International Conference on Numerical Combustion
Rome, Italy, October 14-17, 2025
www.icnc2025.com

Boosting computational efficiency of reactive CFD simulations of multiphase heterogeneous reactors through Dynamic Load Balancing

Mauro Bracconi*[1], Daniele Micale[1], Matteo Maestri[1]
*lead presenter: mauro.bracconi@polimi.it
[1] Laboratory of Catalysis and Catalytic Processes, Politecnico di Milano, Italy

1. Introduction

This work presents a Dynamic Load Balancing (DLB) approach, based on a hybrid MPI/OpenMP parallelization, specifically conceived for detailed reactive Computational Fluid Dynamics simulations of fixed and fluidized reactors, critical technologies in the energy transition (e.g., biomass and waste pyrolysis, chemical looping, …) [1]. This is crucial to address a compelling issue that typically affects parallel computing when employed for the solution of these systems. Here, the distribution of the reactive cells in the computational domain is strongly inhomogeneous due to the geometry of the system in fixed bed reactors and evolving catalyst distribution in fluidized bed reactors. Moreover, the cost of chemical calculations is highly related to the local thermochemical state (e.g. temperature, partial pressures). As a result, the computational load is highly variable both in time and space. This often results in imbalanced conditions among the different subdomain which leads to highly inefficient calculations. The proposed approach, which is based on a hybrid MPI/OpenMP parallelization, improves the computational efficiency of multiscale CFD simulations paving the way for a more efficient exploitation of High-Performance Computing resources and expanding the current boundaries of feasible simulations.

2. Results and discussion

The computational performances of the combined DLB and hybrid parallelization approach have been investigated in fixed and fluidized bed reactor simulations using the catalyticFoam framework [2]. Here, we showcase the simulation of an industrial-scale fluidized bed reactor for Oxidative Methane Coupling. When the simulation is carried out without the DLB (dashed lines in Fig. 1a), the computational cost of the different processors is severely imbalanced leading to a severe computational imbalance. The DLB redistributes the computational burden among the processors leading to even computational cost among the processors (solid lines in Fig. 1a). As a result, the parallel efficiency rises from 19% without DLB (128 MPI) to 91% when DLB and hybrid parallelization (64 MPI x 2 threads) is employed (Fig. 1b). Overall, this results in a 2.1-fold reduction of the simulation time. This work proposes an effective strategy to improve the parallelization efficiency of multiscale CFD frameworks with a concomitant reduction of the computational cost paving the way to fundamentally investigate and design catalytic reactors of industrially relevant scales.

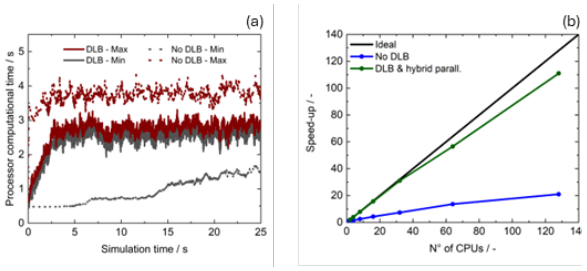

Figure 1. Computational cost of the chemical solution with and without DLB with 128 CPUs (a) and strong scalability analysis (b)

References

[1] M. Bracconi, Intensification of catalytic reactors: A synergic effort of Multiscale Modeling, Machine Learning and Additive Manufacturing. Chem. Eng. Process. - Process Intensif. 2022; 181:109148.
[2] D. Micale et al., Increasing Computational Efficiency of CFD Simulations of Reactive Flows at Catalyst Surfaces through Dynamic Load Balancing. ACS Eng. Au 2024;4(3):312–324