

wed. 04/25/12:

HW: Report is half of the work and the score.

HW 1 often too short.

HW 2 too short

HW 3 nice attempt at reports, with some significant problems: If you change the algorithm itself, you need to discuss in detail!

→ Report text must explain what algorithm you implemented.  
Code listing is just backup.

HW 3 #1 (b) Printing from copyright page  
It is first edition,

05 04 03            14 13 12

might mean 11th printing from 02 = 2002.

(a) → Appendix A!

#3 (a) Ideas for improvement:

- improve output (output max variable, output format)

- implement command-line input

./traj a b n

- some error checking ( $n < 0$ ,  $b \geq a$ ,  $\text{argc} < 4$ ,  
 $n$  must be int,  $n \% np \neq 0$ )

- implement MPI\_Wtime and outputs

- go to double-precision

- implement load-balancing if  $n \% np \neq 0$

(c) Run for several  $n$ ! Ex.  $10^7$ ,  $10^8$ ,  $10^9$

Comment on how close to optimal speedup you are.

Describe the table!

## Ch. 5 MPI\_Gather, etc.

Last time: Traditional C for matrices = multi-dimensional arrays does not work with MPI commands, since we cannot be sure that the data is consecutive in memory.

Therefore, let's program in 1-D arrays only, and use column-oriented storage to also be compatible with Fortran and Matlab for instance

Fortran: BLAS = basic linear algebra subroutines  
LAPACK and other packages, that are actually Fortran code.

Ex.:  $C = AB$  = matrix-matrix product

$C_{ij}$  = dot product of row of A with col of B

= BLAS 1 = level-1 BLAS

= dot product fact. has 1 for loop

$C = \sum$  of outer products of col. of A  
with row of B

= BLAS 2 has 2 for - loops

$C = AB$  = BLAS 3 = has 3 for loops

$n = 8192$  for dim. of all matrices

BLAS 1	BLAS 2	BLAS 3
11,000 sec.	1,100 sec.	100 sec.

↑  
Could be faster depending  
on loop ordering!

Key reason:

BLAS actually uses  
block algorithms!

with blocks that are  
optimal for the cache!

This should motivate us to use  
a data structure that is compatible.

double \*l\_A /\* l\_A ∈ ℝ<sup>n × l\_n</sup> \*/

l\_A = (double \*) malloc((n \* l\_n) \* sizeof(double))

$$\text{Ex.: } A_{ij} = \frac{1}{i+j+1} \quad 0 \leq i, j < n$$

Hilbert matrix

for (j = 0; j < n; j++)

for (i = 0; i < n; i++)

$$A[i + n * j] = 1.0 / ((double)(i + j + 1))$$

The  $i$  index is in the inner loop because in that case the memory of  $A$  is accessed consecutively.

Since the code will load several elements of  $A$  into cache at the same time.

Ex. MPI\_Scatter(A, n \* l\_n, MPI\_DOUBLE,

l\_n = n / mp, l\_A, n \* l\_n, MPI\_DOUBLE,

A = [ l\_A | l\_A | ... | l\_A ]<sup>0, MPI\_COMM\_WORLD</sup>

HW 4 : Some ideas :

#1 (a) MPI\_Gather from  $l_A$  to  $A \rightarrow$  above

(b) dot product  $\rightarrow$  last time

(c) double Enorm (double \*  $l_x, \dots$ )

return sqrt (parallel\_dot( $l_x, l_x, -$ ))

$\rightarrow$  include math.h and link -lm

(d)  $y = Ax$  for  $x, y \in \mathbb{R}^n$ ,  $A \in \mathbb{R}^{n \times n}$

and requirement that  $A$  is given

as  $l_A$ ,  $x$  as  $l_x$ , and output must be  $y$  as  $l_y$

$$y = \begin{bmatrix} l_y \\ l_y \\ l_y \end{bmatrix} = Ax = \begin{bmatrix} l_A & l_A & l_A \end{bmatrix} \begin{bmatrix} l_x \\ l_x \\ l_x \end{bmatrix}$$