

Tu. 05/08/12:

HW 4 — some comments on report:

For tables and figures, you should think of writing in 4 steps:

- Introduce each: "Table 3.1 shows the results of a serial run of ---."
- Define the contents: "It shows timing data in seconds for 1, 2, 4, 8 nodes using 1, 2, 4, or 8 processes per node."
- Describe the contents: "As we double the number of processes, the times"
- Draw conclusions or make comparison: "This shows the scalability of the algorithm/code/implementation/hardware."

Memory prediction:

You can focus on large arrays immediately, here only $L \cdot A \in \mathbb{R}^{n \times 2 \cdot n}$. But also, can just do it for

1-process code! Reason: We want to use this to determine which n to use \Rightarrow So, need to know if it fits on 1 node! If it fits there, it will fit on any number of nodes.

We want $n = 2^v$. Ex.: $n = 8192 = 2^{13}$

$$\begin{aligned} \text{memory for } A &= 8 \text{ B/double} * n^2 \text{ double} = 8n^2 \text{ B} \\ &= 8 \cdot (2^{13})^2 \text{ B} = 2^3 2^{26} = 2^{29} \text{ B} = 2^{19} \text{ kB} = 2^9 \text{ MB} \\ &= 0.5 \cdot 2^{10} \text{ MB} = 0.5 \text{ GB} \end{aligned}$$

$$N = 8192 \Rightarrow M = 0.5 \text{ GB}$$

$$N = 16 \text{ k} \Rightarrow M = 2 \text{ GB}$$

$$N = 32 \text{ k} \Rightarrow M = 8 \text{ GB}$$

HW5: Specify the hardware used and software versions.
→ Put this at end of Sec. 1.

Starting point for HW6 is HW5 =
serial C code for $Ax()$ — all other parts
of the code should already be parallel!

Step 0: Restate serial code, but with l_u instead of u , l_v for v , and l_j for j
 (For debugging, have no " j " defined, so that you can be sure that all j are now l_j .)

```

for (l_j = 0; l_j < l_N; l_j++)
    for (i = 0; i < N; i++) {
        tmp = 4.0 * l_u [i + N * l_j]
        if (l_j > 0) tmp -= u [i + N * (l_j - 1)]
        ;
        l_v [i + N * l_j] = tmp
    }
    
```

if this is run on 1 proces, it will work just like the serial code.

But if we want to run on $p > 1$ proceses, then "if ($l_j > 0$)" needs work.

$l_j == 0$ is only a subdomain boundary, not necessarily on $\partial\Omega$. \Rightarrow For proceses $id > 0$, we need to have an "else" case, since it is actually

in interior of Ω and u_{ij-1} makes sense

If (ij) is on proces id and we have $l_j == 0$, then where is u_{ij-1} ? It is on proces $id-1$. \Rightarrow Need to receive it from there!

Step 1 - some confusing code

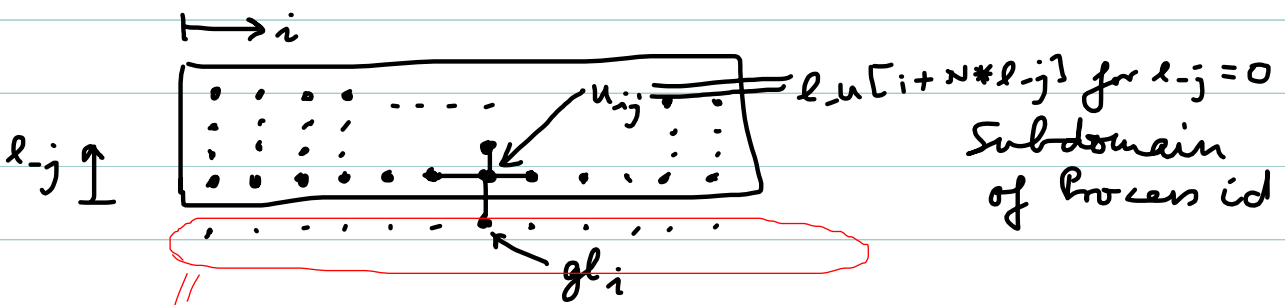
```

for ( l-j = 0, ..... )
  for ( i = 0; ..... ) {
    tmp = 4.0 * l_u [ i + N * l-j ]
    if ( l-j > 0 )
      tmp -= l_u [ i + N * ( l-j - 1 ) ]
    else { /* l-j == 0 */
      j = l-j + id * l_N
      if ( j > 0 )
        tmp -= gl [ i ]
      else /* on  $\partial\Omega$  */
        /* do nothing since u = 0 */
    }
  }

```

Notice that we need to store the data from Process $id-1$ = id_{left} somewhere.

Terminology of "ghost points"



vector of $gl \in \mathbb{R}^N$ on Process id =
 = one portion of length N of l_u on Process $id-1$

Notice that our data structure for l_u is exactly, so that these N values are consecutive in l_u namely the last N values on Process $id-1$

(\Rightarrow) So, each process needs to send these last N values of l_u to Process $id_{right} = id + 1$

\Rightarrow MPI_Send (&(l_u[N*(l_N-1)]), N, MPI_DOUBLE,

Code above has too complicated of an if-statement

Also, we want to start preparing for doing useful work (= calculations) while waiting for communications \Rightarrow want to clearly identify those portions of code that do or do not depend on any variable that needs to be received.

Step 2a - rough outline :

```
for (l_j = 1; l_j < l_N - 1; ... )
    for (i = 0; ...
        tmp -= l_u[i + N*(l_j - 1)]
```

```
l_j = 0
for (i = 0; ... )
    j = l_j + id * l_N
    if (j > 0) tmp -= gl[i]
```

```
l_j = l_N - 1
for (i = 0; ... )
    j = l_j + id * l_N - 1
    if (l_i < l_N - 1) tmp -= l_u[i + N*(l_j + 1)]
```

Step 2 - write clean code (w/ communications)

= with \downarrow tag

```
MPI_Recv ( gl, N, MPI_DOUBLE, idleft, 1, ... ) ←  
MPI_Recv ( gr, N, MPI_DOUBLE, idright, 2, ... ) ←  
MPI_Send ( &(l_u[N*(l_N-1)]), N, DOUBLE, idright, 1 ←  
MPI_Send ( l_u, N, MPI_DOUBLE, idleft, 2, ... ) ←  
    ||  
    &(l_u[0])
```

```
for ( l_j = 1; l_j < l_N - 1; ... )
```

```
    for ( i = 0; ...
```

```
        j = l_j + id * l_N
```

```
        tmp = 4.0 * l_u[i + N * l_j]
```

```
        tmp -= l_u[i + N * (l_j - 1)]
```

```
        if (i > 0) tmp -= l_u[i - 1 + N * l_j]
```

```
        if (i < N) tmp -= l_u[i + 1 + N * l_j]
```

```
        tmp -= l_u[i + N * (l_j + 1)]
```

```
        l_v[i + N * l_j] = tmp
```

```
    }
```

```
    l_j = 0
```

```
    for ( i = 0; i < N; i++ ) {
```

```
        j = l_j + id * l_N
```

```
        tmp = 4.0 * l_u[i + N * l_j]
```

```
        if (j > 0) tmp -= gl[i]
```

```
        : as above
```

```
    l_j = l_N - 1
```