Many industrial applications have stringent requirements regarding mixing, thermal inertia and transfer quality between multiple phases. In this context, dense gas-particle flows have exceptional characteristics and are often applied for the development of innovative processes for energy transition: biomass gasification plants, chemical loop combustion, fluidized bed solar receivers, hydrogen combustion and heat storage on particles.

To design, improve, or better understand the behavior of these complex industrial processes, numerical simulations are the perfect tool combining local-scale modeling with industrial-scale geometries. However, for optimal accuracy, the use of refined meshes and High Performance Computing is mandatory.

Neptune_cfd is a multiphase flow simulation software, based on a multi-fluid (Eulerian) approach, developed in the framework of the NEPTUNE project, financially supported by CEA, EDF, IRSN and Framatome. It builds upon the open-source software code_saturne for its HPC capabilities.

In 2015, neptune_cfd performances had been assessed on meshes of 108 million cells but this proved to be insufficient for some industrial applications in terms of spatial resolution. For these reasons, from 2018 to 2020, we participated in so-called *Grand-Challenges* to push neptune_cfd to its limits while putting at stake National and European pre-exascale supercomputers.

In 2018, we simulated an industrial fluidized-bed reactor with a 1 billion cells mesh on the supercomputer Olympe from CALMIP regional mesocenter and on EDF supercomputer Gaïa. This study was pursued in 2019 at IDRIS with an 8 billion cells mesh on up to all Jean-Zay 60 000 CPU cores. This actually was the limit for manageable production simulations. In 2020, we reproduced the 8 billion cells simulation and stepped-up to 64 billion cells on the Prace CEA-TGCC Joliot-Curie supercomputer to obtain highly detailed profiling data and to measure neptune_cfd speed-up up to 300 000 cores.

Over these studies, we assessed the sensitivity of neptune_cfd computations to parallel I/O, to node depopulation and process CPU pinning, to network topology between nodes, and to processor architecture, while measuring the code performance with highly detailed profilings. From a HPC standpoint, this data serves to further improve neptune_cfd scaling whose efficiency plateaus at 85% on 60 000 cores. Neptune_cfd has proved to be a severe but useful benchmark tool for new supercomputers, which detects and help to overcome unnoticed erratic hardware behaviors.

The performed studies were also important for scientific advances in the modeling of industrial particle-laden flows: such refined simulations yield a database of accurate results from which one can develop so-called “subgrid models”. These allow to preserve a refined physics resolution even on coarser meshes: a single highly-resolved simulation allows for many accurate and less expensive future simulations.

This series of challenges underlines the strong complementarity of Tiers-2 regional computing center with national Tiers-1 and Prace European Tiers-0 systems to validate and improve performances and stability of supercomputers and simulation software at large-scales. Over these 4 years, we increased the size of possible simulations by a factor of 500, from meshes of 108 million to 64 billion cells, thanks to a close collaboration between HPC center support teams, software developers and expert users.