

# HPC Challenge Awards

## Class 2 - Productivity

Cleve Moler

The MathWorks, Inc.

SC|06, November 14, 2006

Copyright (c), 2006, The MathWorks, Inc.

# Parallel Computing with MATLAB

Distributed Computing Toolbox

Version 3.0 (2006b) is available now.

Version ?.? (2007a) is in development.

# Parallel Model

Distributed memory multiprocessors

Single Program, Distributed Data

All processors run the same MATLAB program

# **“Think Matrices, Not Messages”**

**A MATLAB process is a “lab”**

**numlabs**

**labindex**

**parfor**

**dcolon**

**darray**

**gop**

**labSend**

**labReceive**

**labSendReceive**

# MathWorks HPC Lab

Manufactured by Rackable Systems  
16 dual socket, dual core Opteron 285s  
2.6 GHz  
4 Gbytes/node  
Gigabit Ethernet  
Linux  
AMD ACML math library  
MPICH2

numlabs = 1 ... 64

**Stream**

```
% HPC_STREAM
% Each lab works independently.

m = 2000000;
b = randn(m,1);
c = randn(m,1);
alpha = randn;

% Repeat the test n times
n = 10;
t = zeros(n,1);
for k = 1:n
    tic
    a = b + alpha*c;
    t(k) = toc;
end

% Performance uses the best time
perf = 24*m/min(t)/1.e9;
disp(perf)
```

**P>> numlabs**

**16**

**P>> stream**

<b>1:</b>	<b>2.4539</b>
<b>2:</b>	<b>2.4784</b>
<b>3:</b>	<b>2.4260</b>
<b>4:</b>	<b>2.4303</b>
<b>5:</b>	<b>2.4236</b>
<b>6:</b>	<b>2.4758</b>
<b>7:</b>	<b>2.4194</b>
<b>8:</b>	<b>2.4306</b>
<b>9:</b>	<b>2.4353</b>
<b>10:</b>	<b>2.4486</b>
<b>11:</b>	<b>2.4409</b>
<b>12:</b>	<b>2.4608</b>
<b>13:</b>	<b>2.4682</b>
<b>14:</b>	<b>2.4516</b>
<b>15:</b>	<b>2.4315</b>
<b>16:</b>	<b>2.4341</b>



**P>> numlabs**

**64**

**P>> stream**

**1: 1.5969**  
**2: 1.6194**  
**3: 1.6019**

**.**

**.**

**.**

**12: 1.7027**

**.**

**.**

**.**

**43: 1.5944**

**.**

**.**

**.**

**62: 1.6796**

**63: 1.6250**

**64: 1.6019**

$A \setminus b$

# Matrix Computations

M

Productivity

Sparse matrices

Improving performance

ScaLAPACK

Performance

Expertise

# Redistribute on the fly

M

1d unblocked noncyclic

Rows or columns

Efficient for common array operations

ScaLAPACK


2d block cyclic

Scales to many processors

# Distributed colon operation

`dcolon(1, 1, 10)`


[1 2 3 4 5 6 7 8 9 10]



[1 2 3] [4 5 6] [7 8] [9 10]

# Distributed arrays

`A = darray(Aloc, [10, 10], 2)`



```
function B = mat2scal a(A)
```

```
%  
% A = MATLAB dcol on column di stri buti on  
%  
% 11 12 | 13 14 | 15 16 | 17 | 18 | 19  
% 21 22 | 23 24 | 25 26 | 27 | 28 | 29  
% 31 32 | 33 34 | 35 36 | 37 | 38 | 39  
% 41 42 | 43 44 | 45 46 | 47 | 48 | 49  
% 51 52 | 53 54 | 55 56 | 57 | 58 | 59  
% 61 62 | 63 64 | 65 66 | 67 | 68 | 69  
% 71 72 | 73 74 | 75 76 | 77 | 78 | 79  
% 81 82 | 83 84 | 85 86 | 87 | 88 | 89  
% 91 92 | 93 94 | 95 96 | 97 | 98 | 99  
%  
% B = ScaLAPACK di stri buti on  
%  
% 11 12 . 17 18 | 13 14 . 19 | 15 16  
% 21 22 . 27 28 | 23 24 . 29 | 25 26  
% .....  
% 51 52 . 57 58 | 53 54 . 59 | 55 56  
% 61 62 . 67 68 | 63 64 . 69 | 65 66  
% .....  
% 91 92 . 97 98 | 93 94 . 99 | 95 96  
% -----  
% 31 32 . 37 38 | 33 34 . 39 | 35 36  
% 41 42 . 47 48 | 43 44 . 49 | 45 46  
% .....  
% 71 72 . 77 78 | 73 74 . 79 | 75 76  
% 81 82 . 87 88 | 83 84 . 89 | 85 86
```

