Parallel & Cluster Computing: Linear Algebra

Paul Gray, University of Northern Iowa
Henry Neeman, University of Oklahoma
Charlie Peck, Earlham College
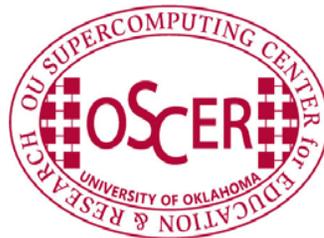
Tuesday October 2 2007
University of Oklahoma

Preinvented Wheels

Many simulations perform fairly common tasks; for example, solving systems of equations:

$$\mathbf{A}\mathbf{x} = \mathbf{b}$$

where $\mathbf{A}$ is the matrix of coefficients, $\mathbf{x}$ is the vector of unknowns and $\mathbf{b}$ is the vector of knowns.

$$
\begin{bmatrix}
a_{1,1} & a_{1,2} & a_{1,3} & \cdots & a_{1,n} \\
a_{2,1} & a_{2,2} & a_{2,3} & \cdots & a_{2,n} \\
a_{3,1} & a_{3,2} & a_{3,3} & \cdots & a_{3,n} \\
\vdots & \vdots & \vdots & \ddots & \vdots \\
a_{n,1} & a_{n,2} & a_{n,3} & \cdots & a_{n,n}
\end{bmatrix}
\begin{bmatrix}
x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_n
\end{bmatrix}
=
\begin{bmatrix}
b_1 \\ b_2 \\ b_3 \\ \vdots \\ b_n
\end{bmatrix}
$$
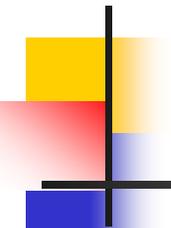
Scientific Libraries

Because some tasks are quite common across many science and engineering applications, groups of researchers have put a lot of effort into writing ***scientific libraries***: collections of routines for performing these commonly-used tasks (e.g., linear algebra solvers).

The people who write these libraries know a lot more about these things than we do.

So, a good strategy is to use their libraries, rather than trying to write our own.
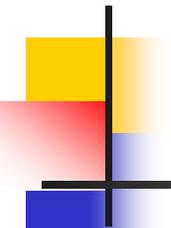
Solver Libraries

Probably the most common scientific computing task is solving a system of equations

$$\mathbf{A}\mathbf{x} = \mathbf{b}$$

where **A** is a matrix of coefficients, **x** is a vector of unknowns, and **b** is a vector of knowns.

The goal is to solve for **x**.

Solving Systems of Equations

Don'ts:

- **<u>Don't</u>** invert the matrix ($\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$). That's much more costly than solving directly, and much more prone to numerical error.

- **<u>Don't</u>** write your own solver code. There are people who devote their whole careers to writing solvers. They know a lot more about writing solvers than we do.

Solving Do's

Do's:

- **Do** use standard, portable solver libraries.

- **Do** use a version that's tuned for the platform you're running on, if available.

- **Do** use the information that you have about your system to pick the most efficient solver.

All About Your Matrix

If you know things about your matrix, you maybe can use a more efficient solver.

- Symmetric: $a_{i,j} = a_{j,i}$

- Positive definite: $\mathbf{x}^T\mathbf{A}\mathbf{x} > 0$ for all $\mathbf{x} \neq 0$ (e.g., if all eigenvalues are positive)

- Banded:

0 except on the bands
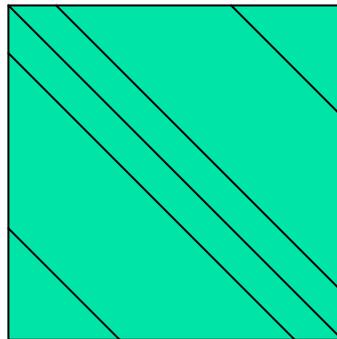
Tridiagonal:
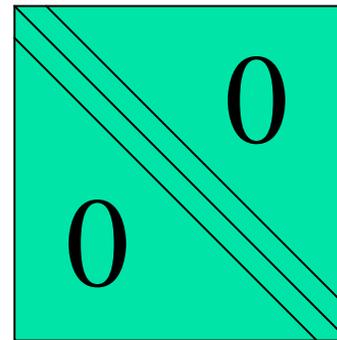
and …

