... */
/* Tell pOSKI we will call SpMV 500 times (workload hint) */
poski_TuneHint_MatMult(A_tunable, OP_NORMAL, α, x_view, β, y_view,500);
/* Tell pOSKI we think the matrix has 8x8 blocks (structural hint) */
poski_TuneHint_Structure(A_tunable, HINT_SINGLE_BLOCKSIZE, 8, 8);
```

```
/* Ask pOSKI to tune */
poski_TuneMat(A_tunable);
```

```
for( i = 0; i < 500; i++ )
    poski_MatMult(A_tunable, OP_NORMAL, α, x_view, β, y_view);</pre>
```

How to call pOSKI: Implicit Tuning

- Ask library to infer workload (optional)
 - Library profiles all kernel calls
 - May periodically re-tune

```
poski_mat_t A_tunable = poski_CreateMatCSR( ... );
/* ... */
```

```
for( i = 0; i < 500; i++ ) {
    poski_MatMult(A_tunable, OP_NORMAL, α, x_view, β, y_view);
    poski_TuneMat(A_tunable); /* Ask pOSKI to tune */
}</pre>
```

Performance on Intel Sandy Bridge

- Jaketown: i7-3960X @ 3.3 GHz
- #Cores: 6 (2 threads per core), L3:15MB
- pOSKI SpMV (Ax) with double precision float-point
- MKL Sparse BLAS Level 2: *mkl_dcsrmv()*



Avoiding Communication in Sparse Linear Algebra

- Computational intensity of one SpMV ≤ 2, limits performance
- k-steps of typical iterative solver for Ax=b or Ax=λx
 - Does k SpMVs with starting vector (eg with b, if solving Ax=b)
 - Finds "best" solution among all linear combinations of these k+1 vectors
 - Many such "Krylov Subspace Methods"
 - Conjugate Gradients, GMRES, Lanczos, Arnoldi, ...
- Goal: minimize communication in Krylov Subspace Methods
 - Assume matrix "well-partitioned," with modest surface-to-volume ratio
 - Parallel implementation
 - Conventional: O(k log p) messages, because k calls to SpMV
 - New: O(log p) messages optimal
 - Serial implementation
 - Conventional: O(k) moves of data from slow to fast memory
 - New: O(1) moves of data optimal
- Lots of speed up possible (modeled and measured)
 - Price: some redundant computation, numerical stability issues

• Replace k iterations of $y = A \cdot x$ with $[Ax, A^2x, ..., A^kx]$



- Example: A tridiagonal, n=32, k=3
- Works for any "well-partitioned" A

• Replace k iterations of $y = A \cdot x$ with $[Ax, A^2x, ..., A^kx]$



• Replace k iterations of $y = A \cdot x$ with $[Ax, A^2x, ..., A^kx]$



• Replace k iterations of $y = A \cdot x$ with $[Ax, A^2x, ..., A^kx]$



• Replace k iterations of $y = A \cdot x$ with $[Ax, A^2x, ..., A^kx]$



• Replace k iterations of $y = A \cdot x$ with $[Ax, A^2x, ..., A^kx]$



- Replace k iterations of $y = A \cdot x$ with $[Ax, A^2x, ..., A^kx]$
- Sequential Algorithm



- Replace k iterations of $y = A \cdot x$ with $[Ax, A^2x, ..., A^kx]$
- Sequential Algorithm



- Replace k iterations of $y = A \cdot x$ with $[Ax, A^2x, ..., A^kx]$
- Sequential Algorithm



- Replace k iterations of $y = A \cdot x$ with $[Ax, A^2x, ..., A^kx]$
- Sequential Algorithm



- Replace k iterations of $y = A \cdot x$ with $[Ax, A^2x, ..., A^kx]$
- Parallel Algorithm



- Example: A tridiagonal, n=32, k=3
- Each processor communicates once with neighbors

- Replace k iterations of $y = A \cdot x$ with $[Ax, A^2x, ..., A^kx]$
- Parallel Algorithm



- Example: A tridiagonal, n=32, k=3
- Each processor works on (overlapping) trapezoid

Same idea works for general sparse matrices

Partitioning by rows → Graph partitioning

Processing left to right → Traveling Salesman Problem





What about multicore?

- Two kinds of communication to minimize
 - Between processors on the chip
 - Between on-chip cache and off-chip DRAM
- Use hybrid of both techniques described so far
 - Use parallel optimization so each core can work independently
 - Use sequential optimization to minimize off-chip
 DRAM traffic of each core

Speedups on Intel Clovertown (8 core) Test matrices include stencils and practical matrices See SC09 paper on bebop.cs.berkeley.edu for details



Minimizing Communication of GMRES

Classical GMRES for Ax=b