# Scaled speedup

If

$$p = \text{numlabs}$$

$$n \sim p^{1/3}$$


then, with perfect speedup

$$\text{time} \sim \text{constant}$$

$$\text{megaflops} \sim p$$



$$n = 4000 * p^{1/3}$$




# Main program

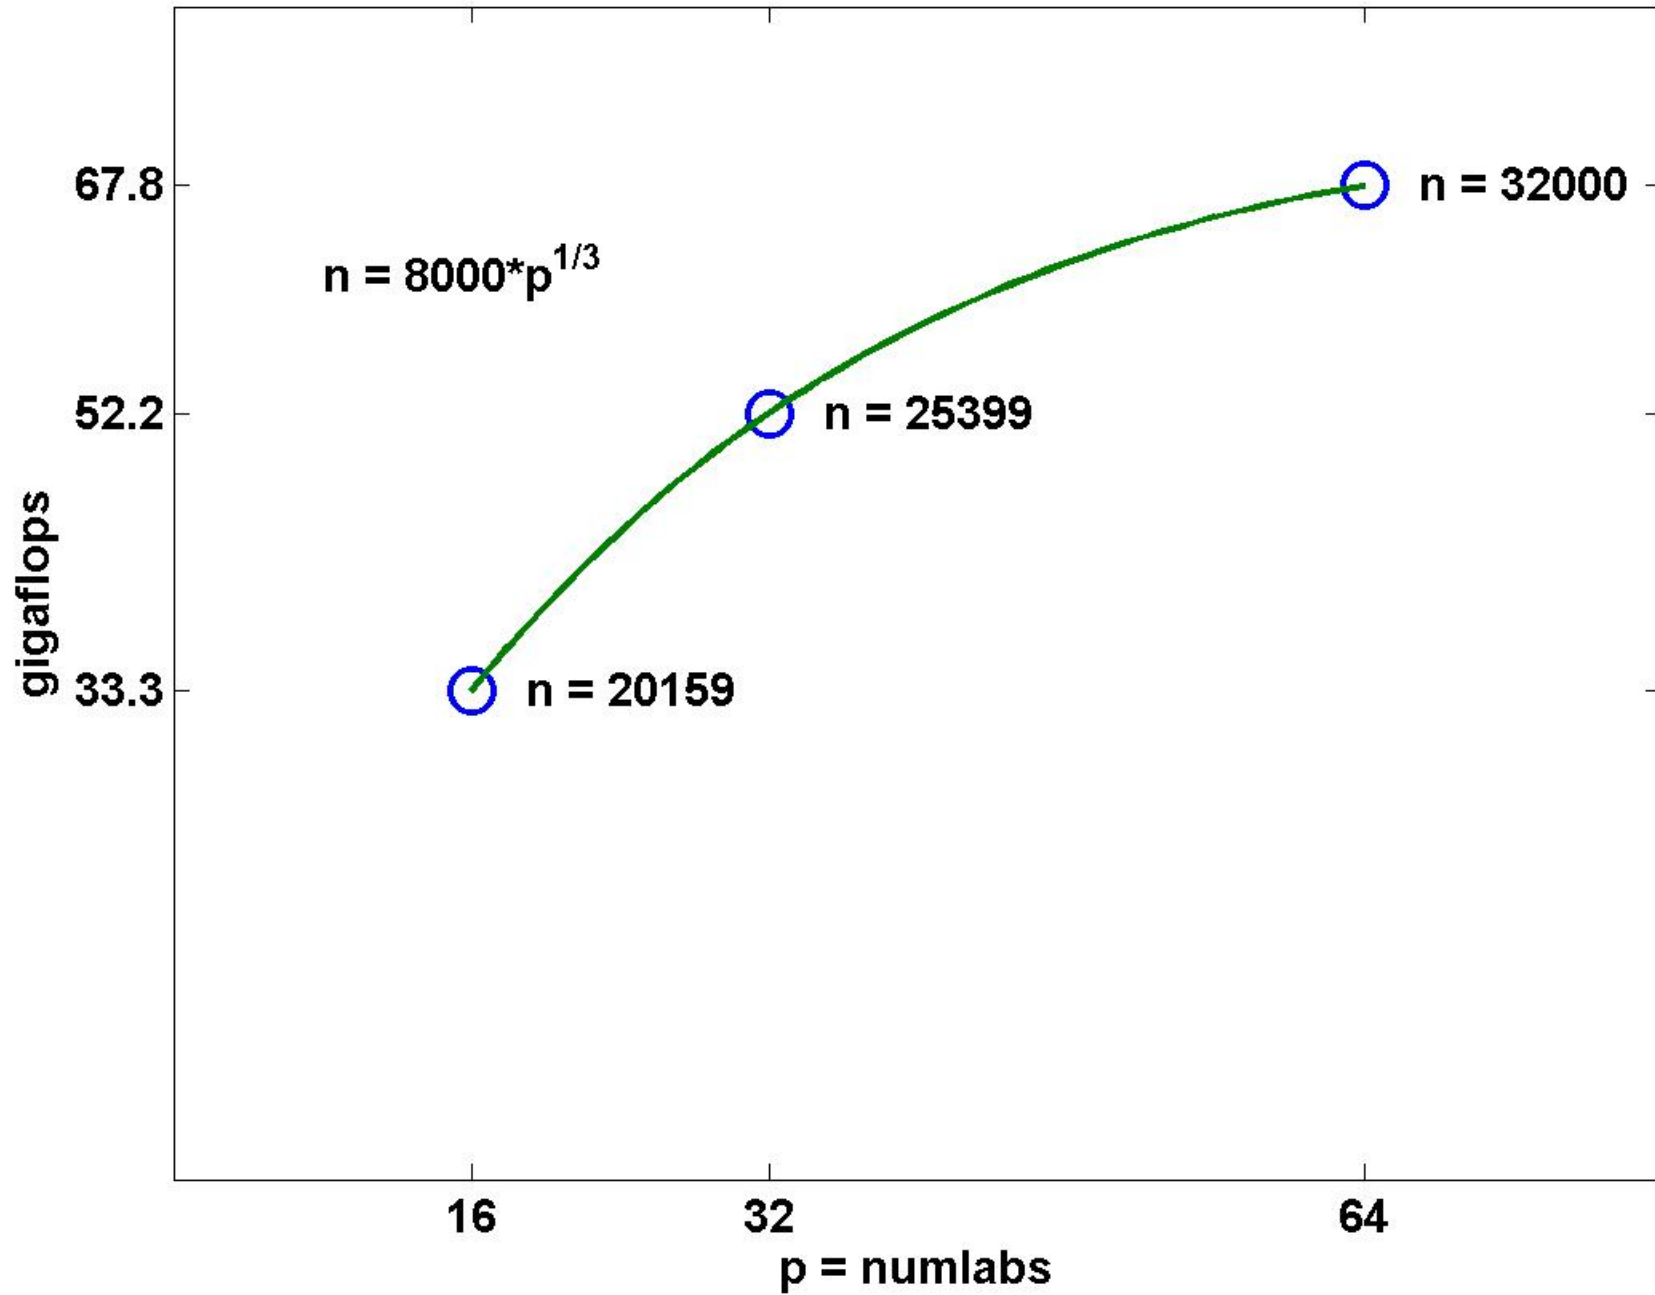
```
% HPC solve linear equations

n = ceil(4000*numl abs^(1/3));
A = randn(n, n, darray);
b = ones(n, 1);
tic
    x = A\b;
t = toc;
gf = (2/3*n^3 + 2*n^2)/t/1.e9;
r(1) = norm(A*x-b, inf)/(eps*norm(A, 1)*n)
r(2) = norm(A*x-b, inf)/(eps*norm(A, 1)*norm(x, 1))
r(3) = norm(A*x-b, inf)/(eps*norm(A, inf)*norm(x, inf)*n)
if labindex == 1;
    fprintf('%6d %10.4f %10.4f %14.4e\n', n, t, gf, max(r));
end
```