Sparse Matrices

A *__sparse matrix__* is a matrix that has mostly zeros in it. "Mostly" is vaguely defined, but a good rule of thumb is that a matrix is sparse if more than, say, 90-95% of its entries are zero. (A non-sparse matrix is *__dense__*.)
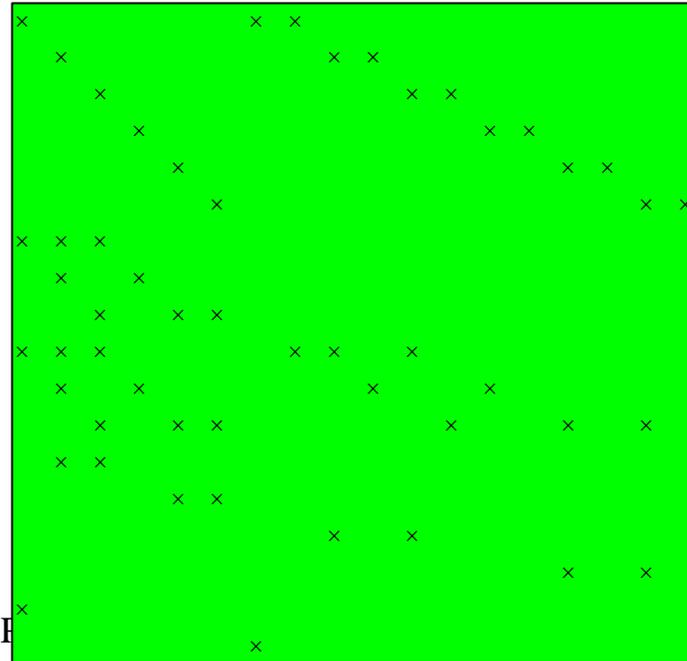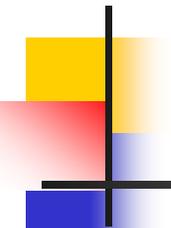
Linear Algebra Libraries

- BLAS [1],[2]
- ATLAS[3]
- LAPACK[4]
- ScaLAPACK[5]
- PETSc[6],[7],[8]
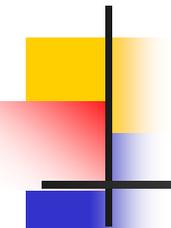
BLAS

The **<u>Basic Linear Algebra Subprograms</u>** (BLAS) are a set of low level linear algebra routines:

- Level 1: Vector-vector (e.g., dot product)
- Level 2: Matrix-vector (e.g., matrix-vector multiply)
- Level 3: Matrix-matrix (e.g., matrix-matrix multiply)

Many linear algebra packages, including LAPACK, ScaLAPACK and PETSc, are built on top of BLAS.

Most supercomputer vendors have versions of BLAS that are highly tuned for their platforms.
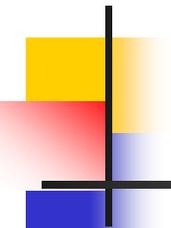
ATLAS

The **Automatically Tuned Linear Algebra Software** package (ATLAS) is a self-tuned version of BLAS (it also includes a few LAPACK routines).

When it's installed, it tests and times a variety of approaches to each routine, and selects the version that runs the fastest.

ATLAS is substantially faster than the generic version of BLAS.

And, it's free!
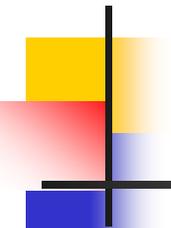
Goto BLAS

In the past few years, a new version of BLAS has been released, developed by Kazushige Goto (currently at UT Austin).

This version is unusual, because instead of optimizing for cache, it optimizes for the ***Translation Lookaside Buffer*** (TLB), which is a special little cache that often is ignored by software developers.

Goto realized that optimizing for the TLB would be more effective than optimizing for cache.

