


Parallel Numerics

Exercise 3: Vector–Vector Operations & P2P Communication II

1) Data Dependency Graphs & DAGs



”A *directed acyclic graph* (DAG) is a directed graph that has no positive cycles, that is, no cycles consisting exclusively of forwards arcs. A DAG can be used to represent a parallel algorithm, as we proceed to show. [...] Each node represents an operation performed by an algorithm, and the arcs are used to represent data dependencies. In particular, an arc $(i, j) \in A$ indicates that the operation corresponding to node j uses the results of the operation corresponding to node i . An operation could be elementary (e.g., an arithmetic or a binary Boolean operation [...]), or it could be a high-level operation like the execution of a subroutine.”

Bertsekas, Tsitsiklis: Parallel and Distributed Computation, Prentice–Hall International Editions, 1989

- i) What is the difference between the data dependency graph and the direct acyclic graphs used to illustrate the program execution? How can one identify an iterative algorithm if only the data dependency graph is given?
- ii) We want to sum up 2^k values. Give a scetch/evaluation scheme for this problem and a sequential algorithm (pseudocode). What is the minimal number of execution steps to be done? Assume that the cpu supports only unary and binary operations.

- iii) Give an optimal parallel algorithm (not in pseudocode, just thoughts) for the problem from exercise *ii*) using a butterfly scheme (see the illustration above). Hereby, in every step node 0 sums up the first two entries, node 1 sums up the values three and four and so on. What is the minimal number of time steps required? What is the maximal number of processors one can use for such a problem? What happens if more processors are used?

2) BLAS — Basic Linear Algebra Subroutines

Make yourself familiar with the different BLAS levels. Set up a table having the columns blas level, description, complexity and example. Why are BLAS routines preinstalled on every computer and, usually, not linked directly (static) into the programs using them?

blas level	description	complexity	example
1	$\alpha \cdot a$ $\alpha \in \mathbb{R}, a \in \mathbb{R}^d$ $\alpha \cdot \beta$ $\beta \in \mathbb{R}$
2
...			

3) P2P Communication II: Blocking vs Non-blocking

- i) In MPI many descriptions use the term *buffer*. Make yourself familiar with this term and explain the difference between the following types of buffers: Program variables, user-level buffer (assigned to MPI) and hardware/system buffer.
- ii) MPI offers a blocking and a non-blocking send and receive command. What does this mean with respect to a variable used as buffer. Make yourself familiar with the semantics of the blocking and non-blocking P2P operations.
- iii) What is the advantage of combining blocking and non-blocking operations?

4) Non-blocking Operations in Action

The following question is taken from the exam of WT 2005/06. Given is the following source code running on three processors:

```
...
double    a,b;
int       myRank,noOfProcessors;
MPI_Status status;

MPI_Init(...);
MPI_Comm_size( MPI_COMM_WORLD, &noOfProcessors );
MPI_Comm_rank( MPI_COMM_WORLD, &myRank );

a = myRank+1;

if (myRank==0) {
    for (int i=1; i<noOfProcessors; i++) {
        MPI_Recv(&b,1,MPI_DOUBLE,i,0,MPI_COMM_WORLD,&status);
        a+=b;
    }
}
else {
    MPI_Send(&a,1,MPI_DOUBLE,0,0,MPI_COMM_WORLD);
}
...
```

- i) What value is stored within the variable `a` on every processor after the source code extract has finished?
- ii) Give a technical term for the type of communication implemented by this source code fragment.
- iii) What may happen if one uses non-blocking communication within the source code not adding additional statements?

5) Vector–Vector Operations

In Scientific Computing vector-vector operations are frequently occurring operations. The dimension n of the vectors is often very large. Hence, parallelisation of the basic linear algebra operations can be advantageous.

In this exercise we will parallelise the adding and the dot-product of two vectors \vec{a} and \vec{b} . The result is a vector \vec{r} or a number z , respectively:

$$\vec{r} = \vec{a} + \vec{b}, \quad z = \vec{a} \cdot \vec{b} = \sum_{i=1}^n a_i b_i$$

Write a parallel program for adding two vectors as follows:

- process 0 allocates three arrays **a**, **b**, **r** with dimension **LENGTH** using `(float*)malloc(LENGTH*sizeof(float))`
- process 0 stores random numbers in the arrays **a**, **b** using the function `rand()`
- partition these vectors in p equal parts (p is the number of started processes), and calculate the adding of the parts by the p processes separately (process 0 has to send the parts to the other processes respectively)
- process 0 collects the parts and stores the result in array **r**

What is the speedup for calculating the sum of two vectors with **LENGTH=10.000.000** with 2,4,8 processes? Use the function `MPI_Wtime()` to determine the time spent in the loop for calculating the sum.

Write a similar program to calculate the dot-product. What is the speedup? In this program process 0 has to add p numbers in the end. How can you parallelize this addition? Is it advantageous?