Increasing computational efficiency of reactive CFD simulations of fixed and fluidized bed reactors through Dynamic Load Balancing

Daniele Micale[1], Mauro Bracconi[1*], Matteo Maestri[1]

[1] *Laboratory of Catalysis and Catalytic Processes, Department of Energy, Politecnico di Milano, Via La Masa 34, 20156 Milano, Italy*

**mauro.bracconi@polimi.it*

Highlights

- First-principles-based multiscale models are needed for advanced reactor analysis
- DLB reduces computational cost in reactive CFD simulations of catalytic reactors
- DLB and hybrid MPI/OpenMP parallelization enhance computational efficiency

1. Introduction

Detailed multiscale models are nowadays a crucial tool to fundamentally understand, analyze, design existing and intensified reactors [1]. They rely on a first-principles-based approach, consisting of solving the characteristic governing equations at each scale. In this view, Computational Fluid Dynamics (CFD), able to properly account for the transport phenomena at the reactor and catalyst scales, is coupled with the accurate description of the chemical kinetics granted by microkinetic models or kinetic Monte Carlo (kMC) simulations [2]. However, the CFD of reactive flows is challenging and computationally expensive since the detailed description of homogeneous and heterogeneous chemistry involves the solution of stiff numerical problems along with complex meshes characterized by tens/hundreds of millions of cells. Therefore, reactive CFD simulations are usually carried out in High Performing Computing systems by geometrically dividing the computational mesh into subdomains assigned to different processors. Two issues arise and require additional care when parallel computing is employed in reactive simulations [3]. First, the distribution of the reactive cells in the computation domain can be strongly inhomogeneous due to the geometry of the system (e.g. packed bed reactor). Second, the cost of chemical calculation is highly related to the local thermochemical state (e.g., partial pressure, temperature, coverages). As a result, the computational load is highly variable across the subdomains often resulting in imbalanced conditions which leads to highly inefficient calculations.
To tackle this problem, this work aims at developing a Dynamic Load Balancing (DLB) approach, specifically conceived for detailed CFD simulations of reactive flows at catalyst surfaces [4], which on-the-fly redistributes the calculations among the processors obtaining an even computational load and a concomitant higher efficiency. This results in an optimized distribution of the computational burden associated with the chemistry solution across processors, thereby minimizing overheads. 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. Methods

Our approach employs DLB combined with a hybrid parallelization technique, integrating both MPI and OpenMP protocols and it has been implemented in the catalyticFoam framework [2]. During each time step, the DLB receives the local thermochemical state and computational costs of the chemical calculations, represented by the solution of an ODE system in the previous step. By using these costs as an indicator of the computational load of the processors, it assigns an ensemble of ODE systems to each processor taking advantage of the MPI protocol for data exchange. Therefore, at the beginning of the integrations, each processor holds a certain number of ODEs which can be different from its number of computational cells. Then, in each processor, the solution of the ODEs is further subdivided among an ensemble of threads according to a shared memory parallelization approach based on the OpenMP

protocol. In doing so, it is possible to consider the optimal number of processors required by the transport operators, managed by MPI, and to add performance boost by further subdividing the expensive ODE solution into multiple threads.

3. Results and discussion

The computational performances of the combined DLB and hybrid parallelization approach have been investigated in fixed and fluidized bed reactor simulations [4]. To exemplify the effectiveness of the approach, we show the simulation of an industrial-scale fluidized bed reactor for Oxidative Methane Coupling (Figure 1a). The system has been initially simulated using a pure MPI parallelization, with and without the DLB procedure. When the simulation is carried out without the DLB (dashed lines in Figure 1b), the computational cost of the different processors is severely imbalanced. The cost of the most overloaded processors is more than eight times larger than the one of the most underloaded revealing a poor computational efficiency. This is due to the very different thermochemical states in the two processors which leads to a severe imbalance penalizing the efficiency of the calculation. Figure 1b shows the same trends when the DLB is employed. The DLB redistributes the computational burden among the processors leading to even computational cost among the processors (solid lines in Figure 1b). This is achieved by exchanging from 4 to 14% of the cells during the simulation between the processors to improve the computational efficiency of the multiscale CFD simulations. Indeed, a strong scalability analysis (Figure 1c) has been carried out revealing that, by increasing the number of CPUs, the speed-up monotonically increases as expected with and without DLB. However, the simulation without DLB (blue line) is strongly hindered by overheads related to the imbalanced chemical calculations, while the combined DLB and hybrid parallelization grants a higher speed-up (green line). 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. Overall, this results in a 2.1-fold reduction of the simulation time as shown in Figure 1d.

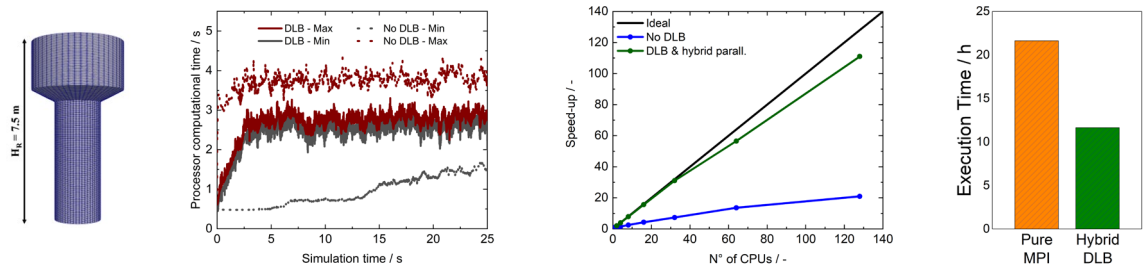

(a) (b) (c) (d)

Figure 1 (a) Computational domain; (b) Computational cost of the chemical solution with and without DLB with 128 CPUs; (c) Strong scalability analysis; (d) Overall computational cost without DLB and with hybrid DLB approach

4. Conclusions

This work proposes an effective strategy to improve the parallelization efficiency of multiscale CFD frameworks with a concomitant reduction of the computational cost. This paves the way to fundamentally investigate and design catalytic reactors and processes of industrially relevant scales by fully exploiting the potential of HPC facilities.

References

[1] M. Bracconi, Chem. Eng. Process. - Process Intensif. 181 (2022) 109148.
[2] D. Micale et al., Chemie-Ingenieur-Technik. 94 (2022) 634–651.
[3] J. Muela et al., Comput. Fluids. 190 (2019) 308–321.
[4] D. Micale et al., ACS Eng Au 2024, in press. doi:10.1021/acsengineeringau.3c00066

Keywords
Multiscale model, Reactive CFD simulations, Dynamic Load Balancing, Reactor simulations