ATLAS vs. BLAS Performance

BETTER

Aspen Systems Linux Pentium4 Xeon 2 GHz BLAS Performance (gcc -O3)

ATLAS DGEMM: 2.76 GFLOP/s = 69% of peak

Generic DGEMM: 0.91 GFLOP/s = 23% of peak
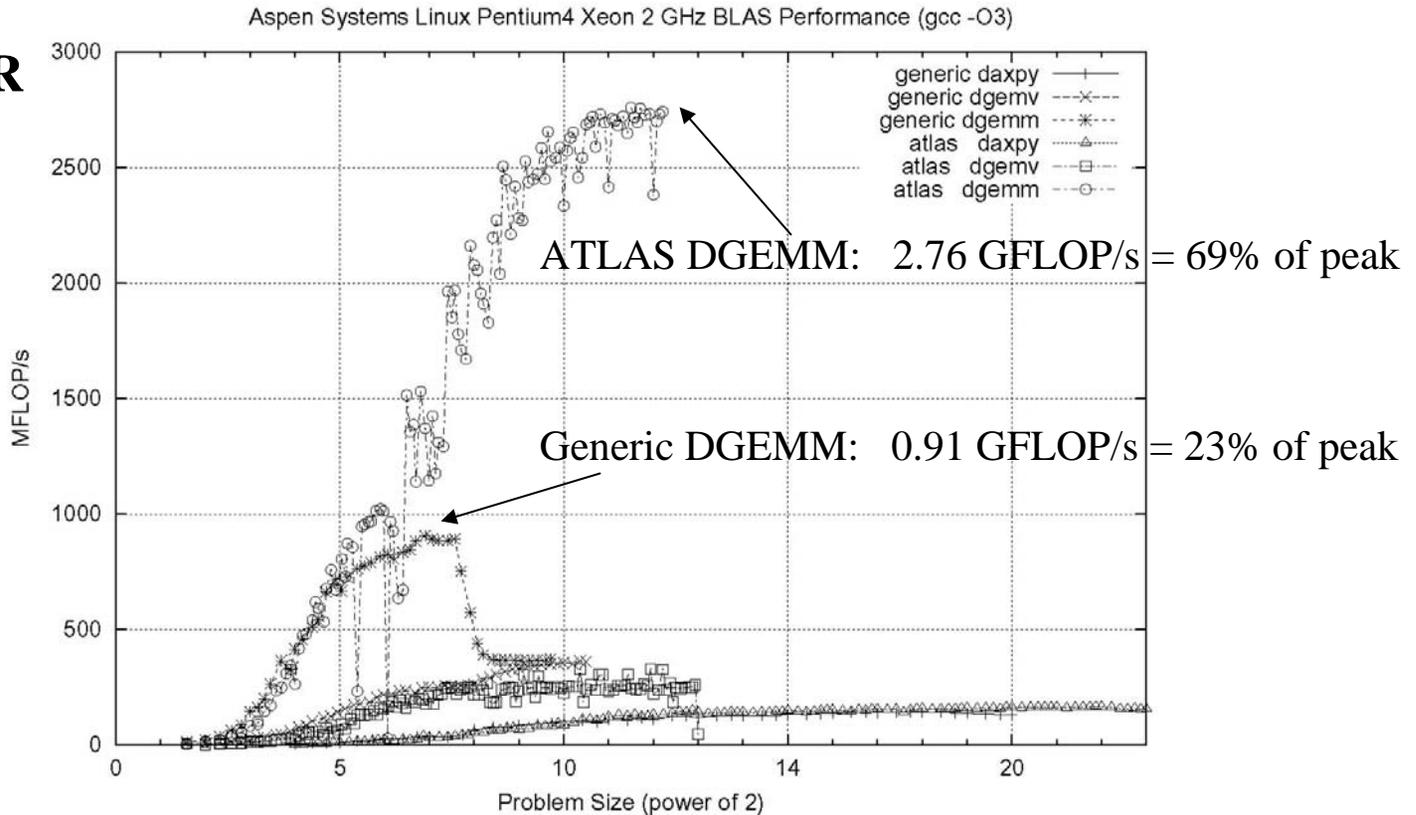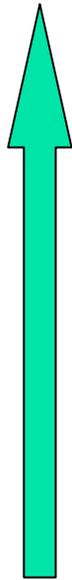
DGEMM: **D**ouble precision **GE**neral **M**atrix-**M**atrix multiply
DGEMV: **D**ouble precision **GE**neral **M**atrix-**V**ector multiply

LAPACK

LAPACK (Linear Algebra PACKage) solves dense or special-case sparse systems of equations depending on matrix properties such as:

- Precision: single, double
- Data type: real, complex
- Shape: diagonal, bidiagonal, tridiagonal, banded, triangular, trapezoidal, Hesenberg, general dense
- Properties: orthogonal, positive definite, Hermetian (complex), symmetric, general

LAPACK is built on top of BLAS, which means it can benefit from ATLAS or Goto BLAS.

