MPI – Part 2

1. In this task you will design and study the properties of an algorithm to compute a matrix-vector product $y = Ax$ built on a 1D distribution by columns of the matrix A. Process P_i owns the block of colums A_i, as well as the subvector x_i. The resulting vector y will be scattered among all processes, with process P_i being the owner of subvector y_i. Assuming 4 processes, the partitioned operation looks like:

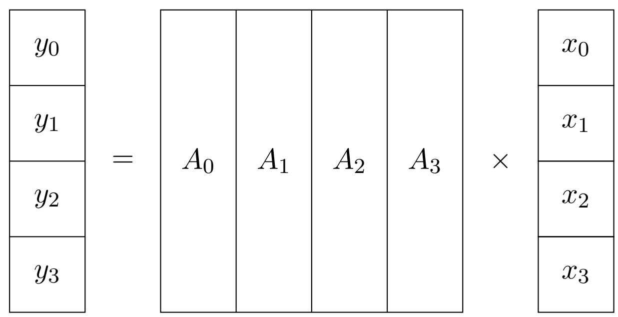

   a) Sketch the parallel algorithm as we did in class. Notice: An implementation is not requested.

      1) $t_i = A_i * x_i$ ($t_i \in R^n$: partial result)
      2) $y_i =$ `ReduceScatter(` t_i `)`

   b) Which collective communication operations does the algorithm utilize?

      – Reduce-Scatter

   c) Give the parallel cost for the algorithm, that is a lower bound for $T_p(n)$ taking into account both compute time and communication time.

      – $T_p(n) = log_2(p)\alpha + \frac{p-1}{p}n(\beta+\gamma) + 2\frac{n^2}{p}\gamma \approx log_2(p)\alpha + n(\beta+\gamma) + 2\frac{n^2}{p}\gamma$

   d) Study the strong and weak scalability properties of the algorithm as we did in class.

      – $T_1(n) = 2n^2\gamma$
      – $S_p(n) = \frac{T_1(n)}{T_p(n)} = \frac{2n^2\gamma}{log_2(p)\alpha + n(\beta+\gamma) + 2\frac{n^2}{p}\gamma}$
      – $E_p(n) = \frac{S_p(n)}{p} = \frac{2n^2\gamma}{plog_2(p)\alpha + np(\beta+\gamma) + 2n^2\gamma} = \frac{1}{1+\frac{plog_2(p)}{2n^2}\frac{\alpha}{\gamma}+\frac{p}{2n}\frac{(\beta+\gamma)}{\gamma}}$
      – Strong scalability: $\lim_{p\to\infty} E_p(n) = 0$
      – Weak scalability:
        M: local memory; Mp: combined memory;
        n_M: largest problem solvable with p processes $\to n_M^2 = Mp$
        $\lim_{p\to\infty} E_p(n_M) = \frac{1}{1+\frac{log_2(p)}{2M}\frac{\alpha}{\gamma}+\frac{\sqrt{p}}{2\sqrt{M}}\frac{(\beta+\gamma)}{\gamma}} = 0$