

# Programming Multi-Core Systems with OpenMP

## Programming Scalable Systems with MPI

### 1 Getting started with OpenMP

Of course, your first program should be *Hello World*, but as we are using OpenMP we want each thread to say *hello*. Write a simple program that starts a parallel section where each thread says hello and prints out his/her thread number as well as the total number of threads. You can adapt the code shown during the lecture, in particular if you are not a C expert.

Compile the code as shown during the lecture and run it with different numbers of threads on your own computer as well as on the SURFsara Lisa cluster. Do not forget to activate the compiler's OpenMP support and to switch on compiler optimizations of at least level `-O2`. Use the same compiler options throughout the entire assignment.

For running code on a cluster with batch job submission system, as the Lisa cluster, you need to explicitly set the number of threads in the code using the `omp_set_num_threads()` routine.

**Directions:** Report on your code and on the experiments.

### 2 Accelerating computations with OpenMP

Adapt the code that computes the element-wise product of two vectors as shown during the lecture. Run the compiled code for different vector lengths as well as for different numbers of repetitions of the element-wise vector product.

What speedups over sequential execution of the same code can you achieve on your own computer as well as on a single node of the SURFsara Lisa cluster?

**Directions:** Report on your code and on the experiments.

### 3 Parallelizing reduction operations with OpenMP

Extend your code developed for the previous assignment to compute the scalar product of two vectors instead of the element-wise product. The scalar product is the sum of all element-wise products. Compute the scalar product with a single loop and avoid an intermediate vector to store the element-wise products.

Devise three parallel implementations of your scalar product code:

1. using a critical section,
2. using the reduction clause,
3. using your own reduction scheme that neither uses a critical section nor the reduction clause.

**Directions:** Compare all three parallel implementations in terms of speedup over sequential execution for different vector sizes and different numbers of repetitions as in the previous assignment.

---

## 4 Overlapping communication and computation in MPI

During the lecture you learned that overlapping (costly) communication with (productive) computation is crucial in message passing programming to achieve good performance.

**Exercise 1** Improve the time iteration loop of the wave equation, as shown on slide 58, such that it even better overlaps communication and computation.

**Exercise 2** Devise a variant of the time iteration loop of the wave equation that exploits the split-phase communication features of MPI for improved performance.

**Directions:** Use pseudo code in analogy to the lecture slides, explain your solutions in detail and argue how the goal of overlapping communication with computation while retaining functional correctness is achieved.

## 5 Collective communication in MPI

Collective communication is a form of structured communication, where instead of a dedicated sender and a dedicated receiver all MPI processes of a given communicator participate. A simple example is *broadcast* communication where a given MPI process sends a message to all other MPI processes in the communicator. Write a C+MPI function

```
int MYMPI_Bcast( void *buffer,          /* INOUT : buffer address */
                 int count,           /* IN    : buffer size */
                 MPI_Datatype datatype, /* IN    : datatype of entry */
                 int root,            /* IN    : root process (sender) */
                 MPI_Comm communicator) /* IN    : communicator */
```

that implements broadcast communication by means of point-to-point communication. Each MPI process of the given communicator is assumed to execute the broadcast function with the same message parameters and root process argument. The root process is supposed to send the content of its own buffer to all other processes of the communicator. Any non-root process uses the given buffer for receiving the corresponding data from the root process.

**Exercise 1** Implement the above function MYMPI\_Bcast using point-to-point communication. Assume a fully-connected network topology, ie every node may send messages to any other node at the same cost.

**Exercise 2** A particular advantage of collective communication operations is that their implementation can take advantage of the underlying physical network topology without making an MPI-based application program specific to any such topology. For your implementation of broadcast assume a 1-dimensional ring topology, where each node only has communication links with its two direct neighbours with (circularly) increasing and decreasing MPI ranks. While any MPI process may, nonetheless, send messages to any other MPI process, messages may need to be routed through a number of intermediate nodes. Communication cost can be assumed to be linear in the number of nodes involved. Aim for an efficient implementation of your function on an (imaginary) ring network topology.

**Directions:** Write real C+MPI code for this assignment, and make sure that it compiles. No experiments are required, but describe your code in detail: how it is adapted to the given network topology and how many atomic communication events (sending a message from one node to a neighbouring node) are required to complete one broadcast.

**Submission deadline: January 29, 2016**