

---

Exercise 4: distributed memory parallelization using MPI

---

Deadline: Jan 30, 2024, 16:00

Status:

Wednesday 24<sup>th</sup> January, 2024, 16:44

Supervisor: Prof.Dr. G. Haase,

`gundolf.haase@uni-graz.at`

---

## 1 Preliminaries

### 1.1 MPI

The **M**essage **P**assing **I**nterface has been introduced in the early 90-ties (i.e., after Rocky V) and it is still the standard environment for distributed parallel computing. It covers about 140 functions, available in F77, C and C++. Already 6 functions allow to write parallel codes. Most of the other functions are based on these 6. We will be mainly concerned with the following functions.

Basic functions	<b>MPI_Init</b> <b>MPI_Finalize</b> <b>MPI_Send</b> <b>MPI_Recv</b> <b>MPI_Comm_rank</b> <b>MPI_Comm_size</b>
additional functions	<b>MPI_Barrier</b> <b>MPI_Bcast</b> <b>MPI_Gather</b> <b>MPI_Scatter</b> <b>MPI_Reduce</b> <b>MPI_Allreduce</b>

The current standard is MPI-3.1 although this basic course will use only a small subset of its functionality.

### 1.2 Online help

First of all the MPI Homepage and especially the overview of the MPI functions should be consulted. We will refer frequently to these web pages during the course.

The description for MPI-functions in C and Fortran (nad Fortran 2008) can be found **here**. The c++ bindings were deprecated.

We have to distinguish between the MPI standard and its implementations. The most commonly used implementations are MPICH and OpenMPI (that is not OpenMP !!). All three are available as packages under LINUX but only one of them should be used in order to avoid confusion wrt. paths to executables, libraries and headers. We will refer to the latter one, see the man pages of OpenMPI.

### 1.3 Getting started on a (pool of) LINUX-workstations/PCs

First, open a shell and type

```
mpirun
```

If MPI is not available then you have to install additional packages (in Ubuntu) via

```
sudo apt-get install openmpi-bin openmpi-doc libopenmpi-dev
```

or install it from the scratch (just for fun).

Check whether the ssh-deamon is running

```
ps -ax |grep sshd
```

If not you have to install it too.

In order to avoid the password request for each parallel process started (think of 64 parallel processes) you have to create secure authentication keys for your account.

```
ssh-keygen
cp ~/.ssh/id_rsa.pub ~/.ssh/authorized_keys
chmod go-rwx ~/.ssh/authorized_keys
# locally
ssh-add
# remote
ssh-copy-id your_username@192.0.2.0
```

See also my hints.

## 1.4 Installing the example code

Copy and unpack the provided supporting material for the Jacobi template<sup>1</sup> into a folder and unpack it.

---

<sup>1</sup><http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/jacobi.template.zip>

## 2 Your first parallel code

- E1 Compile the program in `template`<sup>2</sup> and set in the *Makefile* the variable `COMPILER` to `COMPILER=GCC_` . Adapt *Makefile* and *../GCC-default.mk* to your needs and paths. (1 pts)  
Compile and link the code

```
make
```

Start the program with 4 processes

```
make run
```

or directly via

```
mpirun -np 4 ./main.GCC_ (ensure that you use the right mpirun)
```

If you 'mpirun ...' report some error "... not enough slots .." then use the option '-oversubscribe', i.e., `mpirun --oversubscribe -np 4 ./main.GCC_`

The following MPI functions require a *communicator* as parameter. This communicator describes the *group* of processes which are to be covered by the corresponding MPI function. By default, all processes are collected in `MPI_COMM_WORLD` which is one of the constants supplied by MPI for predefined datatypes, error classes, collective operations etc.. We restrict the examples to those global operations. For this purpose, create special MPI-type variable `MPI_Comm icom = MPI_COMM_WORLD`; which is used as parameter !

- E2 Write Your first parallel program by implementing (1 pts)  
**MPI\_Init** and **MPI\_Finalize**

compile the program and start 4 processes

```
mpirun -np 4 ./main.GCC_
```

- E3 Determine the number of your parallel processes and the local process rank by using (1 pts)  
the routines

**MPI\_Comm\_rank** and **MPI\_Comm\_size** .

Let the root process (0==rank) write the number of running processes. Start with different numbers of processes.

- E4 The file *greetings.cpp*<sup>3</sup> includes a routine (1 pts)  
**Greetings(MPI\_Comm const &icom)**

that prints the names of the hosts your processes are running on. Call that routine from your main program and change the routine such that the output is ordered with respect to the process ranks. Study the routines

**MPI\_Send** and **MPI\_Recv**

with respect to *tags* and *ranks*.

