

Parallel Numerics

Exercise 4: Matrix-Matrix Operations & P2P Communication III

1 DAGs for Function Evaluation

Given is $f(x_1, x_2, x_3, x_4) = x_1^2 + x_2x_3 - x_1x_3 - x_1x_2 = x_2x_3 + x_1(x_1 - x_3 - x_2)$.

Give the DAGs for both evaluation schemes in parallel and compare the number of required time steps.

2 Amdahl's Law

Given is an algorithm working on a fixed amount of data. f is the fraction of a calculation that is sequential and, thus, cannot benefit from parallelisation. The algorithm has execution time t_1 .

- What is the execution time t_p , if exactly this problem is ported to a parallel machine with p processors.
- What is the speedup?
- What is the efficiency?
- What happens if the number of processors is increased ($\lim_{p \rightarrow \infty}$)?

3 Gustafson's Law (Gustafson-Barsis' Law)

You are given a parallel algorithm running on p nodes that needs the time t_p to finish. Again, f is the fraction of a calculation that is sequential and, thus, cannot benefit from parallelisation.

- Derive the time t_1 , one processor would need to solve this problem.
- What is the speedup?
- What is the efficiency?
- What happens, if the number of processors is increased ($\lim_{p \rightarrow \infty}$)?

- e) Compare these formulas to Amdahl's law. In today's problems, often not the execution but the pure amount of data is the limiting factor for computations. Thus, an algorithm is parallelised because one wants to compute bigger problems, and a speedup is only a (nice) side-effect. What speedup law is the appropriate one for this scenario?

4 P2P Communication III: Synchronised Operations

Given is the following pseudocode fragment:

```
1     # node 0:
2     int a = 0;
3     send a to node 1
4     a = 1;
5     send a to node 1
6
7     #node 1:
8     int b;
9     receive b from node 0
10    printf("%i", b);
11    receive b from node 0
12    printf("%i", b);
```

- a) Implement the send and receive operations using blocking and non-blocking operations. What might be the difference in the results?
- b) Define the MPI term *synchronised*.
- c) Implement the send and receive operations using synchronised non-blocking operations. Insert a wait operation directly after the next integer assignment on node 0. What might be the difference in the results?
- d) Implement the send and receive operations using synchronised blocking operations. What might be the drawback?
- e) Make yourself familiar with the MPI terms *buffered send* and *ready send*.

5 Matrix-Matrix Multiplication

The product $A \cdot B$ of two $N \times N$ -matrices A and B is calculated as follows:

$$C = A \cdot B \quad \text{where} \quad (C)_{ij} = \sum_{k=1}^N (A)_{ik}(B)_{kj}.$$

A program that calculates this product has to use three independent loops (over i, j, k). The execution speed of the program depends on the arrangement of these loops.

- a) Give a sketch of the data dependency graph for one element of this problem.

- b) Assume N is a multiple of \sqrt{p} . All the three matrices are split up into p quadratic submatrices. Every node has to compute one submatrix of the result matrix C . Illustrate this fact!
- c) Pick out three different nodes. What parts of A and B are required by them to compute the result?
- d) Have a look at the node computing the upper left submatrix of C . Interpret the multiplication formula $(C)_{ij} = \sum_{k=1}^{\sqrt{p}} (A)_{ik}(B)_{kj}$ block-wise. Let's assume every summand of the result is computed in one time step. What different parts of A and B are required subsequently? How many time steps are required?
- e) Every node is allowed to communicate with the node that computes the left, right, top or bottom submatrix of C only. Assume this cartesian topology is cyclic. Furthermore, every node is allowed to hold only one submatrix of A and B at one time. Derive a communication scheme.