# lu(A)



# lu(A)



FFT

# FFT main program

```
function xfft(n)

% Create a complex row vector, distributed over the labs.
x = rand(1, n, darray()) + rand(1, n, darray())*i;

% Time the forward FFT
tic
y = fft1(x);
t = toc;

% Performance in gigaflops
perf = 5*n*log2(n)/t/1.e9;

% Error is scaled by roundoff level
z = ifft1(y);
err = norm(x-z, inf)/(log2(n)*eps);

fprintf('%8d %8d %8.3f %8.3f\n', numlabs, n, perf, err)
```

```

function x = fft1(x)
% FFT1 Distributed one-dimensional finite Fourier transform.
%   FFT1(x) where x is a distributed row vector
%           whose length is a multiple of numlabs.

n = length(x);
m = n/numlabs;
d = distribute(m(x);

% Reshape into two-dimensional array with numlab rows
x = reshape(x, numlabs, m);

% Transpose
x = redistribute(x.', d);

% Local one-dimensional FFT
xloc = fft(local(x));

% Twiddle factors
omega = exp(2*pi*i*(lindex-1)/n);
xloc = (omega.^(0:m-1)') .* xloc;
x = darray(xloc, d);

% Transpose again
x = redistribute(x.', d);

% More local one-dimensional FFTs
xloc = fft(local(x));

% Assemble final result
x = darray(xloc, d);
x = redistribute(x.', d);
x = reshape(x, 1, n);

```

FFT, length =  $2^n$