---

<sup>2</sup><http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/jacobi.template.zip>

<sup>3</sup>[http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/first.template/html/greetings\\_8cpp.html](http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/first.template/html/greetings_8cpp.html)

### 3 Global Operations

Let some Double Precision vector  $\underline{x}$  be stored blockwise disjoint, i.e., distributed over all processes  $s$  ( $s=0,\dots,P-1$ ) such that  $\underline{u} = (\underline{x}_0^T, \dots, \underline{x}_{P-1}^T)^T$ . See suggestions of function interfaces<sup>4</sup>.

- E5** Write a routine (1 pts)

**DebugVector(xin, icommm)**

that prints the Double Precision vector  $\underline{x}$ . Start the program with several processes.  
 $\implies$  All processes will write their local vectors, i.e., one has to look carefully for the data of process  $s$ .

Improve the routine **DebugVector** such that process 0 reads the number (from terminal) of that process which should write its data next. Use

**MPI\_Bcast**

to broadcast this information and let the processes react appropriately. If necessary use **MPI\_Barrier** to synchronize the output.

- E6** Write a routine for computing the global scalar product (1 pts)

**par\_scalar(x, y, icommm)**

of two Double Precision vectors  $x$  and  $y$  of local length  $n$ . Use

**MPI\_Allreduce** with the operation **MPI\_SUM**.

- E7** Determine **and exchange** global minimum and maximum of the vector  $\underline{x}$  ! Use (1 pts)

**MPI\_Gather** , **MPI\_Scatter** , **MPI\_Bcast** and **MPI\_Sendrecv** .

How can you reduce the amount communication?

Hint: Compute, first, local min./max. and afterwards let some process determine the global quantities.

Alternatively, you can use **MPI\_Allreduce** and the operations **MPI\_Minloc**/**MPI\_Maxloc**.

- E8** Each of your 4 MPI processes ( $\text{numprocs} := 4$ ) owns a double precision vector  $x$  with 20 elements initialized as (1 pts)

$$x[i] := \text{myrank} * 100 + (i \% 5) * 10 + i \text{ .}$$

You have to exchange vector elements with MPI ranks  $r \in [0, 3]$  . All elements with index  $i \in [r * 5, (r + 1) * 5 - 1]$  have to be exchanged with rank  $r$ .

Check the resulting vector.

You could use a bunch of **MPI\_Sendrecv** calls but you better use the more general function **MPI\_Alltoall** .

Try also the In-place option by using **MPI\_IN\_PLACE** with **MPI\_Alltoall** .

Have a look at the more general function **MPI\_Alltoallv** .

---

<sup>4</sup>[http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/glob/html/globals\\_func.html](http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/glob/html/globals_func.html)

## 4 Data exchange

Download the template<sup>5</sup> containing the functions for the setting up the geometry and the exchange of data between MPI processes.

For a start we use Matlab to describe manually the decomposition of a square into sub-squares, see script square\_4.m<sup>6</sup>. The resulting list of finite elements together with the node coordinates is stored in square\_4.txt<sup>7</sup>. The assignment of elements to a subdomain (= MPI rank)  $r = 0, \dots, 3$  is stored in file square\_4\_sd.txt<sup>8</sup>.

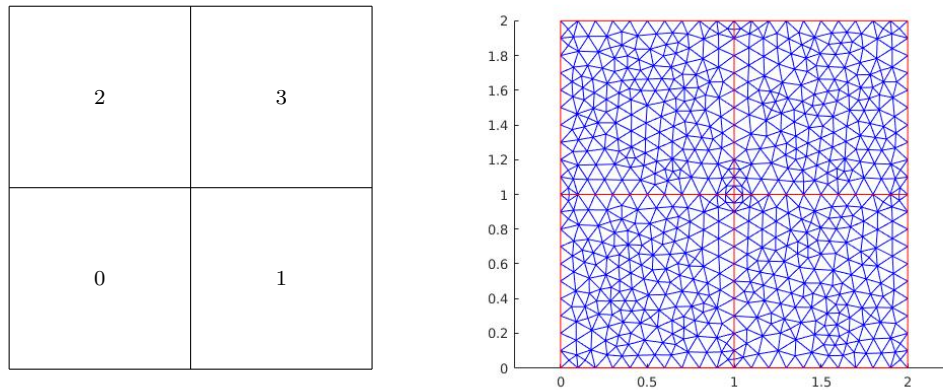


Figure 1: MPI rank per subdomain (left) and discretization (right).

Please note that both text files contain the original Matlab numbering starting with 1 which will be changed automatically to C-numbering (starting with 0) in the C++-constructor of the mesh. The class ParMesh reads these files with its constructor

**ParMesh const mesh("square",icomm);**

for the provided files (*only for 4 MPI processes!*) and initializes also the MPI parallelization for data exchange on the interfaces of the subdomains.

