Parallel Programming

Prof. **Paolo Bientinesi**

pauldj@aices.rwth-aachen.de

WS 17/18

RWTH AACHEN UNIVERSITY

Proposed exercises (1/3): Collectives

Each of the p processes participating in the computation owns a local data buffer.
The length of the buffer local to process P_i is n_i, a random positive integer;
the buffer contains double-precision floating point numbers.

The file `HW3-exercise1.c` provides you with the skeleton of the computation.
For the computation to take place, every process needs the contents of the longest data buffer. If there exist two or more processes whose buffer is the longest, then the entire computation is restarted with new buffer lengths.

Fill out the blanks in `HW3-exercise1.c`.
Minimize the number of collectives used.

Proposed exercises (2/3): Point to Point

The p processes participating in the computation are logically ordered as a ring.
Process P_i, (with $i \in [0, \ldots, p-1]$), owns a local data buffer v_i containing n integers.
In order to make progress, P_i needs to compute

$$v_i^{\text{next}} := f(v_{i-1}^{\text{now}}) + f(v_i^{\text{now}}) - v_{i+1}^{\text{now}}.$$

The function $f : \mathbb{R}^n \to \mathbb{R}^n$ is expensive and does not overwrite its argument.
The buffers v_i^{now} and v_i^{next} are distinct.

Write a program that performs one step of the computation (from v_i^{now} to v_i^{next}), aiming to minimize the execution time. Add short comments explaining the ideas.

Proposed exercises (3/3)

1. Implement your own Scatter: Use the same algorithm as we did for the Broadcast (Minimum Spanning Tree), but the root only sends half of its buffer. Assume the buffer size to be a multiple of the number of processes.

2. Implement your own Scatter: Assume that the processes are organized as a nRows $\times$ nCols grid. First Scatter by rows, then Scatter by columns. Assume the buffer size to be a multiple of nRows and nCols.
   Compare the final result with the solution of the previous exercise. Are there differences? If so, why? How would you obtain the exact same answer?
   If not, why?

3. Implement a Broadcast of vector as a Scatter + (nProcs-1) steps of "pass to the right". Assume the vector size to be a multiple of the number of processes.