```
for i=1 to k
  w = A * v(i-1)
  MGS(w, v(0),...,v(i-1))
     ... Modified Gram-Schmidt
     ... to make w orthogonal
  update v(i), H
     \dots H = matrix of coeffs
     ... from MGS
 endfor
 solve LSQ problem with H for x
Communication cost =
  k copies of A, vectors from
  slow to fast memory
```

Communication-Avoiding GMRES, ver. 1

W = [v, Av, A²v, ..., A^kv] [Q,R] = TSQR(W) ... "Tall Skinny QR" ... new optimal QR discussed before Build H from R solve LSQ problem with H for x

Communication cost = O(1) copy of A, vectors from slow to fast memory

Let's confirm that we still get the right answer ...



Minimizing Communication of GMRES (and getting the right answer)

Communication-Avoiding GMRES, ver. 2

```
W = [v, p_1(A)v, p_2(A)v, ..., p_k(A)v]
    ... where p_i(A)v is a degree-i polynomial in A multiplied by v
    ... polynomials chosen to keep vectors independent
 [Q,R] = TSQR(W)
    ... "Tall Skinny QR"
    ... new optimal QR discussed before
 Build H from R
    ... slightly different R from before
 solve LSQ problem with H for x
Communication cost still optimal:
  O(1) copy of A, vectors from
  slow to fast memory
```



Speed ups on 8-core Clovertown

CA-GMRES = Communication-Avoiding GMRES



Paper by Mohiyuddin, Hoemmen, D. in Supercomputing09



Summary of what is known, open

• GMRES

- Can independently choose k to optimize speed, restart length r to optimize convergence
- <u>Need to "co-tune" Akx kernel and TSQR</u>
- Know how to use more stable polynomial bases
- Proven speedups
- Can similarly reorganize other Krylov methods
 - Arnoldi and Lanczos, for $Ax = \lambda x$ and for $Ax = \lambda Mx$
 - Conjugate Gradients (CG), for Ax = b
 - Biconjugate Gradients (BiCG), CG Squared (CGS), BiCGStab for Ax=b
 - Other Krylov methods?
- Preconditioning how to handle MAx = Mb



What is a sparse matrix?



 How much infrastructure (for code creation, tuning or interfaces) can we reuse for all these cases?

Sparse Conclusions



- Fast code must minimize communication
 - Especially for sparse matrix computations because communication dominates
- Generating fast code for a single SpMV
 - Design space of possible algorithms must be searched at run-time, when sparse matrix available
 - Design space should be searched automatically
- Biggest speedups from minimizing communication in an entire sparse solver
 - Many more opportunities to minimize communication in multiple
 SpMVs than in one
 - Requires transforming entire algorithm
 - Lots of open problems
- For more information, see bebop.cs.berkeley.edu


STRUCTURED GRID MOTIF

Source: Sam Williams





- Applying the finite difference method to PDEs on structured grids produces **stencil operators** that must be applied to all points in the discretized grid.
- Consider the 7-point Laplacian Operator
- Challenged by bandwidth, temporal reuse, efficient SIMD, etc... but trivial to (correctly) parallelize
- most optimizations can be independently implemented, (but not performance independent)









Lattice Boltzmann Methods

- LBMHD simulates charged plasmas in a magnetic field (MHD) via Lattice Boltzmann Method (LBM) applied to CFD and Maxwell's equations.
- To monitor density, momentum, and magnetic field, it requires maintaining two "velocity" distributions
 - 27 (scalar) element velocity distribution for momentum
 - 15 (Cartesian) element velocity distribution for magnetic field
 - = 632 bytes / grid point / time step
- Jacobi-like time evolution requires ~1300 flops and ~1200 bytes of memory traffic





Lattice Boltzmann Methods

- Challenged by:
 - The higher flop:byte ratio of ~1.0 is still bandwidth-limiting
 - TLB locality (touch 150 pages per lattice update)
 - cache associativity (150 disjoint lines)
 - efficient SIMDization
- easy to (correctly) parallelize
- explicit SIMDization & SW prefetch are dependent on unrolling
- Ultimately, 2 of 3 machines are bandwidth-limited



*collision() only



Lattice Boltzmann Methods

- Distributed Memory & Hybrid
- MPI, MPI+pthreads, MPI+OpenMP (SPMD, SPMD², SPMD+Fork/Join)
- Observe that for this large problem, auto-tuning flat MPI delivered significant boosts (2.5x)
- Extending auto-tuning to include the domain decomposition and balance between threads and processes provided an extra 17%
- 2 processes with 2 threads was best (true for Pthreads and OpenMP)





DELIVERING AUTOTUNING WITH SEJITS

Source: Shoaib Kamil



What is SEJITS?

- Goal: Let non-expert programmers quickly write their algorithms in an easy-to-use language, but still get high performance

 First example: Python
- By using common "patterns" to write algorithms, and hints about tuning opportunities, enable system to autotune
- SEJITS = Selective Embedded Just-in-time Specialization

Delivering Autotuning via SEJITS





Several examples exist now: Structured Grids/Stencils CA-Conjugate Gradient Tuned SpMV over other semirings

Summary



- "Design spaces" for algorithms and implementations are large and growing
- Finding the best algorithm/implementation by hand is hard and getting harder
- Ideally, we would have a database of "techniques" that would grow over time, and be searched automatically whenever a new input and/or machine comes along
- Lots of work to do...