Object **mesh** contains basic communication routines and information as

- Number of MPI processes: `int const numprocs = mesh.NumProcs();`
- MPI rank of this process: `int const myrank = mesh.MyRank();`
- inner product: `double ss = mesh.dscapr(x1,x1);`
- Vector accumulation on interfaces: `mesh.VecAccu(x1);`

<sup>5</sup><http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/accu.template.zip>

<sup>6</sup>[http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/accu.template/html/square\\_4m\\_source.html](http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/accu.template/html/square_4m_source.html)

<sup>7</sup>[http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/accu.template/square\\_4.txt](http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/accu.template/square_4.txt)

<sup>8</sup>[http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/accu.template/square\\_4\\_sd.txt](http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/accu.template/square_4_sd.txt)

- E9 Check the vector accumulation `mesh.VecAccu(...)` with your own vector taking into account that you know the interface coordinates. The coordinates can be accessed via `mesh.GetCoords();` . (2 pts)  
See also Figure 2 for a scheme of the available local node numbering.  
Check also the inner product `mesh.dscapr(...)` .
- E10 Write a Method `VecAccu` for class `Parmesh` that adds interface data of an *integer* vector. (2 pts)
- E11 Write a Method `GlobalNodes()` for class `Parmesh` that determines the global number of nodes in the mesh. (2 pts)
- E12 Write a Method `Average` for class `Parmesh` that averages (arithmetic mean) interface data of a double precision vector instead of adding them. (2 pts)
- E13 Generate other domains for parallel data exchange by copying `square_4.m` to a different name and describe your subdomains in a similar way. (2 pts)  
Try also more/less than 4 subdomains and start your code appropriately, see target `run` in *Makefile*.




Figure 2: 4 subdomains in local numbering.

## 5 Iterative Solvers

Download the Jacobi template<sup>9</sup> containing the functions for matrix generation and sequential solvers or use your own version from the shared memory part of the course.

As model problem, we consider the homogeneous Dirichlet boundary value problem ( $\mathbf{u}(x) = 0 \quad \forall x \in \partial\Omega$ ) for the Poisson equation in the unit square  $\Omega := (0, 1)^2$  in its weak formulation:

Find  $\mathbf{u} \in \mathbf{H}_0^1(\Omega)$  such that

$$\int_{\Omega} \nabla^T \mathbf{u}(x) \nabla \mathbf{v}(x) dx = \int_{\Omega} \mathbf{f}(x) \mathbf{v}(x) dx \quad \forall \mathbf{v} \in \mathbf{H}_0^1(\Omega) . \quad (1)$$

We use linear finite elements for the discretization and achieve the linear system of equations

$$K \cdot \underline{u} = \underline{f} . \quad (2)$$

### 5.1 Jacobi solver

Let us denote the diagonal of matrix  $K$  by  $D = \text{diag}(K)$ . Now, we can formulate the  $\omega$ -Jacobi iteration

$$\underline{u}^{k+1} = \underline{u}^k + \omega \cdot D^{-1} \cdot (\underline{f} - K \cdot \underline{u}^k) \quad , \quad k = 0, 1, 2, \dots . \quad (3)$$

You will find a sequential version of the Jacobi solver in the directory *jacobi.template* extended with the parallel mesh and its functions provided already in the previous section 4.

Once you have an object `mesh` of class `Parmesh` the following matrix generation steps require **no communication**:

- Allocation and patten determination of local sparse matrix:  
`FEM_Matrix SK(mesh);`
- Calculation of matrix entries (for Laplacian):  
`SK.CalculateLaplace(fv);`
- Setting values for local vectors, depending e.g. on node coordinates:  
`mesh.SetValues(...)`
- Incorporating boundary conditions (here, for Dirichlet):  
`SK.ApplyDirichletBC(...)`

The only function remaining for MPI adaption is `JacobiSolve(SK, fv, uv );`

**E14** Implement an MPI parallel version of the sequential code in `JacobiSolve`<sup>10</sup> with the interface (4 pts)

```
void JacobiSolve(Parmesh const & mesh, CRS_Matrix const &SK,
                vector<double> const &f, vector<double> &u)
```

which requires the use of `mesh.dscapr(...)` and `mesh.VecAccu(...)`. The latter one is needed also in accumulation of the matrix diagonale.

<sup>9</sup><http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/jacobi.template.zip>

<sup>10</sup>[http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/jacobi.template/html/jacsolve\\_8h.html#a4b66d6c4de6f3c7b26bbf39d0f7054f0](http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/par/jacobi.template/html/jacsolve_8h.html#a4b66d6c4de6f3c7b26bbf39d0f7054f0)

\_\_\_\_\_This document will be extended by further advices, links, etc. \_\_\_\_\_