LAPACK Example

```fortran
REAL,DIMENSION(numrows,numcols) :: A
REAL,DIMENSION(numrows)         :: B
REAL,DIMENSION(numrows)         :: X
INTEGER,DIMENSION(numrows)      :: pivot
INTEGER :: row, col, info, numrhs = 1

DO row = 1, numrows
  B(row) = …
END DO
DO col = 1, numcols
  DO row = 1, numrows
    A(row,col) = …
  END DO
END DO
CALL sgesv(numrows, numrhs, A, numrows, pivot, &
&          B, numrows, info)
DO col = 1, numcols
  X(col) = B(col)
END DO
```
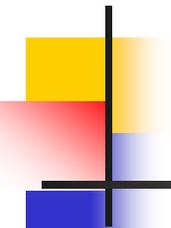
LAPACK: a Library and an API

LAPACK is a library that you can download for free from the Web:

`www.netlib.org`

But, it's also an Application Programming Interface (API): a definition of a set of routines, their arguments, and their behaviors.

So, anyone can write an implementation of LAPACK.
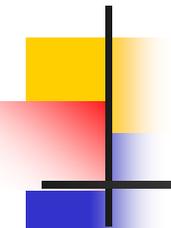
It's Good to Be Popular

LAPACK is a good choice for non-parallelized solving, because its popularity has convinced many supercomputer vendors to write their own, highly tuned versions.

The API for the LAPACK routines is the same as the portable version from NetLib (www.netlib.org), but the performance can be much better, via either ATLAS or proprietary vendor-tuned versions.

Also, some vendors have shared memory parallel versions of LAPACK.
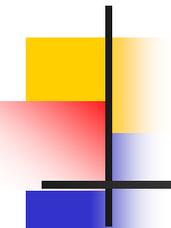
LAPACK Performance

Because LAPACK uses BLAS, it's about as fast as BLAS. For example, DGESV (Double precision General SolVer) on a 2 GHz Pentium4 using ATLAS gets 65% of peak, compared to 69% of peak for Matrix-Matrix multiply.

In fact, an older version of LAPACK, called LINPACK, is used to determine the top 500 supercomputers in the world.
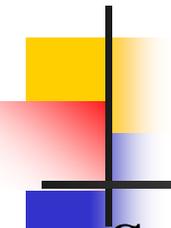
ScaLAPACK

<u>ScaLAPACK</u> is the distributed parallel (MPI) version of LAPACK. It actually contains only a subset of the LAPACK routines, and has a somewhat awkward Application Programming Interface (API).

Like LAPACK, ScaLAPACK is also available from

`www.netlib.org`.
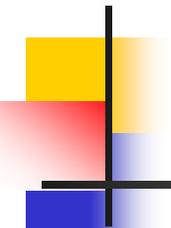
PETSc

PETSc (Portable, Extensible Toolkit for Scientific Computation) is a solver library for sparse matrices that uses distributed parallelism (MPI).

PETSc is designed for general sparse matrices with no special properties, but it also works well for sparse matrices with simple properties like banding and symmetry.

It has a simpler, more intuitive Application Programming Interface than ScaLAPACK.
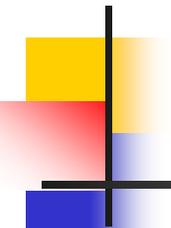
Pick Your Solver Package

- Dense Matrix
  - Serial: LAPACK
  - Shared Memory Parallel: vendor-tuned LAPACK
  - Distributed Parallel: ScaLAPACK
- Sparse Matrix: PETSc

To Learn More

http://www.oscer.ou.edu/

http://www.sc-conference.org/

Thanks for your attention!

Questions?